

Runtime monitoring of contract regulated web services ^{*}

Alessio Lomuscio¹, Wojciech Penczek^{2,3}, Monika Solanki¹ and Maciej Szreter²

¹ Department of Computing, Imperial College London, UK

² Institute of Computer Science, PAS, Poland

³ University of Podlasie, Poland

Abstract. We investigate the problem of locally monitoring contract regulated behaviours in web services. We encode contract clauses in service specifications by using extended timed automata. We propose a *non intrusive* local monitoring framework along with an API to monitor the fulfilment (or violation) of contractual obligations. We illustrate our methodology by monitoring a service composition scenario from the vehicle repair domain, and report on the experimental results.

1 Introduction

Web services (WS) are now considered one of the key technologies for building new generations of digital business systems. Industrial strength distributed applications can be built across organisational boundaries using services as basic building blocks. When services are combined, a significant challenge is to regulate the business interactions between them. Service level agreements (SLAs) provide a useful mechanism to establish agreed levels of service provision when interactions are invoked within certain parameters. Although SLAs are useful, they can represent only basic agreements of service provision. Applications running complex, human-like activities require more general and sophisticated declarative specifications certifying legal-like agreements among the parties. In an environment where previously unknown services are dynamically discovered and binded, their composition is required to be underpinned by binding agreements.

A useful concept from the legal domain in this sense is the one of *contract* as found in human societies. Should a contract be broken by one of the parties, additional rights and/or obligations (e.g., penalties to be paid) may be applicable to some party. Contracts may not simply prescribe certain actions depending on certain states, but may go as far as to specify timing constraints (e.g, deadlines), or more sophisticated measures (such as the number of actions per temporal interval, as in some QoS agreements).

In this paper, we study the problem of monitoring runtime behaviours of *contract regulated web services*. While contracts are usually negotiated offline, it is of interest to monitor at runtime whether interactions between WS are

^{*} Partly supported under the grant „Nowe technologie informacyjne dla elektronicznej gospodarki społeczeństwa informacyjnego oparte na paradygmacie SOA”.

complying to the contracts stipulated between the parties. Runtime monitoring of web services (abbreviated RMCS) is concerned with the actual, rather than possible state transitions occurring in the system. A runtime monitor continuously checks the executions against given set of specification properties. In the case of contract-based web services we are interested in monitoring at runtime whether the contracts the web services are supposed to adhere to are violated in a given run of the system, and if so, whether some recovery action is performed.

Monitoring complex interactions such as the ones above is non-trivial. The key issue is the one of scalability. It is relatively easy to envisage a methodology whereby contracts and possible behaviours are explicitly stored in memory and the stream of events at local WS level is matched at runtime against the envisaged contract-compliant runs. However with many complex contracts to be verified and several WS present in a system the approach is unlikely to be effective in any scenario where the range of possible behaviours is large.

In this paper we put forward a “symbolic” solution to the problem above. We represent both all possible behaviours and the contractually-correct ones as an appropriate timed automata [1] at local web-service level. Specifically we present a local contract runtime monitor (CRM) based on the symbolic toolkit Verics [5], a symbolic model checker for timed-automata. CRM checks the local service’s execution at runtime against the symbolic representations provided, and reports back to the service (or directly to the engineer) any mismatch, or *violation*, between the contract-compliant behaviours originally prescribed and the ones actually received in the input stream. Note that differently from other lines of research we do not wish to monitor the overall service composition here. Instead we focus on a single service and aim to monitor continuously its executions, i.e., the change of its local variables and actions. This is of relevance to several application areas where individual enterprises wish to monitor whether any of their executions violates existing contracts, or service level agreements.

The significant advantage of the approach is that we do not need to keep the whole state space of the possible and the contract-compliant behaviours in memory but we can simply call the timed-automata engine at runtime to match moves against the stream of events coming from the input. Because of the requirements of the setting the approach extends conventional timed automata with additional constraints to allow the specification of compliance and violation of contracts. As discussed below the memory footprint of the CRM is also very attractive as is its performance. Additionally timed-automata offer us a natural formalism to work with any timing properties of interest.

The paper is structured as follows: in Section 2 we briefly introduce the formalism of timed automata as used here. Section 3 presents our monitoring framework. We analyse a motivating case study in 4 and discuss the monitoring results. Section 5 presents related work and conclusions.

2 Monitoring via Timed Automata

Let \mathbb{N} denote the set of naturals (including 0), \mathbb{Z} - the set of integers, \mathbb{Q} - the set of rational numbers, \mathbb{R} (\mathbb{R}_+) - the set of (non-negative) reals, and V be a finite set of integer variables. By a (*variable*) *valuation* we mean any total mapping $\mathbf{v} : V \rightarrow \mathbb{N}$. We extend the mapping \mathbf{v} to expressions of $Ex(V)$ in the usual way. The satisfaction relation (\models) for the boolean expressions is also standard. Given a variable valuation \mathbf{v} and an instruction $\alpha \in Ins^L(V)$, we denote by $\mathbf{v}(\alpha)$ the valuation \mathbf{v}' , obtained after executing α at \mathbf{v} , which is defined as follows:

- if $\alpha = \epsilon$ then $\mathbf{v}' = \mathbf{v}$,
- if $\alpha = (v := ex)$, then $\mathbf{v}'(v) = \mathbf{v}(ex)$ and $\mathbf{v}'(v') = \mathbf{v}(v')$ for all $v' \in V \setminus \{v\}$,
- if $\alpha = \alpha_1 \alpha_2$, then $\mathbf{v}' = (\mathbf{v}(\alpha_1))(\alpha_2)$.

Let $\mathcal{X} = \{x_1, \dots, x_{n_x}\}$ be a finite set of real-valued variables, called *clocks*. The set of *clock constraints* over \mathcal{X} and V , denoted $\mathcal{C}(\mathcal{X}, V)$, is defined by the grammar: $\mathbf{cc} ::= true \mid x_i \sim c \mid x_i \otimes x_j \sim c \mid x_i \otimes x_j \sim v \mid x_i \otimes v \sim c \mid v \otimes w \sim x_i \mid \mathbf{cc} \wedge \mathbf{cc}$, where $x_i, x_j \in \mathcal{X}$, $v, w \in V$, $c \in \mathbb{N}$, $\otimes \in \{+, -\}$, and $\sim \in \{\leq, <, =, >, \geq\}$. Let \mathcal{X}^+ denote the set $\mathcal{X} \cup \{x_0\}$, where $x_0 \notin \mathcal{X}$ is a fictitious clock representing the constant 0. A *clock-to-clock assignment* A over \mathcal{X} is a function $A : \mathcal{X} \rightarrow \mathcal{X}^+$. $Asg(\mathcal{X})$ denotes the set of all the assignments over \mathcal{X} . By a *clock valuation* we mean a mapping $\mathbf{c} : \mathcal{X} \rightarrow \mathbb{R}_+$. The satisfaction relation (\models) for a clock constraint $\mathbf{cc} \in \mathcal{C}(\mathcal{X}, V)$ under a clock valuation \mathbf{c} and a variable valuation \mathbf{v} is defined as $(\mathbf{c}, \mathbf{v}) \models (x_i \otimes v \sim c)$ iff $\mathbf{c}(x_i) \otimes \mathbf{v}(v) \sim c$, and similarly for the other cases.

In what follows, the set of all the pairs (\mathbf{c}, \mathbf{v}) , composed of a clock and a variable valuation, satisfying a clock constraint \mathbf{cc} is denoted by $\llbracket \mathbf{cc} \rrbracket$. Given a clock valuation \mathbf{c} and $\delta \in \mathbb{R}_+$, by $\mathbf{c} + \delta$ we denote the clock valuation \mathbf{c}' such that $\mathbf{c}'(x) = \mathbf{c}(x) + \delta$ for all $x \in \mathcal{X}$. Moreover, for a clock valuation \mathbf{c} and an assignment $A \in Asg(\mathcal{X})$, by $\mathbf{c}(A)$ we denote the clock valuation \mathbf{c}' such that for all $x \in \mathcal{X}$ we have $\mathbf{c}'(x) = \mathbf{c}(A(x))$ if $A(x) \in \mathcal{X}$, and $\mathbf{c}'(x) = 0$ if $A(x) = x_0$. Finally, by \mathbf{c}^0 we denote the *initial* clock valuation, i.e., the valuation such that $\mathbf{c}^0(x) = 0$ for all $x \in \mathcal{X}$. In this paper we assume a slightly modified definition of timed automata with discrete data [17], which extend the standard timed automata of Alur and Dill in the following way:

Definition 1. A timed automaton with discrete data (*TADD*) is a tuple $\mathcal{A} = (\Sigma, L, l^0, V, \mathcal{X}, \mathcal{E}, \mathcal{I})$, where

- Σ is a finite set of labels (actions),
- L is a finite set of locations,
- $l^0 \in L$ is the initial location,
- V is the finite set of integer variables,
- \mathcal{X} is the finite set of clocks,
- $\mathcal{E} \subseteq L \times \Sigma \times Bool(V) \times \mathcal{C}(\mathcal{X}, V) \times Ins^L(V) \times Asg(\mathcal{X}) \times L$ is a transition relation, and
- $\mathcal{I} : L \rightarrow \mathcal{C}(\mathcal{X}, \emptyset)$ is an invariant function.

The invariant function assigns to each location a clock constraint (without integer variables⁴) expressing the condition under which \mathcal{A} can stay in this location.

The semantics of a TADD \mathcal{A} is given below.

Definition 2. *The semantics of $\mathcal{A} = (\Sigma, L, l^0, V, \mathcal{X}, \mathcal{E}, \mathcal{I})$ for an initial valuation $\mathbf{v}^0 : V \rightarrow \mathbb{Z}$ is a labelled transition system $\mathcal{S}(\mathcal{A}) = (Q, q^0, \Sigma_S, \rightarrow)$, where:*

- $Q = \{(l, \mathbf{v}, \mathbf{c}) \mid l \in L \wedge \mathbf{v} \in \mathbb{Z}^{|V|} \wedge \mathbf{c} \in \mathbb{R}_+^{|\mathcal{X}|} \wedge \mathbf{c} \models \mathcal{I}(l)\}$ is the set of states,
- $q^0 = (l^0, \mathbf{v}^0, \mathbf{c}^0)$ is the initial state,
- $\Sigma_S = \Sigma \cup \mathbb{R}_+$ is the set of labels,
- $\rightarrow \subseteq Q \times \Sigma_S \times Q$ is the smallest transition relation:
 - for $a \in \Sigma$,
 $(l, \mathbf{v}, \mathbf{c}) \xrightarrow{a} (l', \mathbf{v}', \mathbf{c}')$ iff there exists a transition $t = (l, a, \beta, \mathbf{cc}, \alpha, A, l') \in \mathcal{E}$ such that $\mathbf{v} \models \beta$, $(\mathbf{c}, \mathbf{v}) \models \mathbf{cc}$, $\mathbf{v}' = \mathbf{v}(\alpha)$, $\mathbf{c} \models \mathcal{I}(l)$, and $\mathbf{c}' = \mathbf{c}(A) \models \mathcal{I}(l')$ (action transition),
 - for $\delta \in \mathbb{R}_+$,
 $(l, \mathbf{v}, \mathbf{c}) \xrightarrow{\delta} (l, \mathbf{v}, \mathbf{c} + \delta)$ iff $\mathbf{c} \models \mathcal{I}(l)$ and $\mathbf{c} + \delta \models \mathcal{I}(l)$ (time transition).

Intuitively, in the initial state all the variables are set to their initial values, and all the clocks are set to zero. Then, at a state $q = (l, \mathbf{v}, \mathbf{c})$ the system can either execute an action or time transition.

2.1 TADD Semantics for RMCS

Inspired by related work in the formal representation of states of compliance and violation [10], we partition the set of global states Q of $\mathcal{S}(\mathcal{A})$ for $\mathcal{A} = (\Sigma, L, l^0, V, \mathcal{X}, \mathcal{E}, \mathcal{I})$ into two subsets G and R such that $G \cap R = \emptyset$ ⁵. The set G represents *green* (or *ideal*) states, whereas R represents the *red* (or *non-ideal*) ones. Intuitively, G contains the states of compliance and R contains the states of violation with respect to the contract, i.e., the whole set of clauses being included. Figure 1 illustrates the intuition behind the semantics.

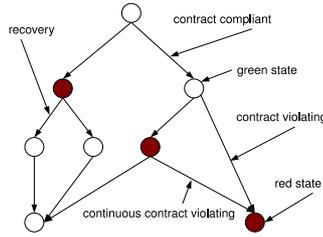


Fig. 1. Partitioning of states and transitions in a TADD

⁴ To ensure the monotonicity of the timed successor relation.

⁵ This partition is obtained “location-wise” from a partition of the set of locations L .

Based on the above partitioning each action transition (q, a, q') of $\mathcal{S}(\mathcal{A})$ can be one of the following four types of transitions:

- **Contract compliant:** between green and green states, i.e., $q, q' \in G$. These transitions occur when the observed behaviour is in compliance with the prescribed behaviour of the contract.
- **Contract violating:** between green and red states, i.e., $q \in G$ and $q' \in R$. These transitions occur when the observed behaviour violates the prescribed behaviour of the contract.
- **Recovery:** between red and green states, i.e., $q \in R$ and $q' \in G$. These transitions occur when a recovery action is taken by the service after a violation of the prescribed behaviour is recorded.
- **Continuous contract violating:** between red and red states, i.e., $q, q' \in R$. The transitions occur when no recovery results from a previous violation.

We say that there is a *step* from state q_1 to q_2 in \mathcal{A} if $q_1 \xrightarrow{\delta_1} q'_1 \xrightarrow{a} q'_2 \xrightarrow{\delta_2} q_2$, for some states $q'_1, q'_2 \in Q$, $\delta_1, \delta_2 \in \mathbb{R}_+$, and $a \in \Sigma$.

2.2 Querying of TADD for RMCS

In our approach described below, we rephrase the problem of local monitoring of executions against contract compliant behaviours into the following model checking problem: For a given TADD \mathcal{A} and a pair (Q_1, Q_2) of sets of global states of $\mathcal{S}(\mathcal{A})$, we check whether there are two states $q_1 \in Q_1$ and $q_2 \in Q_2$ such that there is a step from q_1 to q_2 . If so, we denote the step as $Q_1 \rightsquigarrow Q_2$. Below we explain the methodology for monitoring contracts represented by a TADD:

- The transition relation is first encoded into a propositional formula. Then for each step, this propositional formula is conjuncted with the encodings of a pair of sets of states (Q_1, Q_2) given as an input.
- We start by making a query about a step from the source set Q_1 to the set of the red states $(Q_1 \rightsquigarrow R)$. To this aim the input (Q_1, R) is encoded. If the resulting formula is satisfiable, then “non compliance” is reported.
- If the resulting formula is not satisfiable, then the input (Q_1, Q_2) is encoded. Again, satisfiability of the resulting formula is tested. Depending on the result, either “compliance” or “invalid transition” is reported

Our tool uses MiniSAT [6] for checking satisfiability, but any standard SAT-solver capable of processing propositional formulas in the conjunctive normal form can be applied.

3 Runtime monitoring framework

Our architecture for local monitoring, RMCS, is illustrated in Figure 2. Agents implementing WS are the primary entities within our framework. Service behaviour and contracts associated with them may be specified at a high level using

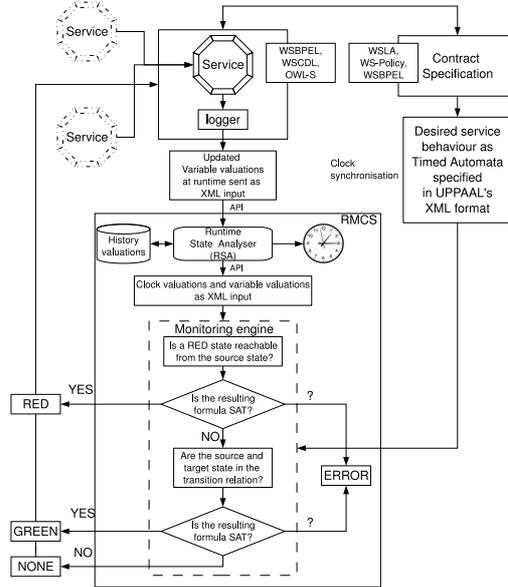


Fig. 2. The general architecture and methodology

WS standards, e.g., WSBPEL [13] and contracts, e.g., WSLA [7]. The TADD specification for the service is engineered from these interface representations. A significant feature of our framework is that we do not place any restriction on service implementation in terms of development infrastructure and execution platforms. Central to our framework is a *non intrusive* approach to monitoring. The mechanism works independently of service execution. The module responsible for linking the service to the monitoring mechanism is the “logging framework”. Each service to be monitored is associated with a logger. The logger records a “snapshot” of the variables of interest that are to be monitored. Snapshots may be finely grained (i.e., every change in valuation is recorded), or coarse (i.e, recorded after every pre-specified or random number of changes). Snapshots may also be time bound, i.e., taken after a specific time interval. Each snapshot captures variable valuation as they are generated, updated by the service or received from partners. Every snapshot is passed to the runtime state analyser using a dedicated API provided by the logging framework.

TADDs for services: The specification of service behaviour used by RMCS is the TADD representation described in Section 2. We use the XML format generated by the model checker UPPAAL for representing the TADD. Our choice is motivated by the fact that UPPAAL provides a user friendly GUI. This is of great help to system engineers when modelling the TADDs. Secondly, the XML representation format can be modified easily in order to take into account any extensions to the TADD model. As illustrated in Figure 3, the TADD specification encodes all possible desired behaviours for a service. Typically, the full set of behaviours for a contract regulated service can be derived from:

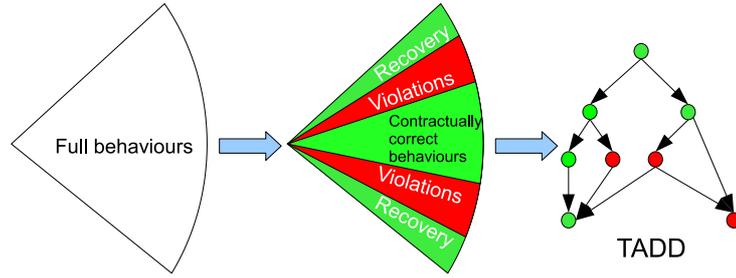


Fig. 3. Set of behaviours for a service

- its contractually compliant behaviours. These behaviours encapsulate contractual obligations for the service.
- behaviours that are classified as violations of the contract.
- behaviours that define a recovery from incurred violations.

There is a one-to-one correspondence between variables defined in the TADD and the service implementation in terms of types and names i.e., variables names and their types across the two representations are kept identical for simplicity. The logging framework passes execution snapshots to the RSA as an XML data structure.

Runtime State Analyser (RSA): The runtime state analyser interfaces with the logger for receiving snapshots of the latest variable valuations generated by the service. Snapshots are passed to the RSA via the logging framework. RSA is also responsible for updating clocks by querying the system hardware, in accordance with the granularity of a *tick* chosen by the service. A tick can be defined in terms of seconds, minutes, hours or days i.e., clock values may be captured every second, minute or day or any other interval chosen by the service. Clocks may also be updated based on resets and assignments defined in the TADD for the service. The monitoring engine reports back any resets or assignments made to the clocks, along with reporting the results, e.g., if the monitoring engine reports that a clock x is reset, the current valuation of the clock is discarded and the clock is re-initialised. Clock resets and updates are significant especially when a recovery action is taken against a violation of contract. In such scenarios one would like to start the monitoring again with clock valuation recorded before the violation occurred. The clock valuations once recorded are then added to the variable valuation snapshot received from the logger. RSA is also responsible for storing the history of service executions and passing the augmented snapshots to the monitoring engine.

The runtime information passed to the monitoring engine from the RSA consists of one or several steps. A step is a pair of consecutive snapshots, represented as “source” and “target” states. The states define immediately previous (source) and current (target) clock and variable valuations recorded for the service. An example step for the case study in section 4 for the component

RepairCompany, clock x and variable *maxRepairRequestTime* can be informally shown as: $(x = 3, \text{maxRepairRequestTime} = 7) \rightarrow (x = 5)$.

Any component of a source or a target such as a clock valuation, or a variable valuation can be omitted. Thus each set can range from containing only the system states (no values given at all) to a single state (every component is specified).

3.1 The monitoring engine

The monitoring engine is the core component responsible for testing the conformance of runtime service behaviour presented as an input from the RSA, against the prescribed TADD specification of the service. Each execution step passed to the engine is encoded and its conformance to the TADD specification is tested by means of the model checking approach described in Section 2.2. Our SAT-based verification method does not need to construct the complete model for \mathcal{A} , which could be unfeasible for both the explicit-state as well as BDD-based methods. Instead, the timed automaton is encoded as a propositional formula, but testing of its satisfiability is postponed until the concrete source and target states of an execution step are provided. This significantly reduces the computational cost as information about concrete states prunes the state space to be searched. The engine monitors if the service has taken an execution step from the source set to the target set of states in accordance with its prescribed TADD. In addition, it checks if it is possible to reach a target red state from a given set of source states. In the general case the system consists of several components represented by automata; if at least one component of a location reachable as a result of the transition is red, then this fact is reported.

Monitoring results: The engine checks at runtime whether the stream of execution steps received as inputs from the RSA, conforms with its symbolic representation of all possible behaviours. For each execution step, the answer returned by the monitoring engine is one of the following facts:

- **GREEN** - the step is conforming with the specification, i.e., there is a contract compliant transition between the source and target states.
- **RED** - a red state is reached as a target of the transition given, i.e., a contract has been violated as a result of the transition. This also signifies the fact that the inputs do not comply with the extended format of the TADD for the service.
- **NONE** - the step is not conforming with the specification, i.e., there is no such transition, neither contract compliant or otherwise.
- **ERROR** - the specification given does not mirror the observed transition so it amounts to an error.

Results reported at runtime may be analysed in several ways. In case of contract compliant transitions, the service can continue executing as per the orchestrated workflow. For contract violating transitions, the service administrator may impose on the service to execute one of the prescribed recovery transition. In other

cases the administrator may choose to override the violations reported and allow the service to carry on the execution. For a continuous contract violating transition being reported, the service may be stopped. Finally, the outputs generated may be stored in a log file for future offline analysis.

4 A vehicle repair contract: case study

We now present a description of a case study followed by a detailed discussion on the local monitoring and analysis of one of the agents in the composition.

We consider a service composition scenario that defines a repair contract between a client (C) and a vehicle repair company (RC). A repair contract specifies details concerning a particular repair, i.e., the type of repair to be performed, price, dates, pickup and delivery locations etc. For simplicity we only model the behaviour of RC . Table 1 identifies some of the contract clauses governing the actions taken by RC , the deadlines against which the contracts are monitored, if the clause can be violated, and, if a violation is recorded, whether any recovery is possible. Note that in some cases RC may take an “offline” action, in response to a violation from which no recovery may be possible. For example consider clause 6: “For any violation take recovery action within $maxRecoveryTime$ - number of days”. If the recovery action is not taken, C may take an offline legal action against RC .

The informal behaviour of RC is described as follows. When RC receives a request from C to undertake a repair job, it sends a repair proposal. In response, C sends an acceptance or rejection message. If accepted, RC sends a contract initiation message to C . RC then waits for the vehicle to arrive, failing which it sends two reminders to C . If the vehicle fails to arrive, it takes an offline action. As per the contract, RC is *obliged* to assess the damage, repair the vehicle and send a report to C . On receiving the report, C is *obliged* to send payment to RC . If the payment is not sent, RC sends two reminders to C and then takes an offline action.

The actions taken by RC in response to messages sent by C are monitored to meet the deadlines set for various activities as per the contract. Failure to meet deadlines is considered a violation of the contractual obligations. In some cases a recovery from the violation may be possible.

4.1 Monitoring the runtime behaviour of the Repair Company

The full set of behaviours of the RC is represented by a TADD⁶. As described in Section 4, deadlines for various activities are decided during contract negotiation between the parties. Deadlines are defined in terms of number of days. For example consider a contract clause to be monitored: *If C sends a damaged vehicle to RC, RC assesses the damage to the vehicle within 3 days -clause (3) in table 1. A snippet of the TADD for the clause is shown in the Figure 4.* Figure

⁶ The complete TADD for the example is too large to be shown here.

clause	Contract regulated actions	Deadline	Violation	Recovery
1	Receives a repair request by C	5 days	-	-
2	Sends a repair proposal to C	7 days	-	-
3	Assess damage to the vehicle	3 days	yes	yes
4	Execute repair	30 days	yes	yes
5	Send repair report to C	5 days	yes	yes
6	For any violation take recovery action	3 days	yes	no (take offline action)

Table 1. Some contract regulated actions for RC

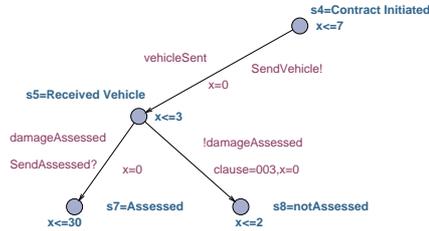


Fig. 4. TA specification of clause (3)

4 describes the timeline in number of days for clause (3), a snapshot passed to RSA at $x = 0$ from the logger when a vehicle for repair arrives, snapshots sent to the monitoring engine by the RSA and the results from monitoring. As per the contract, once a damaged vehicle has arrived the damage has to be assessed within 3 days. A snapshot is again sent by the logger to the RSA at $x = 5$. The snapshots taken at $x = 0$ and at $x = 5$ are sent by the RSA as a pair - or as a “step” to RMCS. The results returned by the monitoring engine are $\{RED, reset, 003\}$. *RED* signifies that a violation has occurred, i.e., the damage was not assessed within the deadline, *reset* indicates that the clock has been reset and 003 indicates the index of the clause that has been violated.

4.2 Experimental results and Discussion

In order to validate our methodology, we implemented the above case study and monitored several runtime execution steps for the service. To provide an indication of how many variables the toolkit can monitor at the same time we scaled the example described above parametrizing the number of cars in the contract. As one clock and one integer variable are associated with every car, numbers of clocks and int variables grow respectively. Notice that the bigger values can these vars have, the more bits are needed for encoding them.

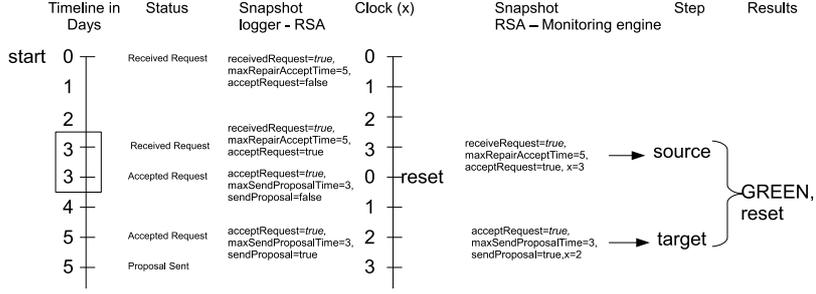


Fig. 5. Runtime valuations for clause (2)

The client is interested in getting x cars repaired. The request for all these repairs is included as a single contract.

Step	state	Explanation
step 1	source	RC waits to receives the request for repairing cars.
	target	RC receives the request for repairing x cars. We show here an example of the clock and variable valuations for three cars.
step 3	source	RC accepts the request for repairing x cars.
	target	RC sends repair proposals for repairing x cars.

Table 2. Explanation of trace contents for steps 1 and 3

step nr	nr of cars	nr of int variables	nr of clocks	$N_c/Nvars$	time [s]	answer
1	10	10	10	6779/16528	<1	YES
	20	20	20	17738/43455	<1	
	300	300	300	265741/652852	4.3	
3	10	10	10	6743/16431	<1	NO
	30	30	30	26781/65822	<1	
	300	300	300	265811/653052	5.4	

Table 3. The experimental results. Parameters of the example are described in the text; size of encoding: $N_c/Nvars$ is the number of clauses/Boolean variables in the result CNF formula; time refers to checking this formula using the tool Minisat.

Table 2 explains the contents of traces for contract clauses 1 and 3 (see Table 1). Table 3 presents experimental results. It can be stated that the approach performs extremely well against explicit approaches, which, although more immediate in their construction, typically fail to scale due to their memory footprint. This phenomenon could be even more visible if we could have a network of automata instead of a single automaton defining a contract. The experiments show the approach can monitor effectively several hundreds of variables. This

is sufficient for very complex monitoring of key aspects of a service. We did not optimise the monitoring process in any way; we expect our results to improve significantly by tailoring the approach to a particular problem we wish to monitor.

We found the only timeconsuming step of our methodology to be the construction of the automaton representing all behaviours. However this only needs to be done once, tools to assist the user in the design exist, and it can then be used for all monitoring purposes. Additionally it is to be noted that for complex applications, a representation of the service composition in an automata-based framework (or something equivalent) is expected to be produced during the design phase, so the construction above may in practice be derivable from existing formalisations of the composition under analysis.

5 Related work and conclusions

In this paper we presented a symbolic approach based on timed automata for the runtime monitoring of contract regulated agent based WS. Monitoring service behaviour has been an active area of research. Several efforts have investigated various formalisms and frameworks for the monitoring of functional and non-functional properties of services. The monitoring problem has been considered for several formalisms in papers [16, 2, 4, 14, 12, 11, 9, 3]. Timed automata have been used in earlier work such as [8] on monitoring and fault diagnosis of systems, while [15] presents an approach which also uses timed automata for monitoring SLAs. The aims of the above approaches are however quite different from our objectives in this paper. However [8, 15] are not concerned with local monitoring of contract-based executions.

An attractive feature of our approach over those mentioned above is that histories and pending contracts are not stored in memory during the monitoring. This positively impacts the scalability of the approach and is particularly useful when monitoring multiple and long running contracts between several services. As a case study we presented the monitoring of contracts for a repair company. Although the TADD for the service is not large enough to exploit the full capabilities of RMCS, we believe it is still sufficiently significant to demonstrate the methodology and scope of the proposed approach. Experiments demonstrate larger scenarios would be handled just as well by the technique.

While verification is still an aspect of systems validation we are not aware of symbolic attempts to the runtime monitoring of these notions. It seems to us that it may be of interest to investigate whether this could be achieved in ways related to the technique presented here.

Much work remains to be done. An important part of our future work is the translation to TADDs from high level specification standards such as WSBPEL. Developing such a translation is non trivial as most standards do not support the explicit representation of timing constraints on prescribed activities. These standards therefore need to be augmented with such support. Additionally, we

are interested in developing an interactive compiler for services specified in WS-BPEL to be compiled into our TADD representation.

References

1. R. Alur. Timed Automata. In *Proceedings of the 11th International Conference on Computer Aided Verification (CAV'99)*, volume 1633 of *LNCS*, pages 8–22. Springer-Verlag, 1999.
2. Fabio Barbon, Paolo Traverso, Marco Pistore, and Michele Trainotti. Run-time monitoring of instances and classes of web service compositions. In *ICWS '06: Proceedings of the IEEE International Conference on Web Services*.
3. Luciano Baresi, Carlo Ghezzi, and Sam Guinea. Smart monitors for composed services. In *ICSOC '04: Proceedings of the 2nd International Conference on Service Oriented Computing*. ACM.
4. Domenico Bianculli and Carlo Ghezzi. Monitoring conversational web services. In *IW-SOSWE '07: 2nd international workshop on Service oriented software engineering*. ACM.
5. P. Dembiński, A. Janowska, P. Janowski, W. Penczek, A. Pólrola, M. Szreter, B. Woźna, and A. Zbrzezny. VerICS: A tool for verifying Timed Automata and Estelle specifications. In *Proc. of the 9th Int. Conf. on Tools and Algorithms for the Construction and Analysis of Systems (TACAS'03)*, LNCS. Springer-Verlag.
6. N. Eén and N. Sörensson. MiniSat. <http://minisat.se/MiniSat.html>.
7. Alexander Keller and Heiko Ludwig. The WSLA Framework: Specifying and Monitoring Service Level Agreements for Web Services. *J. Netw. Syst. Manage.*, 2003.
8. M. Krichen and S. Tripakis. Black-box conformance testing for real-time systems. In *11th International SPIN Workshop on Model Checking of Software (SPIN'04)*, Barcelona, Spain, LNCS.
9. Zheng Li, Yan Jin, and Jun Han. A runtime monitoring and validation framework for web service interactions. In *ASWEC '06: Proceedings of the Australian Software Engineering Conference (ASWEC'06)*. IEEE Computer Society.
10. A. Lomuscio and M. Sergot. Deontic interpreted systems. *Studia Logica*, 75(1):63–92, 2003.
11. G. Mahbub, K.; Spanoudakis. Run-time monitoring of requirements for systems composed of web-services: initial implementation and evaluation experience. In *ICWS'05, IEEE International Conference on Web Services*.
12. Carlos Molina-Jimenez, Santosh Shrivastava, Ellis Solaiman, and John Warne. Contract representation for run-time monitoring and enforcement. *cec*, 2003.
13. OASIS Web service Business Process Execution Language (WSBPEL) TC. Web service Business Process Execution Language Version 2.0, 2007.
14. Marco Pistore, F. Barbon, Piergiorgio Bertoli, D. Shaparau, and Paolo Traverso. Planning and monitoring web service composition. In *AIMSA*, pages 106–115, 2004.
15. F. Raimondi, J. Skene, L. Chen, and W. Emmerich. Efficient monitoring of web service slas. Technical report, UCL, London, 2007.
16. Monika Solanki. *A Compositional Framework for the Specification, Verification and Runtime Validation of Reactive Web Service*. PhD thesis, De Montfort University, Leicester, UK, October 2005.
17. A. Zbrzezny and A. Pólrola. SAT-based reachability checking for timed automata with discrete data. *Fundamenta Informaticae*, 79(3-4):579–593, 2007.