Enacting declarative languages using LTL : avoiding errors and improving performance

Citation for published version (APA):

Pesic, M., Bosnacki, D., & Aalst, van der, W. M. P. (2009). *Enacting declarative languages using LTL : avoiding errors and improving performance*. (Computer science reports; Vol. 0914). Technische Universiteit Eindhoven.

Document status and date:

Published: 01/01/2009

Document Version:

Publisher's PDF, also known as Version of Record (includes final page, issue and volume numbers)

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.tue.nl/taverne

Take down policy

If you believe that this document breaches copyright please contact us at:

openaccess@tue.nl

providing details and we will investigate your claim.

Enacting Declarative Languages using LTL: Avoiding Errors and Improving Performance

Maja Pešić¹, Dragan Bošnački², and Wil M.P. van der Aalst¹

¹ Department of Mathematics and Computer Science,
Eindhoven University of Technology,
P.O. Box 513, NL-5600 MB, The Netherlands
{m.pesic,w.m.p.v.d.aalst}@tue.nl

² Department of Biomedical Engineering,
Eindhoven University of Technology,
P.O. Box 513, NL-5600 MB, The Netherlands
dragan@win.tue.nl

Abstract. In our earlier work we have proposed using the declarative language DecSerFlow for modeling, analysis and enactment of web service processes. DecSerFlow uses constraints, which are formally specified as Linear Temporal Logic (LTL) formulas, to implicitly define possible executions of a model: any execution that satisfies all constraints is possible. A finite representation of all possible executions is retrieved as an automaton generated from LTL-based constraints. Standard algorithms for creating Büchi automata from LTL formulas cause errors when be applied to DecSerFlow due to important semantical differences. On the one hand, LTL handles *infinite traces* where each element of the trace can refer to *zero or more propositions*. On the other hand, executions of DecSerFlow models are *finite* sequences of *single events*. In this paper we describe how both LTL and automata generation algorithms can be adjusted to fit two properties of DecSerFlow. Since adjustments for finite traces have already been proposed by other researchers, this paper focuses on the modifications needed to handle occurrences of single events. Besides eliminating errors caused by the described mismatch, the proposed adjustments also improve the performance of the model checking algorithms dramatically. This is important as for the enactment of declarative languages like DecSerFlow new problems related to model checking techniques are generated after each step in the process.

1 Introduction

In our previous work [1] we have proposed using DecSerFlow for specification, verification, monitoring and orchestration (i.e., execution) of web services in the context of business processes. DecSerFlow uses constraints to implicitly specify in which order tasks can be executed. Constraints are rules specified on two levels. First, there is a graphical representation which is similar to other graphical process modeling approaches: tasks are represented as rectangles and constraints as special lines between tasks. Second, Linear Temporal Logic (LTL) [5] is used

Fig. 1. A DecSerFlow model

for the formal specification of the semantics of these constraints. The graphical representation of a DecSerFlow model improves the readability of models, while (often complex and unreadable) LTL expressions enable verification, monitoring and deadlock-free execution.

Figure 1 shows an illustrative example of a DecSerFlow model. This model contains tasks A , B and C , and two *response* constraints. The first constraint specifies that each successfully completed execution of task A must eventually be followed by at least one successfully completed execution of task B , which is formally defined with the LTL formula $\Box(A_c \Rightarrow \Diamond B_c)$. The second constraint specifies that each successfully completed execution of task B must eventually be followed by at least one successfully completed execution of task A , which is formally defined by the LTL formula $\Box(B_c \Rightarrow \Diamond A_c)$. Note that DecSerFlow uses the notion of events related to execution of a task: scheduling, starting, completing, canceling, delegating, etc. For example, starting, canceling and completing the execution of some task T is denoted by T_s , T_x and T_c , respectively.

The DecSerFlow language is supported by the DECLARE system, which enables modeling, verification and enactment of LTL-based models, e.g., DecSerFlow models. Note that this system cannot directly be used to enact web services. Instead it can be used for manual execution of declarative models, where the user manually executes each step in the process. DECLARE is an open source system and it can be downloaded from <http://declare.sf.net>.

Using LTL for formalization of constraints makes DecSerFlow a truly declarative process modeling language, which increases flexibility [1]. On the one hand, procedural process models lack flexibility because they explicitly specify all possible executions (i.e., orderings of tasks). On the other hand, DecSerFlow models are more flexible because all possible executions are specified implicitly, as all executions that satisfy all constraints in the model. A finite representation of all possible executions is obtained from LTL specifications of DecSerFlow constraints: by means of generating an automata for LTL formulas. Algorithms for generating automata that represent exactly all traces that satisfy an LTL formula have been developed in the field of model checking [5]. In DecSerFlow, we use automata generated from LTL specifications of constraints for model verification, execution monitoring and ensuring deadlock-free execution of models of web services [13].

Typically, LTL and model checking techniques are used to verify properties of a given model. DecSerFlow uses LTL to define possible executions, as se-

Fig. 2. LTL for DecSerFlow

quences of executed events. For example, one possible execution of the model shown in Figure 1 is executing task C three times, which is specified as $\sigma = C_s, C_c, C_s, C_c, C_s, C_c$. There are *two important differences* between the ‘standard’ LTL and the LTL applied to DecSerFlow models, as illustrated in Figure 2.

The first difference between standard LTL and DecSerFlow LTL is the length of execution traces (i.e., words). On the one hand, standard LTL considers infinite traces, as shown in Figure 2(a). On the other hand, Figure 2(b) shows that the execution of a web service process eventually terminates. Hence, the infinite semantics of standard LTL [5] cannot be applied to DecSerFlow models. In order to apply LTL to finite traces, we adopt a simple and efficient approach originally proposed by Giannakopoulou et al. [7].

The second difference between standard LTL and the DecSerFlow LTL is the semantics of elements in a trace. Standard LTL assumes that one element of the trace can refer to more than one *proposition*. For example, it is possible to monitor two properties: (P_1) the motor temperature is higher than 80 degrees and (P_2) the speed of the turbine is higher than 150 km/h. As Figure 2(a) shows, each element of the trace could then refer to: (1) none of the two properties, i.e., neither P_1 nor P_2 hold, (2) only property P_1 holds, (3) only property P_2 holds, or (4) properties P_1 and P_2 both hold. In the case of execution traces of DecSerFlow models we assume that only one property holds at one moment, i.e., each of the elements of the trace refers to exactly one event, as shown in Figure 2(b).

Due to these differences, using standard LTL and automata generation algorithm may cause errors when it comes to verification, monitoring and deadlock-free execution of DecSerFlow models. In this paper we show how the *semantics* of standard LTL and the automata generation can be adjusted for declarative languages like DecSerFlow. Besides for elimination of errors, the proposed adjustments improve the performance of the algorithm, both with respect to the size of the automata and the processing time. We will use the model shown in Figure 1 as a running example.

The remainder of this paper is organized as follows. We start by describing how LTL and automata generated from LTL are used in DecSerFlow(sections 2 and 3). In Section 4 we sketch how the finite trace semantics can be applied to LTL and generated automata, by using the approach described in [7]. Section 5 describes how the LTL semantics and the algorithm for generating automata must be changed in order to be applicable to sequences of *single events*. The results of experiments testing the performance of proposed changes are presented in Section 6. Finally, related work is discussed in Section 7 and Section 8 concludes the paper.

2 LTL and DecSerFlow

DecSerFlow uses LTL to formally specify constraints. A well-formed LTL formula can use standard logical operators ($!$, $\wedge$ and $\vee$) and several additional temporal operators: $\circ$ (next), U (until), W (weak until), V (release), $\square$ (always) and $\diamond$ (eventually).

Given a finite set of atomic propositions P , every $p \in P$ is a well-formed LTL formula. If Φ and Ψ are well-formed LTL formulas, then **true**, **false**, $!\Phi$, $\Phi \wedge \Psi$, $\Phi \vee \Psi$, $\square\Phi$, $\diamond\Phi$, $\circ\Phi$, $\Phi U \Psi$, $\Phi V \Psi$ and $\Phi W \Psi$ are also well-formed LTL formulas. An interpretation of an LTL formula is a set of infinite traces $\sigma = \sigma_1, \sigma_2, \dots$ over 2^P (sets of propositions). We write $\sigma^{i \rightarrow}$ for the suffix of σ starting at position i , i.e., $\sigma^{i \rightarrow} = \sigma_i, \sigma_{i+1}, \dots$. The semantics of LTL is defined as follows:

- $\sigma \models p$ iff $p \in \sigma_1$, for $p \in P$
- $\sigma \models \Phi \vee \Psi$ iff $(\sigma \models \Phi) \vee (\sigma \models \Psi)$
- $\sigma \models !\Phi$ iff $\sigma \not\models \Phi$
- $\sigma \models \circ\Phi$ iff $\sigma^{2 \rightarrow} \models \Phi$
- $\sigma \models \Phi \wedge \Psi$ iff $(\sigma \models \Phi) \wedge (\sigma \models \Psi)$
- $\sigma \models \Phi U \Psi$ iff $(\exists_{1 \leq i} : (\sigma^{i \rightarrow} \models \Psi \wedge (\forall_{1 \leq j < i} : \sigma^{j \rightarrow} \models \Phi)))$

Also, abbreviations are used:

- $\Phi \Rightarrow \Psi$ for $!\Phi \vee \Psi$
- $\diamond\Phi$ for **true** $U\Phi$
- $\Phi W \Psi$ for $(\Phi U \Psi) \vee (\square\Phi)$
- **true** for $\Phi \vee !\Phi$
- $\square\Phi$ for $!\diamond\Phi$
- $\Phi V \Psi$ for $!(\Phi U !\Psi)$
- **false** for $!\mathbf{true}$

DecSerFlow does not directly use LTL for constraint specification. Instead, constraints are created from constraint templates. Each template has a unique name and graphical representation and its semantics is formalized by an LTL formula. DecSerFlow has more than twenty constraint templates, which are used to create constraints in models. However, new templates can be easily added to the language. Hence, any LTL formula can be used in DecSerFlow. For the detailed description of the full list of DecSerFlow templates we refer the reader to [1]. A constraint in a DecSerFlow model inherits the name, graphical representation and LTL formula from its template such that the template's formal parameters are replaced by real tasks from the model. For example, parameters X and Y from the *response* template are replaced by tasks A and B for one, and tasks B and A for the other *response* constraint in the sDecSerFlow model shown in Figure 1.

Note that in the DECLARE tool it is possible to edit the language by adding, modifying and removing templates. Moreover, DECLARE supports multiple LTL-based languages (DecSerFlow, CigDec, ConDec, etc) [13].

3 Automata and DecSerFlow

A widely exploited property of LTL is the fact that for every LTL formula a Büchi automaton can be generated, such that the language (i.e., all accepted traces) of this automaton exactly represents all traces that satisfy the formula. In the field of model-checking, various algorithms are proposed for generating automata from LTL formulas. Some examples of these algorithms can be found in [5, 8, 3, 4]. These algorithms are based on expanding a graph node into a set of nodes,

which will eventually become states of the automaton [5]. Each node has several fields [5]. Field *ID* is a unique identification for the node; Field *INCOMING* contains the set of nodes that lead to this node; Field *NEW* contains the set of formulas that must hold in the current node but have not yet been processed, i.e., this node must make these formulas true; Field *OLD* contains the set of formulas that have already been processed; Field *NEXT* contains the set of formulas that must hold at all immediate successors of this node.

In this paper we will not describe the algorithm in detail. For detailed descriptions of available algorithms we refer the interested reader to the existing literature, e.g., [5, 8]. Various algorithms for creating a Büchi automaton from an LTL formula are built on the same idea [5], but apply various optimization features. In the remainder of this paper we will use the algorithm presented in [8] as the representative of these algorithms, and we will refer to this algorithm as to the **BASIC** algorithm. The **BASIC** algorithm consists of several basic steps:

1. The original LTL formula is rewritten to a normal form.
2. A graph of nodes is created by the recursive method *expand*, which is briefly sketched in Algorithm 3.1. Formulas from the field *NEW* are processed one by one (line 14), by breaking down the formulas to the level of propositions. Formulas aUb , aVb and $a \vee b$ are broken down by creating two new nodes (following special rules) and expanding them further (lines 24-27) and formula $a \wedge b$ by adding a and b to the field *NEW* and further expanding the current node (lines 29-31). While processing a literal f (lines 16-23), a conflict occurs and the current node is discarded if the field *OLD* of the current node already contains $\neg f$ (lines 17 and 18). When there is no conflict, f is added to the field *OLD* and the current node is further expanded (lines 20 and 21). Expanding a specific node finishes when all formulas have been processed, i.e., the *NEW* field is empty (lines 2-12). If an equal node is already in the *GRAPH*, then this (equal) node is updated with data for the current node (lines 3-6). Otherwise, the current node is added to the *GRAPH*, a new node is created and expanded with the current *NODE* in the field *INCOMING* and all formulas from the field *NEXT* of the current node in the field *NEW* of the created node (lines 7-11).
3. A generalized Büchi automaton is created in the following manner: (1) created nodes become states in the automaton, (2) automaton edges are defined by the *INCOMING* field, (3) labels on edges are defined by literals stored in the field *OLD*, and (4) infinite acceptance of states is imposed, ensuring that, whenever a node contains pUq , some successor node will contain q .
4. The generalized Büchi automaton is converted (i.e., de-generalized) into a Büchi automaton, which can be further used for model-checking.

A finite representation of all possible executions of a DecSerFlow model (i.e., all executions that satisfy all constraints in the model) is obtained by generating the *model automaton* from the *model formula*. The model formula F for a model with n constraints with LTL formulas $f_1, f_2, \dots, f_n$ is a conjunction of LTL formulas for all constraints from the model, i.e., $F = f_1 \wedge f_2 \wedge \dots \wedge f_n$. For

Algorithm 3.1 BASIC algorithm: expanding the graph of nodes

```
1: function expand(NODE, GRAPH)
2:   if  $NODE.NEW = \emptyset$  then {*finished processing node*}
3:     if  $\exists A \in GRAPH : equal(NODE, A)$  then {*equivalent node already processed*}
4:       update_existing(NODE, A); {*just update the existing node*}
5:       return GRAPH;
6:     else
7:        $GRAPH \leftarrow GRAPH \cup \{NODE\}$ ; {*add NODE to GRAPH*}
8:        $NEWNODE \leftarrow createNewNode()$ ; {*create NEWNODE*}
9:        $NEWNODE.INCOMING \leftarrow \{NODE\}$ ;
10:       $NEWNODE.NEW \leftarrow NODE.NEXT$ ;
11:      return expand(NEWNODE, GRAPH); {*expand NEWNODE*}
12:     end if
13:   else
14:      $f \leftarrow getFormula(NODE.NEW)$ ; {*get the next formula for processing*}
15:      $NODE.NEW \leftarrow NODE.NEW \setminus \{f\}$ ;
16:     if  $(f \in P \vee !f \in P) \vee (f = \text{true} \vee f = \text{false})$  then {*f is a literal*}
17:       if  $(f = \text{false}) \vee (!f \in NODE.OLD)$  then {*a contradiction in NODE*}
18:         return GRAPH {*discard current NODE*}
19:       else
20:          $NODE.OLD \leftarrow NODE.OLD \cup \{f\}$ 
21:         return expand(NODE, GRAPH)
22:       end if
23:     end if
24:     if  $f = aUb, aVb, \text{ or } a \vee b$  then
25:        $NODE1 \leftarrow create1(f, NODE)$ ; {*depending the current operator in f*}
26:        $NODE2 \leftarrow create2(f, NODE)$ ; {*depending the current operator in f*}
27:       return expand(NODE2, expand(NODE1, GRAPH));
28:     end if
29:     if  $f = a \wedge b$  then
30:        $NODE.NEW \leftarrow NODE.NEW \cup \{a, b\}$ ;
31:       return expand(NODE, GRAPH);
32:     end if
33:   end if
34: end expand;
```

example, the *model formula* for the DecSerFlow model shown in Figure 1 is a conjunction of formulas for the two *response* constraints, as shown in Figure 3(a). Figure 3(b) shows the Büchi automaton generated by the BASIC algorithm for this model formula. Hence, the language (i.e., all accepted traces) of this automaton represents all possible executions of the DecSerFlow model from Figure 1. Automaton states are represented by ovals such that a single border marks a non-accepting and a double border accepting state. Transitions are represented as directed labeled arcs between states. The arc without a source state marks the initial state.

Fig. 3. Retrieving a finite representation of all possible executions of the model shown in Figure 1

Automata generated from DecSerFlow models can be used for multiple purposes [1, 11]. For example, the *model automaton* and automata generated for each constraint can be used to monitor the state of a model instance (we refer to one execution of a model as to one instance of the model) and constraints during execution. This is done by checking if the trace satisfies the model, i.e., if the automaton accepts the trace. Naturally, when processing the instance state the model automaton is used, and when processing states of constraints automata generated from LTL specifications are used. Note that the generated automata are non-deterministic, and we say that an automaton accepts a trace if the trace can be ‘replayed’ on the automaton in such a way that an accepting state is reached [1]. Given the current execution trace of an instance, the instance or constraint state is determined as follows:

- If the trace is accepted by the automaton, then the instance/constraint is *satisfied*.
- The instance/constraint is *temporarily violated* if the current trace is not accepted but it is a prefix of a trace accepted by the automaton. In other traces, the instance/constraint is *temporarily violated* if the current trace can be ‘replayed’ on the automaton, but all possible replay scenarios lead to non-accepting state.
- If the trace is neither accepted by the automaton nor it is a prefix of an accepted trace, then the instance/constraint is (*permanently*) *violated*. In other cases, the instance/constraint is *violated* if the current trace can not be ‘replayed’ on the automaton at all.

Consider, for example, the situation when the DecSerFlow model shown in Figure 1 is executed by executing task A , i.e., the execution trace is

$$\sigma = A_s, A_c. \quad (1)$$

The model automaton shown in Figure 3(b) suggests that σ satisfies the model because it brings the model automaton to the accepting state S_2 by triggering transitions ‘ $!A_c$ ’ and ‘ A_c ’: $S_0 \xrightarrow[(!A_c)]{A_s} S_1 \xrightarrow[(A_c)]{A_c} S_2$.

In addition to state monitoring, the model automaton can be used to ensure a deadlock-free execution and to verify service models. On the one hand, if the service execution would be driven by the model automaton, deadlocks would be eliminated. On the other hand, we have developed a verification procedures that can detect two types of errors based on model automata. First, a dead task is a model task that can never be executed because its completion is never allowed by transition labels. Second, a model has a conflict if the generated automaton is empty, i.e., it has no states and its language is empty. The DECLARE system used the above described procedures for model verification, monitoring states of instances and constraints and ensuring the deadlock-free execution [13].

4 Applying the Finite Traces Semantics

As explained in Section 1, an execution trace of a web service is a *finite* sequence, i.e., $\sigma = p_1, p_2, \dots, p_n$. Using automata generated for infinite traces can create problems if used for finite executions of DecSerFlow models. For example, if the finite semantics of the until (U) operator is considered, trace σ given in (1) does not satisfy the model formula shown in Figure 3(a). This is because event A_c is not followed by event B_c by the end of the trace (note that this is required by constraint $response(A, B)$). Hence, trace σ does not satisfy the model shown in Figure 1. However, as explained in Section 3, the model automaton shown in Figure 3(b) suggests otherwise: trace σ satisfies the model because it brings the model automaton to the accepting state S_2 .

To avoid this type of errors, the algorithm for automata generation must be adjusted to finite traces, as described in [7]. The infinite semantics of LTL is reflected in the fact that i does not have an upper bound in the definition of the until (U) operator: $\exists_{1 \leq i} : \dots$ (cf. Section 2). The manner in which the infinite acceptance is imposed in the standard algorithm is described in Section 3. As described in [7], finite semantics must be reflected in the upper bound n of i in the until operator: $\sigma \models \varphi U \psi$ if and only if $(\exists_{1 \leq i \leq n} : (\sigma^i \models q \wedge (\forall_{1 \leq j < i} : \sigma^j \models p)))$. The main change in the algorithm is the way accepting conditions are imposed: a finite trace is accepting only if it satisfies all required eventualities (i.e., untils). Formulas that still need to be satisfied are stored in the field NEXT. Therefore, only if the NEXT field of a node does not contain any until (U) formulas, the automaton state generated from this node is accepting [7]. In the remainder of this paper we will use $\mathbf{BASIC}^{FIN}$ to denote the algorithm presented in [8] and modified for finite traces as described in [7]. We will refer to the automaton generated by the $\mathbf{BASIC}^{FIN}$ algorithm as to the $\mathbf{BASIC}^{FIN}$ automaton. Figure 4 shows the $\mathbf{BASIC}^{FIN}$ automaton generated for the model formula given in Figure 3(a). Indeed, this automaton suggests that trace σ does not satisfy the model from

Fig. 4. The BASIC^{FIN} model automaton for model shown in Figure 1

Figure 1, because this trace brings the automaton to the non-accepting state S_1 :

$$S_0 \xrightarrow[A_c \wedge !B_c]{A_s} S_0 \xrightarrow[-]{A_c} S_1.$$

5 Applying the ‘Single Event’ Semantics

In standard LTL, traces are defined over 2^P , which means that σ is a sequence of sets of propositions ($\forall_{1 \leq i} : \sigma_i \subseteq P$). Hence, each element σ_i of a trace is a *set of propositions*. The fact that one element of a trace can refer to multiple properties in standard LTL is semantically expressed in the way the proposition is defined: $\sigma \models p$ if and only if $p \in \sigma_1$ (cf. Section 2). As explained in Section 1, execution trace of a web service is a sequence of *single* events/propositions: $\forall_{1 \leq i \leq n} : \sigma_i \in P$. In order to adjust the semantics of LTL to traces where each element refers to exactly one event, i.e. proposition, we must check if the proposition *is* the first element of the trace: $\sigma \models p$ if and only if $p = \sigma_1$.

Automata that consider traces containing sets of propositions can create problems in DecSerFlow. For example, consider the BASIC^{FIN} model automaton shown in Figure 4 and an instance of the model (cf. Figure 1 on page 2) with the execution trace σ given in (1). As explained in Section 4, this automaton suggests that the instance is not satisfied because trace σ brings the automaton to the non-accepting state S_1 . Moreover, because an accepting state is reachable (i.e., accepting state S_0 is reachable from S_1 via transition $A_c \wedge B_c$), this automaton suggests that the instance is temporarily violated. However, because events are triggered *one by one*, transition $A_c \wedge B_c$ can never be taken. Hence, the state of this instance is actually *permanently violated* because an accepting state can no longer be reached.

In order to eliminate this error, we must adjust the algorithm to consider sequences of single events in the following way. The most important change is strengthening the contradiction test in Algorithm 3.1 (lines 16-18). As described in Section 2, literals that belong to the field *OLD* will become labels on the transitions in the generated automaton. Therefore, if the processed formula f is a literal, contradiction occurs and the current node is discarded if negation of f ($!f$) is already in the field *OLD* (lines 13 and 14).

Algorithm 5.1 shows how contradiction requirements must be strengthened in order to reflect the single event property. The additional requirement is: if we are processing a proposition f and another proposition $l \neq f$ is already in the

Algorithm 5.1 Detecting a contradiction in the ‘single events’ algorithm

```
1: function expand(NODE, GRAPH)
2: ...
3: if  $(f \in P \vee !f \in P) \vee (f = \text{true} \vee f = \text{false})$  then
4:   if  $(f = \text{false}) \vee (!f \in \text{node.Old}) \vee (l \in P \wedge f \neq l \wedge l \in \text{node.Old})$  then
5:     return GRAPH
6:   else
7:      $\text{NODE.Old} \leftarrow \text{NODE.Old} \cup f$ 
8:     return expand(NODE, GRAPH)
9:   end if
10: end if
11: ...
12: end expand;
```

field *OLD*, then this is also a contradiction (line 3). Further, the handling of contradictions and regular situations stays the same: the current node is discarded when a contradiction is detected. Otherwise, further expansion of the current node in the graph is continued.

In addition to strengthening the contradiction requirement, labels on transitions can be displayed in a more concise way. If a label contains one positive proposition and an arbitrary number of negative proposition (e.g., $A_c \wedge !B_c \wedge !C_c$), it can be replaced by a shorter label containing only the positive proposition (e.g., A_c). This is because the latter is implied by the former. Note that making labels shorter is not necessary from the semantical and correctness perspective, but it significantly improves the readability.

Figure 5 shows the BASIC_{SE}^{FIN} automata generated for the model formula given in Figure 3(a). This automaton correctly indicates that an instance with trace σ is *permanently violated* because σ is neither accepted by the automaton, nor it is a prefix of a trace accepted by this automaton. Moreover, tasks *A* and *B* can never be executed (i.e., these are dead tasks) because the language of this automaton does not accept traces that contain A_c or B_c . This is because, as soon as either *A* or *B* would be executed, no finite execution would be able to satisfy *both response* constraints from the model shown in Figure 1.

Fig. 5. The BASIC_{SE}^{FIN} model automaton for model shown in Figure 1

5.1 Correctness Arguments

In this section we give only a brief discussion of the correctness of the single event Algorithm 5.1, for brevity denoted also as A_2 . Algorithm 3.1 (in the sequel denoted as A_1) can also be used to generate automata that accept only single-event traces that satisfy a given model formula F . To this end we need to add a restrictive conjunct to F . Thus, to rule out all multiple-event traces, we run A_1 on the formula $R \wedge F$, where $a_1, a_2, \dots, a_n \in P$ are the events that appear in F , and $R = \Box \bigwedge_{i,j} !(a_i \wedge a_j)$ where $1 \leq i \leq n, 1 \leq j \leq n$ and $i \neq j$.

For the correctness of the standard algorithm A_1 we rely on [5, 8]. Hence, we can establish the correctness of the single-event algorithm A_2 by showing the following: for each run $\mathcal{R}_2$ of the algorithm A_2 applied to a given model formula F , there exists a run $\mathcal{R}_1$ of algorithm A_1 applied to the formula $R \wedge F$, such that the automata produced by $\mathcal{R}_2$ and $\mathcal{R}_1$ are equivalent, i.e., they accept the same language. Moreover, to each state of the automaton generated by A_2 there corresponds an equivalence class of states in the automaton generated by A_1 .

To show this, we construct a run $\mathcal{R}_1$ of algorithm A_1 as we trace the run $\mathcal{R}_2$ of algorithm A_2 . In A_2 run $\mathcal{R}_2$ is applied to F' , which is the normal form of F . Run $\mathcal{R}_1$ simulates $\mathcal{R}_2$ in a “stuttering” manner, i.e., multiple steps of $\mathcal{R}_1$ can correspond to a single step of $\mathcal{R}_2$. Algorithm A_1 is actually applied on the normal form of $R \wedge F$, where F is rewritten to its normal form F' and R is rewritten into $R' = \text{false} \vee \bigwedge_{i,j} (!a_i \vee !a_j)$ such that $1 \leq i \leq n, 1 \leq j \leq n$ and $i \neq j$. In $\mathcal{R}_1$ we process R' before F' . It is straightforward to check that a sequence of node transformations as a result of the processing of R' in $\mathcal{R}_1$ leads to insertion of $n - 1$ literals of the form $!a_i$ in the field *OLD* of each generated node [5]. This ensures that at most one positive proposition a_k can be added to the *OLD* field of each node, such that $!a_k \notin \text{OLD}$. Adding the second positive proposition a_l ($l \neq k$) will automatically invoke a contradiction in the original sense because the *OLD* field already contains the negation $!a_l$.

At each point one can prove that the following invariant holds for the parallel execution of $\mathcal{R}_1$ and $\mathcal{R}_2$. Let run $\mathcal{R}_2$ discard its current node n_2 , because field *OLD* contains proposition l (lines 4-5 in A_2) different from the currently processed formula/proposition f . Then field *OLD* of node n_1 , currently considered by run $\mathcal{R}_1$, contains $!f$. As a consequence, $\mathcal{R}_1$ also discards the node (because of contradiction) and in this way rules out a possible multiple event transition in the automaton.

Besides that, one can show that for each action by $\mathcal{R}_2$ that adds a new node n_1 to the set of nodes (lines 7-11), there exists an action of $\mathcal{R}_2$ that adds a corresponding node n_2 . Both n_1 and n_2 contain in *OLD* only one literal without negation and this literal is the same in both nodes. Nodes n_1 and n_2 will be transformed into equivalent states in the resulting automata. Vice versa, each node added by $\mathcal{R}_2$ corresponds to an equivalent node added by $\mathcal{R}_1$.

6 Experiments

We have performed experiments to measure the effects of the ‘single events’ adjustments on the performance of the **BASIC** and **BASIC^{FIN}** algorithms. We have used the testing method based on randomly generated LTL formulas presented in [3, 8]. Each test set consists of F randomly generated LTL formulas of length L with N propositional variables. Temporal operators U and V are generated with the probability P . A formula of length L is generated in the following way:

- $L = 1$ Randomly generate a propositional variable using a uniform distribution.
- $L = 2$ Randomly generate a unary operator from the set $\{!, \circ\}$. Apply the generated unary operator to a random formula of $L = 1$.
- $L > 2$ Randomly generate an operator from the set $\{!, \circ, \vee, \wedge, U, V\}$. The probability to generate either U or V is $\frac{P}{2}$ and $\frac{1-P}{4}$ to generate the other operators. If the chosen operator is unary, it is applied to a random formula of $L = 1$. If the chosen operator is binary, it is applied to two random formulas: (1) one of length S and the other of length $L - S - 1$. S is generated randomly using a uniform distribution between 1 and $L - 2$ inclusive.

We have performed tests on ten sets containing $F = 100$ randomly generated formulas of varying lengths with 5 propositional variables ($N = 5$) and the probability of $\frac{1}{2}$ to select operators U and V . In each test set, formulas had different lengths: $L \in \{5, 10, 15, 20, 25, 30, 35, 40, 45, 50\}$. For each formula we have generated three automata and used them for two types of tests:

- In order to measure the total effects on the DecSerFlow performance, we have compared automata generated by the **BASIC_{SE}^{FIN}** and **BASIC** algorithms. Although we are aware that **BASIC_{SE}^{FIN}** and **BASIC** produce different results, such a comparison is still interesting. This is because we want to show that using special DecSerFlow LTL results in a better performance than what one would expect from standard LTL.
- In order to measure the effect of single events adjustments on the **BASIC** algorithm, we have compared automata generated by the **BASIC_{SE}** and **BASIC** algorithms. Although the **BASIC_{SE}** algorithm can not be used for DecSerFlow because it uses infinite traces, which can introduce errors (cf. Section 4), adjusting the **BASIC** algorithm to the single event could improve performance if it is applied to other application domains where propositions also hold one-by-one. Note that it might seem more appropriate to compare **BASIC_{SE}** with the **BASIC** algorithm applied to the restricted formula like described in Section 5.1. However, this version has even worse performance than the original **BASIC** algorithm because of the increased complexity of the processed formula, which consists of the restrictive addition R and the original LTL formula (cf. Section 5.1).

When evaluating algorithms for generating automata from LTL formulas, it is a common practice to consider the size of the automaton (i.e., number of states and transitions) and the total time needed to generate the automata [3,

8]. In addition to measuring the automata size and processing time, we have measured the number of times the procedure for *expanding* a node was invoked. This is because strengthening the contradiction requirements with the single event property causes the algorithm to abandon more expansion paths in the graph. Hence, the *expand* procedure is typically invoked fewer times.

Figure 6(a) shows the ratio between the results of the BASIC_{SE}^{FIN} algorithm and the BASIC algorithm. Figure 6(b) shows the ratio between the results of the BASIC_{SE} algorithm and the BASIC algorithm. The results show that imposing either the DecSerFlow requirements (i.e., finite traces of single events) or only the *single events* requirement significantly improves the performance by reducing the processing time, automata size, and number of invocations of the *expand* procedure. For example, only 10% of the original time is needed to process a formula of length 25. A lower number of invocations of procedure *expand* results in shorter processing times. The typical automata size is also reduced because transitions referring to more than one property are eliminated.

Fig. 6. The ratio of number of states, number of transitions, number of node expansion and generation time between various algorithms for $F = 100$ randomly generated formulas of various lengths with $N = 5$ and $P = \frac{1}{2}$

Figure 6 shows that effects become more significant as the length of the formula rises up to the values $L = 25$ (for BASIC_{SE}^{FIN}) and $L = 35$ (for BASIC_{SE}). As the length of formulas further rises, the effect begin to drop. Hence, in order to investigate the effects of our changes on even longer formulas, we have tested the formulas of length $L = 45$ and $L = 50$. Again, these two sets contain $F = 100$ randomly generated formulas of lengths with 5 propositional variables ($N = 5$) and the probability of $\frac{1}{2}$ to select operators U and V . Table 1 shows our results. On the one hand, for the set of formulas with length $L = 45$ the BASIC algorithm failed due to the lack of memory, while the same algorithm did not finish processing the set of formulas with length $L = 50$ within 20 hours (after which we stopped the processing). On the other hand, the average processing

Table 1. Average processing time in milliseconds of various algorithms for $F = 100$ randomly generated formulas with $N = 5$ and $P = \frac{1}{2}$

formula length	BASIC	BASIC _{SE}	BASIC _{SE} ^{FIN}
$L = 45$	out of memory	456.757	274.405
$L = 50$	> 72000000.000	4368.51	2131.766

time of formulas with length $L = 45$ was 456.757 milliseconds and 274.405 milliseconds for the BASIC_{SE} and BASIC_{SE}^{FIN} algorithms, respectively. The BASIC_{SE} algorithm processed formulas with length $L = 50$ in 4368.51 milliseconds and The BASIC_{SE}^{FIN} algorithm in 2131.766 milliseconds, on average. Table 1 shows that dramatic speedups are possible when using the modifications proposed in this paper. Hence, these are now used in our DECLARE system.

The improved performance of the BASIC algorithm is important for enactment of declarative models, because the model formula of a DecSerFlow model is generated as a conjunction of all constraints and, hence, it can be much longer than typical formulas used in model checking. If $L(f)$ denotes length of LTL formula f , then the length of model formula $F = c_1 \wedge c_2 \wedge \dots \wedge c_n$ of a DecSerFlow model with n constraints specified with LTL formulas $c_1, c_2, \dots, c_n$ is $L(F) = n - 1 + \sum_{i=1}^n L(c_i)$. For example, the length of model formula F (cf. Figure 3(a)) is $L(F) = 2 - 1 + \sum_{i=1}^2 5 = 11$, because the length of each *response* constraint in the model shown in Figure 1 is $L(c) = 5$.

7 Related Work

In our previous work, we have proposed using DecSerFlow for declarative specification of web service processes [1]. The SCIFF language is another declarative language [2], which is based on abductive logic programming. While DecSerFlow and SCIFF are similar with respect to their declarative approach to process modeling, both languages have some specific advantages. SCIFF has more expressive power and is more efficient while checking constraints on executed traces (i.e., a posteriori) [11]. However, SCIFF cannot ensure a deadlock-free execution at run-time. Moreover, the adjustments proposed in this paper enable detection of more sophisticated model errors and improves efficiency.

DecSerFlow uses LTL for formal specification of constraints and automata generated from LTL formulas for retrieving the final representation of all possible model executions as all executions that satisfy model constraints. LTL is extensively used in the field of model checking and algorithms for automata generation from LTL formulas based on [5] are proposed in this field. Moreover, improving the performance of these algorithms is an important topic in the field [3, 8, 4]. While these approaches improve the performance of the original algorithm, which works with the standard LTL, changes proposed in this paper improve the performance by limiting the LTL to sequences of single events.

The problem of applying finite traces semantics to standard LTL has been addressed by other researchers. Approach presented in [10] considers only safety

properties and generating finite-trace automata for monitoring running programs. Adjustments of the algorithm for automata generation to finite traces is given in [7]. We use this approach because it does not limit the expressiveness of DecSerFlow [7]

Standard LTL considers properties, rather than events, which means that zero or more properties can hold at any point of time. Methods for model checking for event-based systems use several approaches to LTL for sequences of single events. In [12], an approach is proposed for specifying events indirectly in terms of edges. Edges do not relate directly to occurrence of event, but capture changes of truth/false values of atomic propositions. The Tracta model-checking approach [6] for analysis of concurrent systems uses the special kind of LTL for sequences of single actions: Action Linear Temporal Logic (ALTL). However, this approach uses the Büchi automata following the standard automata-theoretic approach to verification and focuses on solving issues related to hierarchical systems using the Compositional Reachability Analysis (CRA) [6]. ALTL is extended for *fluent model checking* in [9]. Here, instead of using each event occurrence, time intervals between action initiation and termination are considered. A special model checking procedure is proposed for fluent actions. This procedure uses Büchi automata, but avoids the need for using the synchronous product operation [9].

8 Conclusions

Using standard LTL and Büchi automata for enacting DecSerFlow models can cause errors and inefficiencies. This is because of two important differences between the standard model checking problem and the execution of web services processes. This paper describes how the standard LTL and algorithm for automata generation can be adapted in order to fit two special properties of DecSerFlow. Both the adjustments for finite traces described by Giannakopoulou et al. in [7], and the adjustments for sequences for single events described in Section 5 must be used in order to avoid errors. Because automata are generated for the DecSerFlow model formula (which is a conjunction of formulas for all constraints), the performance of the algorithm becomes a potential bottle neck. For the special class of problems where properties hold one at a time, results of our experiments show that the proposed adjustments also significantly decrease the processing time and size of automata. Processing times of days are reduced to seconds. Since the enactment of declarative languages like DecSerFlow require the repeated execution of this procedure, this is highly relevant.

References

1. W.M.P. van der Aalst and M. Pesic. DecSerFlow: Towards a Truly Declarative Service Flow Language. In M. Bravetti, M. Nunez, and G. Zavattaro, editors, *International Conference on Web Services and Formal Methods (WS-FM 2006)*, volume 4184 of *Lecture Notes in Computer Science*, pages 1–23, Vienna, Austria, 2006. Springer-Verlag.

2. M. Alberti, F. Chesani, M. Gavanelli, E. Lamma, P. Mello, M. Montali, S. Storari, and P. Torroni. Computational Logic for Run-Time Verification of Web Services Choreographies: exploiting the SOCS-SI tool. In M. Bravetti and G. Zavattaro, editors, *Proceedings of the 3rd International Workshop on Web Services and Formal Methods (WS-FM'06)*, number 4184 in Lecture Notes in Computer Science, pages 58–72, Heidelberg, Germany, 2006. Springer Verlag.
3. M. Daniele, F. Giunchiglia, and M.Y. Vardi. Improved Automata Generation for Linear Temporal Logic. In *Proceedings of the 11th International Conference on Computer Aided Verification (CAV '99)*, pages 249–260, London, UK, 1999. Springer-Verlag.
4. P. Gastin and D. Oddoux. Fast LTL to Büchi Automata Translation. In *Proceedings of the 13th International Conference on Computer Aided Verification (CAV '01)*, pages 53–65, London, UK, 2001. Springer-Verlag.
5. R. Gerth, D. Peled, M.Y. Vardi, and P. Wolper. Simple On-The-Fly Automatic Verification of Linear Temporal Logic. In *Proceedings of the Fifteenth IFIP WG6.1 International Symposium on Protocol Specification, Testing and Verification XV*, pages 3–18, London, UK, 1996. Chapman & Hall, Ltd.
6. D. Giannakopoulou. *Model Checking for Concurrent Software Architectures*. PhD Thesis, University of London, London, United Kingdom, January 1999.
7. D. Giannakopoulou and K. Havelund. Automata-Based Verification of Temporal Properties on Running Programs. In *ASE '01: Proceedings of the 16th IEEE international conference on Automated software engineering*, pages 412–416, Washington, DC, USA, 2001. IEEE Computer Society.
8. D. Giannakopoulou and F. Lerda. From States to Transitions: Improving Translation of LTL Formulae to Büchi Automata. In *Proceedings of the 22nd IFIP WG 6.1 International Conference Houston on Formal Techniques for Networked and Distributed Systems (FORTE '02)*, volume 2529 of *Lecture Notes in Computer Science*, pages 308–326, London, UK, 2002. Springer-Verlag.
9. D. Giannakopoulou and J. Magee. Fluent Model Checking for Event-Based Systems. In *Proceedings of the 9th European Software Engineering Conference held jointly with 11th ACM SIGSOFT international symposium on Foundations of Software Engineering (ESEC/FSE-11)*, pages 257–266, New York, NY, USA, 2003. ACM.
10. T. Latvala. Efficient Model Checking of Safety Properties. In T. Ball and S.K. Rajamani, editors, *In Model Checking Software. 10th International SPIN Workshop*, number 2648 in Lecture Notes in Computer Science, pages 74–88, Heidelberg, Germany, 2003. Springer Verlag.
11. M. Montali, M. Pesic, W.M.P.van der Aalst, F. Chesani, P. Mello, and S. Storari. Declarative Specification and Verification of Service Choreographies. *ACM Transactions on the Web*. To appear.
12. D.O. Pbreveaun and M. Chechik. Events in Linear-Time Properties. In *Proceedings of the 4th IEEE International Symposium on Requirements Engineering (RE '99)*, pages 123–132, Washington, DC, USA, 1999. IEEE Computer Society.
13. M. Pesic. *Constraint-Based Workflow Management Systems: Shifting Control to Users*. PhD Thesis, Eindhoven University of Technology, Eindhoven, The Netherlands, October 2008.