



CITYSPIN

Cyber-Physical Social Systems for
City-wide Infrastructures

Deliverable 5.3: Monitoring infrastructure for linked-data event streams

Authors	:	Claudio Di Ciccio, Alessio Cecconi
Dissemination Level	:	Public
Due date of deliverable	:	31.05.2019 (v1)
Actual submission date	:	July 22, 2019
Work Package	:	5. Process Mining and Monitoring on Linked Data
Type	:	Report
Version	:	1.0

Abstract

This document provides the definition of a monitoring technique for linked-data event streams with respect to declarative constraints. The theoretical base of D5.1 is extended to deal with a streaming scenario and the result of D5.2 used to interpret the input stream. The technique is grounded in temporal logic and automata theory. The basic step is the reshaping of the constraints as reactive elements, in order to respond as soon as possible to the incoming events. A fine grain analysis is conducted on the constraints with measures at the granularity of the singles events. The validation of the constraints is independently achieved for each constraint activation. Especially, leveraging automata theory, this validation can be performed efficiently online. A detailed algorithm to implement the technique is delivered.

The information in this document reflects only the author's views and nor the FFG neither the Project Team is liable for any use that may be made of the information contained therein. The information in this document is provided "as is" without guarantee or warranty of any kind, express or implied, including but not limited to the fitness of the information for a particular purpose. The user thereof uses the information at his/her sole risk and liability.



FFG

Project Funded by FFG – IKT der Zukunft Programme

Project Number: 861213

Start date: 01.10.2017

Duration: 30 months

History

Version	Date	Reason	Revisited by
1.0	31.05.2019	Submission	CDC, AC

Author List

Project Partner	Name(Initial)	Contact Information
WU Wien	Alessio Cecconi (AC)	cecconi@ai.wu.ac.at
WU Wien	Claudio Di Ciccio (CDC)	claudio.di.ciccio@ai.wu.ac.at

Executive Summary

This deliverable expands the analytical tool delivered by WP5 defining a monitoring technique for linked-data event streams. The previous deliverable D5.2 provided the interface to the incoming linked-data from WP4 and a discovery technique for declarative processes. Leveraging this interface, instead of mining rules out of the event log, the monitor checks the compliance of reactive declarative constraints with respect to the event stream is proposed.

The technique extends the theoretical base of D5.1 to deal with a streaming scenario and uses the result of D5.2 used to interpret the input stream. Grounded in temporal logic and automata theory, the basic step is the reshaping of the constraints as reactive elements, in order to respond as soon as possible to the incoming events. A fine grain analysis is conducted on the constraints with measures at the granularity of the single events. The validation of the constraints is independently achieved for each constraint activation. Especially, leveraging automata theory, this validation can be performed efficiently online. A detailed algorithm to implement the technique is delivered.

This deliverable is largely based on the paper “Interestingness of Traces in Declarative Process Mining: The Janus $LTLp_f$ Approach”[1], published into the *16th International Conference on Business Process Management (BPM18)*.

Contents

History	2
Author List	2
Executive Summary	3
Table of Content	4
List of Figures	5
List of Tables	5
1 Introduction	6
1.1 Project goal	6
1.2 Work package goal	6
1.3 Relation to other work packages	6
1.4 Deliverable goal	6
1.5 Deliverable structure	6
2 Preliminaries	7
2.1 Streaming Event Logs	7
2.2 $LTLp_f$	8
2.3 Automata	9
3 Approach	10
3.1 Reactive Constraints	10
3.2 Measure	11
3.3 Reactive Constraint verification	12
3.4 Online verification through past automata reversion	13
4 Algorithm	14
4.1 Computational Cost.	16
5 Concluding remarks and research outlook	16
6 Summary	16

List of Figures

Figure 1 An automaton verifying $\bigcirc b$	9
--	---

List of Tables

Table 1 Some Declare constraints expressed as Reactive Constraints	10
Table 2 Graphical representation of the separated automata set of $a \square \rightarrow (\ominus b \vee \diamond c)$	13

1 Introduction

1.1 Project goal

The goal of the CitySPIN project is to create a platform for cyber-physical social systems in order to facilitate innovative Smart City infrastructure services. The analysis and monitoring of real data stemming from different sources of smart city infrastructures and services is a pillar of the overarching objective of the project.

1.2 Work package goal

The high-level goal of this WP is to enable process mining and process understanding on Linked Data sources that aggregate, enrich, and filter existing raw data. The operations to that extent are meant to be both off-line, in order to carry out process analytics on historical information and provide insights on the existing processes, and on-line, so as to enable adaptive monitoring on facts as they happen within running process instances.

1.3 Relation to other work packages

WP5 cooperates with the other work packages as follows. First and foremost, process mining and adaptive monitoring will be conducted on data stemming from the Scalable Linked Data Platform, designed, realised and implemented in WP4. The tools generated will be integrated and evaluated in the context of WP7. The process analytics and monitoring of real world information will be subject to secure secure/privacy-aware data access (WP6). The showcase of the project outcome including process analytics and monitoring will be based upon use cases and scenarios elicited and illustrated in the context of WP3.

1.4 Deliverable goal

The goal of Deliverable D5.3 is to define a monitoring infrastructure and architecture for linked data event stream. This continuous monitoring of the stream should be sensitive to the information brought by the event data. More specifically, the monitoring consists in a continuous verification of the adherence of the incoming events sequences to user desired declarative rules. A fine granular measure is defined for this purpose. The approach extend the theory of D5.1 and is empowered by the features of D5.2.

1.5 Deliverable structure

Section 2 illustrates the new concept of event log in a streaming scenario and present the logic system and formal verification instruments upon which our technique is based. Section 3 explains the approach in details, guiding from the specification of constraints as reactive elements, showing how to measure their interaction the the event stream, till how to evaluate them and especially how to do it online. While Sec.3 proves the theoretical robustness of the technique, Section 4 shows step by step the algorithm to implement it. Finally, Section 5 remarks the results and the open questions.

2 Preliminaries

In this section, we illustrate the basic concepts upon which our conformance checking approach is based.

2.1 Streaming Event Logs

In Deliverable D5.1 we defined an event log as the collection of ordered events records. A single event represents the execution of a task, a sequence of events is called trace, and a set of traces compose an event log. Nevertheless, the perspective considered was of a static containers.

A stream is an expanding sequence of events, where new events are recorded continuously. It is like observing the event log while it written, during the execution of the process, rather than when it is completed and stored. Moreover without knowing when it is going to end.

Definition 1 (Streaming event log). *An streaming event log is a tuple $L(E, \Sigma, C, \ell, \iota, \leq, \widehat{\Delta}, \widehat{\alpha})$, where*

$E \triangleq \{e_1, \dots, e_\infty\}$ *is a semi-infinite set of constant symbols named events;*

$\Sigma \triangleq \{a_1, \dots, a_{|\Sigma|}\}$ *is a finite non-empty set of constant symbols named activities;*

$C \triangleq \{\rho_1, \dots, \rho_\infty\}$ *is a semi-infinite set of constant symbols named cases;*

$\ell : E \rightarrow \Sigma$ *is the labelling function connecting events to activities;*

$\iota : E \rightarrow C$ *is the surjective instance function associating events to cases;*

$\leq \subseteq E \times E$ *is the reflexive total order relation over events; we shall denote as $<$ the corresponding strict total order, excluding from \leq the (e, e) pairs;*

$\widehat{\Delta} \triangleq \{\Delta_1, \dots, \Delta_{n+m}\} \triangleq \widehat{\Delta}_{\leq} \cup \widehat{\Delta}_{=}$ *is the universe of interpretation, consisting of $\widehat{\Delta}_{\leq} \triangleq \{\Delta_{\leq_1}, \dots, \Delta_{\leq_n}\}$, the set of ordered (interpretation) domains, and $\widehat{\Delta}_{=} \triangleq \{\Delta_{=1}, \dots, \Delta_{=m}\}$, the set of unordered (interpretation) domains, with $m, n \in \mathbb{N}$ and $\Delta_{\leq} \cap \Delta_{=} = \emptyset$. For $\widehat{\Delta}_{\leq}$ a total order relation \leq is defined. For $\widehat{\Delta}_{=}$ only an equivalence relation $=$ is defined instead.*

$\widehat{\alpha} \triangleq \widehat{\alpha}_{\leq} \cup \widehat{\alpha}_{=} \triangleq \{\alpha_1, \dots, \alpha_{n+m}\}$ *is the universe of attribute functions, consisting of $\widehat{\alpha}_{\leq} \triangleq \{\alpha_{\leq_1}, \dots, \alpha_{\leq_n}\}$ where $\alpha_{\leq_k} : E \rightarrow \Delta_{\leq_k} \cup \{\perp\}$ for $1 \leq k \leq n$ and $\widehat{\alpha}_{=} \triangleq \{\alpha_{=1}, \dots, \alpha_{=m}\}$ where $\alpha_{=l} : E \rightarrow \Delta_{=l} \cup \{\perp\}$ for $1 \leq l \leq m$, associating events to attribute values that are of ordered and unordered domains, respectively; the special symbol \perp denotes absence of knowledge about the attribute value.*

The given definition matches the one defined in D5.1, but for the semi-infiniteness of the events and cases set. The event sequence must have a start event because either the process or its monitoring must start at a precise point in time. If the initial event is not known, e.g., because the monitoring of the process started after its beginning, the starting of the monitoring defines the start of the event sequence. Another crucial point is that, despite the infiniteness of events and cases, the single cases are composed of a finite number of events. This is because by definition a process must have a beginning and an ending. Thus it may be execute an infinite amount of time (infinite number of cases), but each execution is composed by a finite number of events.

Thanks to the process ontologies devised in Deliverable D5.2, it is possible to read process events out of linked data. Nevertheless, instead of verifying the declarative constraints through SHACLare, we are going to propose a new technique.

2.2 LTL_{p_f}

To cater for temporal modalities referring to the past, we resort on an extension of the syntax of Linear Temporal Logic (LTL), namely the one of Linear-time Temporal Logic with Past on Finite Traces (LTL_{p_f}). Even in a streaming settings, the finite semantic assumption holds because cases are finite.

Well-formed LTL_{p_f} formulae are built from an alphabet $\Sigma \ni \{a\}$ of propositional symbols and are closed under the boolean connectives, the unary temporal operators \bigcirc (next) and \ominus (previous), the binary temporal operators \bigcup (Until) and \mathbf{S} (Since):

$$\varphi ::= a \mid (\neg\varphi) \mid (\varphi_1 \wedge \varphi_2) \mid (\bigcirc\varphi) \mid (\varphi_1 \bigcup \varphi_2) \mid (\ominus\varphi) \mid (\varphi_1 \mathbf{S} \varphi_2).$$

From these basic operators it is possible to derive:

- classical boolean abbreviations *True*, *False*, \vee , \rightarrow ;
- constant t_{End} , verified as $\neg \bigcirc \text{True}$, denoting the last instant of the trace;
- constant t_{start} , verified as $\neg \ominus \text{True}$, denoting the first instant of the trace;
- $\diamond\varphi$ as $\text{True} \bigcup \varphi$ indicating that φ holds true eventually before t_{End} ;
- $\varphi_1 \mathbf{W} \varphi_2$ as $(\varphi_1 \bigcup \varphi_2) \vee \square\varphi_1$, which relaxes \bigcup as φ_2 may never hold true;
- $\diamond\varphi$ as $\text{True} \mathbf{S} \varphi$ indicating that φ holds true eventually in the past after t_{start} ;
- $\square\varphi$ as $\neg \diamond \neg \varphi$ indicating that φ holds true from the current instant till t_{End} ;
- $\boxminus\varphi$ as $\neg \diamond \neg \varphi$ indicating that φ holds true from t_{start} to the current instant.

We remark that, w.l.o.g., we consider here the non-strict semantics of \bigcup and \mathbf{S} [2].

Given a trace t , a LTL_{p_f} formula φ is satisfied in a given instant i ($1 \leq i \leq n$) by induction of the following:

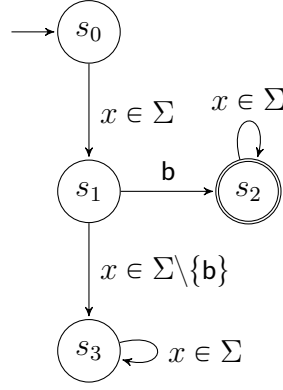
$$\begin{aligned} t, i \models \text{True}; \quad t, i \not\models \text{False}; \quad t, i \models a \text{ iff } t(i) \text{ is assigned with } a; \\ t, i \models \neg\varphi \text{ iff } t, i \not\models \varphi; \quad t, i \models \varphi_1 \wedge \varphi_2 \text{ iff } t, i \models \varphi_1 \text{ and } t, i \models \varphi_2; \\ t, i \models \bigcirc\varphi \text{ iff } i < n \text{ and } t, i+1 \models \varphi; \quad t, i \models \ominus\varphi \text{ iff } i > 1 \text{ and } t, i-1 \models \varphi; \\ t, i \models \varphi_1 \bigcup \varphi_2 \text{ iff } t, j \models \varphi_2 \text{ with } i \leq j \leq n, \text{ and } t, k \models \varphi_1 \text{ for all } k \text{ s.t. } i \leq k < j; \\ t, i \models \varphi_1 \mathbf{S} \varphi_2 \text{ iff } t, j \models \varphi_2 \text{ with } 1 \leq j \leq i, \text{ and } t, k \models \varphi_1 \text{ for all } k \text{ s.t. } j < k \leq i. \end{aligned}$$

In the following, we will classify $\bigcirc, \square, \diamond, \bigcup$ as *future operators*, $\ominus, \boxminus, \diamond, \mathbf{S}$ as *past operators*, and the following pairs of operators as *mirror images*: (i) \bigcirc and \ominus , (ii) \square and \boxminus , (iii) \diamond and \diamond , (iv) \bigcup and \mathbf{S} . We shall also name as *mirror image* of formula φ the temporal formula obtained by replacing all its operators with their mirror images [3], henceforth denoted as φ_M .

Definition 2 (Pure temporal formula [4]). A LTL_{p_f} formula φ is named: **pure past** ($\varphi^\blacktriangleleft$) if it contains only past operators; **pure present** ($\varphi^\blacktriangleright$) if it contains no temporal operators at all; **pure future** ($\varphi^\blacktriangleright$) if it contains only future operators.

For example, $\varphi^\blacktriangleleft = \boxminus(a \mathbf{S} (\diamond b))$, $\varphi^\blacktriangleright = a \wedge b \vee c$, and $\varphi^\blacktriangleright = \square(\diamond a \vee (\bigcirc b) \bigcup c)$ are pure past, pure present, and pure future formulae, respectively.

We argue that separating formulae into ones that refer to the sole past, future, or present, allows for a bi-directional on-line analysis of sub-traces at activation time. The separation theorem, first introduced in [4], proves that such a separation can be obtained.

Figure 1: An automaton verifying $\bigcirc b$

Theorem 1 (Separation theorem (adapted from [4])). *Any propositional temporal formula written with \bigcup , \bigcap , \bigcirc , and \ominus operators can be rewritten as a boolean combination of pure temporal formulae.*

The constructive proof of Theorem 1 in [4] provides a syntactic procedure to pull out \ominus , \bigcap from the scope of \bigcirc , \bigcup in $LTLp_f$ formulae, and vice-versa. It thus provides the base substitution rules such that their recursive application brings to the decomposition of a $LTLp_f$ formula in pure temporal formulae. We capture this notion as follows.

Definition 3 (Temporal separation, separated formula). *Let φ be a $LTLp_f$ formula over Σ . A temporal separation is a function $\mathcal{S} : LTLp_f \rightarrow 2^{LTLp_f \times LTLp_f \times LTLp_f}$. Indicating φ^\bullet , φ° , and φ^\star respectively as pure past, present, and future formulae, defined over Σ as per Def. 2, $\mathcal{S}(\varphi) = \{(\varphi^\bullet, \varphi^\circ, \varphi^\star)_1, \dots, (\varphi^\bullet, \varphi^\circ, \varphi^\star)_m\}$ is such that*

$$\varphi \equiv \bigvee_{j=1}^m (\varphi^\bullet \wedge \varphi^\circ \wedge \varphi^\star)_j.$$

We call separated formula any element in the co-domain of $\mathcal{S}(\varphi)$.

For example, $(\ominus b \vee \diamond c) \equiv ((\ominus b) \wedge (\text{True}) \wedge (\text{True})) \vee ((\text{True}) \wedge (\text{True}) \wedge (\diamond c))$.

2.3 Automata

A (deterministic finite-state) automaton is a rooted finite-state labelled transition system $A = (\Sigma, S, \delta, s_0, S_F) \in \mathcal{A}$, where: Σ is the alphabet; S is a finite non-empty set of states; $\delta : S \times \Sigma \rightarrow S$ is a (total) transition function; s_0 is the initial state; $S_F \subseteq S$ is the set of accepting states. An automaton *accepts* a trace $t = \langle e_1, \dots, e_n \rangle$ if a walk of tuples $\langle (s_0, e_1, s_1), \dots, (s_{n-1}, e_n, s_n) \rangle$, namely a *replay*, exists such that $s_i = \delta(s_{i-1}, e_i)$ for $1 \leq i \leq n$ and $s_n \in S_F$. We shall name s_i in tuple (s_{i-1}, e_i, s_i) as *current state* of the replay at instant i . The set of all traces accepted by A is named *language* of A , denoted as $\mathcal{L}(A)$. Figure 1 depicts an automaton whose language consists of all traces of length $n \geq 2$ having b as its second event. Considering the event log of Example 1, it accepts only $t_4 = \langle d, b, a, e \rangle$. Reportedly approaches exist that build automata verifying any formula of LTL [5], Linear-time Temporal Logic with Past (LTLp) [6, 7], or Linear Temporal Logic on Finite Traces (LTL_f) [8] on traces: such automata accept all and only the traces such that all events satisfy a formula φ . We indicate them as A_φ . The automaton of Fig. 1, e.g., verifies $\bigcirc b$. Notice that automata verification considers whole traces as either satisfying or violating a formula.

Table 1: Some Declare constraints expressed as Reactive Constraints

Constraint	LTL _f expression [9]	RCon	Separation degree
Participation(a)	$\diamond a$	$t_{Start} \square \rightarrow \diamond a$	1
Init(a)	a	$t_{Start} \square \rightarrow a$	1
End(a)	$\square \diamond a$	$t_{End} \square \rightarrow a$	1
RespondedExistence(a, b)	$\diamond a \rightarrow \diamond b$	$a \square \rightarrow (\diamond b \vee \diamond b)$	2
Response(a, b)	$\square(a \rightarrow \diamond b)$	$a \square \rightarrow \diamond b$	1
AlternateResponse(a, b)	$\square(a \rightarrow \diamond b) \wedge \square(a \rightarrow \bigcirc(\neg a \mathbf{W} b))$	$a \square \rightarrow \bigcirc(\neg a \mathbf{U} b)$	1
ChainResponse(a, b)	$\square(a \rightarrow \diamond b) \wedge \square(a \rightarrow \bigcirc b)$	$a \square \rightarrow \bigcirc b$	1
Precedence(a, b)	$\neg b \mathbf{W} a$	$b \square \rightarrow \diamond a$	1
AlternatePrecedence(a, b)	$(\neg b \mathbf{W} a) \wedge \square(b \rightarrow \bigcirc(\neg b \mathbf{W} a))$	$b \square \rightarrow \ominus(\neg b \mathbf{S} a)$	1
ChainPrecedence(a, b)	$(\neg b \mathbf{W} a) \wedge \square(\bigcirc b \rightarrow a)$	$b \square \rightarrow \ominus a$	1

3 Approach

Here we present the concepts upon which our approach is built, beginning with the core notion of Reactive Constraints (RCons). RCons are meant to bear the interestingness semantics, because the role of activation is singled out from the rest of exerted conditions: only if the activation α “triggers” the constraint *and* a LTL_{p_f} formula φ is verified on the trace, then its fulfilment is interesting.

3.1 Reactive Constraints

In Deliverable D5.1 we already suggested that a constraint is composed of an activation component and a reaction one. We are now giving a formal definition for it.

Definition 4 (Reactive Constraint (RCon)). *Given an alphabet Σ , let $\alpha \in \Sigma \cup \{t_{start}, t_{End}\}$ be an activation, and φ be a LTL_{p_f} formula over Σ . A Reactive Constraint (RCon) Ψ is a pair (α, φ) hereafter denoted as $\Psi \triangleq \alpha \square \rightarrow \varphi$. We denote the set of all RCons over Σ as \mathcal{R} .*

In the following, we will assume traces, automata, LTL_{p_f} formulas and RCons to be all defined over a shared alphabet Σ . Constraints of declarative process languages such as Declare can be expressed as RCons, as shown in Table 1. For example, the RCon corresponding to Precedence(d, a) is $a \square \rightarrow \diamond d$. Activations in constraints are identified according to the classification of [10]. RCons can include non-standard Declare constraints such as $a \square \rightarrow (\ominus b \vee \diamond c)$, which imposes that if an event a occurs in a trace, either b immediately precedes it, or c eventually occurs after. We say that a *activates* the RCon.

We remark that although φ is a LTL_{p_f} formula, semantics of RCons detach from classical LTL_{p_f} in that every occurrence of α triggers a new verification of φ on the trace.

Definition 5 (Activator, triggering trace). *Given a trace $t \in \Sigma^*$ of length n and an instant i s.t. $1 \leq i \leq n$, an RCon $\Psi = \alpha \square \rightarrow \varphi$ is activated at i if $t, i \models \alpha$. Event $t(i)$ is then named activator of Ψ . A trace in which at least an activator of Ψ exists, is triggering for Ψ .*

Consider, e.g., the event log from Example 1 and $\Psi = a \square \rightarrow \diamond d$, i.e., Precedence(d, a) in Declare. Ψ is activated in trace $t_4 = \langle d, b, a, e \rangle$ by $t_4(3)$; in trace $t_5 = \langle a, d, a, c, a \rangle$ it is activated

Deliverable 5.3 – v1.0

by $t_5(1)$, $t_5(3)$, and $t_5(5)$; in trace $t_6 = \langle b, c, d, e \rangle$ it is never activated; in trace $t_7 = \langle b, c, a \rangle$ is activated by $t_7(3)$. Therefore t_4, t_5 , and t_7 are triggering for Ψ .

Definition 6 (Interesting fulfilment). *Given a trace $t \in \Sigma^*$ of length n , an instant i s.t. $1 \leq i \leq n$, and an RCon $\Psi = \alpha \square \rightarrow \varphi$, Ψ is interestingly fulfilled at i if $t, i \models \alpha$ and $t, i \models \varphi$. The RCon is violated at instant i if $t, i \models \alpha$ and $t, i \not\models \varphi$. Otherwise, the RCon is unaffected.*

Consider the following example Example 1

Example 1. Let $L = \{t_1^{25}, t_2^{15}, t_3^{10}, t_4^{20}, t_5^5, t_6^{20}, t_7^5, \dots\}$ be a streaming event log, defined on the alphabet of events $\Sigma = \{a, b, \dots, i\}$, where $t_1 = \langle d, f, a, f, c, a, f, b, a, f \rangle$, $t_2 = \langle f, e, d, c, b, a, g, h, i \rangle$, $t_3 = \langle a, d, a, a, a, a, a, a, a, a, a, a, a, a, a, a, a, a, a, a, a, c \rangle$, $t_4 = \langle d, b, a, e \rangle$, $t_5 = \langle a, d, a, c, a \rangle$, $t_6 = \langle b, c, d, e \rangle$, $t_7 = \langle b, c, a \rangle$. We have that $t_1(7) = f$, $t_{1[3:7]} = \langle a, f, c, a, f \rangle$, $t_{2R} = \langle i, h, g, a, b, c, d, e, f \rangle$.

Given a formula $\Psi = a \square \rightarrow \diamond d$. In trace t_4 the RCon is interestingly fulfilled by $t_4(3)$; in trace t_5 it is interestingly fulfilled by $t_5(3)$ and $t_5(5)$, and violated by $t_5(1)$; in trace t_6 it is unaffected at all instants, i.e., neither interestingly fulfilled nor violated; in trace t_7 it is violated by $t_7(3)$. We remark that instants at which a Declare constraint is satisfied can be such that the corresponding RCon is unaffected, i.e., when $t, i \models \varphi$ but $t, i \not\models \alpha$.

As a result each activation of the constraint is independent from the others and can be verified while new events flow into the stream. Either the violation or satisfaction of the constraint is triggered as soon as the, respectively, wrong or right events is recorded into the stream. This include also the end of the trace, if the constraint checks some property that should hold within the entire trace.

3.2 Measure

Because each activator triggers a new check for the RCon, the degree of interestingness of a trace is analysed at the level of events as follows.

Definition 7 (Interestingness degree). *Given a trace $t \in \Sigma^*$ and an RCon $\Psi = \alpha \square \rightarrow \varphi$, we define the interestingness degree function $\zeta : \mathcal{R} \times \Sigma^* \rightarrow [0, 1] \subseteq \mathbb{R}$ as follows:*

$$\zeta(\Psi, t) = \begin{cases} \frac{|\{i : t, i \models \alpha \text{ and } t, i \models \varphi\}|}{|\{i : t, i \models \alpha\}|} & \text{if } |\{i : t, i \models \alpha\}| \neq 0; \\ 0 & \text{otherwise.} \end{cases}$$

Intuitively the interestingness degree measures the percentage of activations leading to (interesting) fulfilment within the trace. For instance, the interestingness degree of $\Psi = a \square \rightarrow \diamond d$ in $t_5 = \langle a, d, a, c, a \rangle$ is $\zeta(\Psi, t_5) = 0.667$, because it is interestingly fulfilled by 2 activators out of 3. All the 20 activators of Ψ in $t_3 = \langle a, d, a, \dots, a, c \rangle$ lead to interesting fulfilment, except one (the first a event). Therefore $\zeta(\Psi, t_3) = 0.950$. Both results give a more detailed information of the RCon behaviour in the trace compared to a strict conformance checking, which in these examples would have return 0 because of the presence of violations.

Also this result, aggregated at the level of the single trace, can be computed online and updated as soon as new events pertaining that trace are recorded into the stream. Next, we introduce the computation of two measures for the whole streaming event log, adapting the classical definition of [11]: *support* measures how often the constraint is (interestingly) satisfied in the whole event log; *confidence* quantifies how the constraint is satisfied in the triggering traces in the event log.

Definition 8 (Support). Given an event log $L = \{t_1, t_2, \dots, t_m\} \in \mathbb{M}(\Sigma^*)$, and an RCon $\Psi \in \mathcal{R}$, we define as support of Ψ on L the function $\sigma : \mathcal{R} \times \mathbb{M}(\Sigma^*) \rightarrow [0, 1] \subseteq \mathbb{R}$ calculated as the average of interestingness degree values of Ψ over all traces $t_j \in L$, with $1 \leq j \leq m$:

$$\sigma(\Psi, L) = \begin{cases} \frac{\sum_{j=1}^m \zeta(\Psi, t_j)}{|L|} & \text{if } |L| > 0; \\ 0 & \text{otherwise.} \end{cases}$$

Definition 9 (Confidence). Given an event log $L = \{t_1, t_2, \dots, t_m\} \in \mathbb{M}(\Sigma^*)$ and an RCon $\Psi \in \mathcal{R}$, let $\tilde{L} = \{\tilde{t}_1, \tilde{t}_2, \dots, \tilde{t}_p\}$ with $1 \leq p \leq m$ be the portion of L that consists of all the traces triggering Ψ . We define the confidence of Ψ on L the function $\kappa : \mathcal{R} \times \mathbb{M}(\Sigma^*) \rightarrow [0, 1] \subseteq \mathbb{R}$ as the average of interestingness degree values of Ψ over the triggering traces $\tilde{t}_j \in \tilde{L}$:

$$\kappa(\Psi, L) = \begin{cases} \frac{\sum_{j=1}^p \zeta(\Psi, \tilde{t}_j)}{|\tilde{L}|} & \text{if } |\tilde{L}| > 0; \\ 0 & \text{otherwise.} \end{cases}$$

As seen above, e.g., the interestingness degree of the RCon of Precedence(d, a), $\Psi = a \square \rightarrow \diamond d$, is $\zeta(\Psi, t_5) = \frac{2}{3} = 0.667$. Averaging the interestingness degree of Precedence(d, a) on all traces (including their multiplicities), we obtain the support $\sigma(\Psi, L)$, which amounts to $\frac{(1 \times 25) + (1 \times 15) + (0.95 \times 10) + (1 \times 20) + (0.67 \times 5) + (0 \times 20) + (0 \times 5)}{100} = 0.728$. The value of Confidence $\kappa(\Psi, L)$ is $\frac{(1 \times 25) + (1 \times 15) + (0.95 \times 10) + (1 \times 20) + (0.67 \times 5) + (0 \times 5)}{80} = 0.910$, thus higher than support, because Ψ is not activated in trace t_6 , hence the lower denominator. ?? shows in detail the interestingness degree, support, and confidence of constraints Precedence(d, a) and $a \square \rightarrow (\ominus b \vee \diamond c)$ on the event log seen in Example 1.

Considering each activation independently to compute interestingness degree, support and confidence, allows for higher resilience to noise in event logs. This is particularly evident in t_3 , in which 19 occurrences of a , the activation, are preceded by d , thus fulfilling the constraint. Only the first a leads to violation. Considering the whole trace as fulfilling or not, as in [12, 13], would lead to an interestingness degree of 0, instead of 0.95, thus decreasing support and confidence too. Because the constraints returned by discovery techniques are those that have a support and confidence above user-defined thresholds, that could lead to a loss of information.

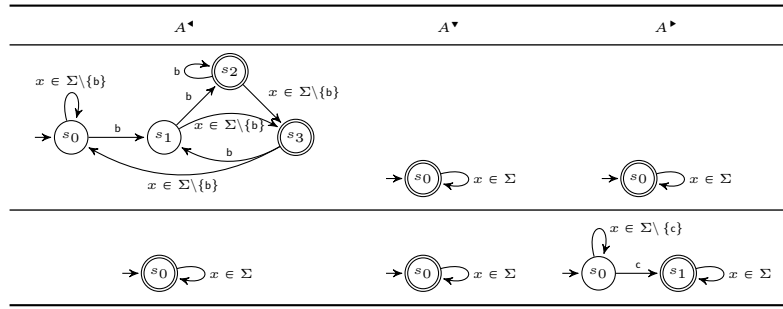
3.3 Reactive Constraint verification

Once an RCon $\Psi = \alpha \square \rightarrow \varphi$ is activated, its fulfilment relies on the verification of the $LTLp_f$ formula φ at the instant of activation. As seen in Theorem 1 it is possible to separate a $LTLp_f$ formula into sub-formulae, each containing either only past, only future, or no temporal operators. Therefore its verification can be decoupled by splitting the trace in two independent sub-traces: one from the beginning to the activator, with which the sole past operators are verified, and one from the activator on, concerning only future operators.

Lemma 1. Given a pure past formula φ^\bullet , a pure present formula φ° , a pure future formula $\varphi^\blacktriangleright$, a trace $t \in \Sigma^*$ of length n and an instant i s.t. $1 \leq i \leq n$, the following hold true:

$$t, i \models \varphi^\bullet \equiv t_{[1,i]}, i \models \varphi^\bullet; \quad t, i \models \varphi^\circ \equiv t_{[i,i]}, i \models \varphi^\circ; \quad t, i \models \varphi^\blacktriangleright \equiv t_{[i,n]}, i \models \varphi^\blacktriangleright.$$

The lemma follows from definition of $LTLp_f$ semantics. For example, evaluating $\varphi^\bullet = \diamond d$ on $t_2 = \langle f, e, d, c, b, a, g, h, i \rangle$ from Example 1 at instant $i = 6$ is equivalent to evaluating it on $t_{2[1:6]} = \langle f, e, d, c, b, a \rangle$ at i , because φ^\bullet concerns only the prefix of t_2 . Instead $\varphi^\blacktriangleright = \diamond c$ at i concerns only the suffix, $t_2[6 : 9] = \langle a, g, h, i \rangle$.

Table 2: Graphical representation of the separated automata set of $a \square \rightarrow (\ominus b \vee \diamond c)$


Theorem 2 (Trace sub-valuation). *Given a $LTLp_f$ formula φ , a trace $t \in \Sigma^*$ of length n and an instant i s.t. $1 \leq i \leq n$, we have that $t, i \models \varphi$ if and only if $t_{[1:i]}, i \models \varphi^\blacktriangleleft$, $t_{[i:i]}, i \models \varphi^\blacktriangleright$, and $t_{[i:n]}, i \models \varphi^\blacktriangleright$ for at least a $(\varphi^\blacktriangleleft, \varphi^\blacktriangleright, \varphi^\blacktriangleright) \in \mathcal{S}(\varphi)$.*

The proof follows from Theorem 1 and Lemma 1. Consider, e.g., the RCon $\Psi = a \square \rightarrow (\ominus b \vee \diamond c)$. It follows that $\varphi = \ominus b \vee \diamond c$ and $\mathcal{S}(\varphi) = \{(\ominus b, True, True), (True, True, \diamond c)\}$. Applying Theorem 2 on $t_2 = \langle f, e, d, c, b, a, g, h, i \rangle$ from Example 1, we have that $t_2, 6 \models \varphi$ if (i) $\langle f, e, d, c, b, a \rangle, 6 \models \ominus b$ or (ii) $\langle a, g, h, i \rangle, 6 \models \diamond c$, aside of *True* formulae which are trivially satisfied. Because the first holds true, we conclude that φ is satisfied by $t_2(6)$.

As seen in Section 2, a formula φ can be verified on a trace t by checking whether t is accepted by automaton A_φ . Given a $LTLp_f$ formula φ and its temporal separation $\mathcal{S}(\varphi)$, we thus introduce the notion of separated automata set, namely a set of triples of automata, each verifying a triple of pure temporal formulae in $\mathcal{S}(\varphi)$.

Definition 10 (Separated automata set (sep.aut.set)). *Given a $LTLp_f$ formula φ we define as separated automata set (sep.aut.set) $\mathcal{A}^{\blacktriangleleft\blacktriangleright\blacktriangleright} \in 2^{\mathcal{A} \times \mathcal{A} \times \mathcal{A}}$ the set of triples $A^{\blacktriangleleft\blacktriangleright\blacktriangleright} = (A^\blacktriangleleft, A^\blacktriangleright, A^\blacktriangleright) \in \mathcal{A} \times \mathcal{A} \times \mathcal{A}$ such that $A^\blacktriangleleft = A_{\varphi^\blacktriangleleft}$, $A^\blacktriangleright = A_{\varphi^\blacktriangleright}$, and $A^\blacktriangleright = A_{\varphi^\blacktriangleright}$ for every $(\varphi^\blacktriangleleft, \varphi^\blacktriangleright, \varphi^\blacktriangleright) \in \mathcal{S}(\varphi)$. We denote as separation degree D the number of triples of $\mathcal{A}^{\blacktriangleleft\blacktriangleright\blacktriangleright}$.*

Considering the latest example $A^{\blacktriangleleft\blacktriangleright\blacktriangleright} = \{(A_{\ominus b}, A_{True}, A_{True}), (A_{True}, A_{True}, A_{\diamond c})\}$, the separation degree is 2. Table 1 lists the separation degrees of some RCons of Declare. In light of the above, we derive from Theorem 2 the following.

Theorem 3 (Trace sub-valuation through automata). *Given a $LTLp_f$ formula φ , its sep.aut.set $A^{\blacktriangleleft\blacktriangleright\blacktriangleright}$, a trace $t \in \Sigma^*$ of length n , and an instant i s.t. $1 \leq i \leq n$, we have that $t, i \models \varphi$ if and only if $t_{[1:i]} \in \mathcal{L}(A^\blacktriangleleft)$, $t_{[i:i]} \in \mathcal{L}(A^\blacktriangleright)$ and $t_{[i:n]} \in \mathcal{L}(A^\blacktriangleright)$ for at least a $(A^\blacktriangleleft, A^\blacktriangleright, A^\blacktriangleright) \in A^{\blacktriangleleft\blacktriangleright\blacktriangleright}$.*

In the example, the application of Theorem 3 entails that $a \square \rightarrow (\ominus b \vee \diamond c)$ is interestingly fulfilled by $t_2(6)$ if $t_{2[1:6]} = \langle f, e, d, c, b, a \rangle \in \mathcal{L}(A_{\ominus b})$ or $t_{2[6:9]} = \langle a, g, h, i \rangle \in \mathcal{L}(A_{\diamond c})$.

3.4 Online verification through past automata reversion

To the best of our knowledge, there is no available technique to build automata that verify $LTLp_f$ formulae. We thus exploit Theorem 2 to rely on the readily available techniques for LTL_f , i.e., without past operators, as described in [8]. To this extent, we rely on mirror images and reversed automata.

Lemma 2. *Let $t \in \Sigma^*$ be a trace of length n and t_R its reverse. Given a pure past formula $\varphi^\blacktriangleleft$, and its mirror image $\varphi_M^\blacktriangleleft$, we have that $t, n \models \varphi^\blacktriangleleft$ iff $t_R, 1 \models \varphi_M^\blacktriangleleft$.*

The proof follows from the semantics of future and past operators of $LTLp_f$ provided in Section 2. For instance, verifying $\varphi^\blacktriangleleft = \diamond d$ on $t_2 = \langle f, e, d, c, b, a, g, h, i \rangle$ from Example 1 at instant

$i = 9$ is equivalent to verifying $\varphi_M^\star = \Diamond d$ on $t_{2R} = \langle i, h, g, a, b, c, d, e, f \rangle$ at $i = 1$. Notice that this holds for sub-traces too, thus verifying φ^\star on t_2 at instant $i = 6$ is equivalent to verifying φ_M^\star over $t_{2[1:6]R} = \langle a, b, c, d, e, f \rangle$ at $i = 1$ in the light of Lemma 1.

It follows from Lemma 2 that any pure past formula can be seen as a pure future one on a reversed trace. Therefore the automaton verifying the mirror image of φ^\star can be used for verification on the reversed trace, as stated in the following.

Corollary 3.1. *Let $A_{\varphi_M^\star}$ be the automaton verifying φ_M^\star . Then $t, n \models \varphi^\star$ iff $t_R \in \mathcal{L}(A_{\varphi_M^\star})$.*

Notice that φ_M^\star is a pure future formula, therefore $A_{\varphi_M^\star}$ can be built by applying the technique of [8] for LTL_f . Furthermore, it is possible to transform the obtained automaton in order to read directly the original trace t thanks to the property of *closure under reversion* of regular languages [14].

Definition 11 (Reversed automaton [14]). *Given a trace $t \in \Sigma^*$, its reverse t_R , and the automaton $A \in \mathcal{A}$, the reversed automaton $\overleftarrow{A} \in \mathcal{A}$ is an automaton such that A accepts t if and only if \overleftarrow{A} accepts t_R , i.e., $t \in \mathcal{L}(A)$ iff $t_R \in \mathcal{L}(\overleftarrow{A})$.*

From Lemma 2 and Corollary 3.1 we derive the following.

Theorem 4 (Valuation through reversed automaton of mirror image). *Let φ^\star be a pure past formula and φ_M^\star its mirror image. Let $A_{\varphi_M^\star} \in \mathcal{A}$ be the automaton verifying φ_M^\star . Given a trace $t \in \Sigma^*$ of length n , we have that: $t, n \models \varphi^\star$ iff $t \in \mathcal{L}(\overleftarrow{A}_{\varphi_M^\star})$.*

Consider the RCon $\Psi = a \square \rightarrow (\ominus b \vee \Diamond c)$ and the pure past formula of its sep.aut.set $\varphi^\star = \ominus b$. It is activated by $t_2(6)$. Its mirror image is $\varphi_M^\star = \circ b$. With automaton $A_{\varphi_M^\star}$, depicted in Fig. 1, we can verify φ_M^\star over trace $t_{2[1:6]R} = \langle a, b, c, d, e, f \rangle$ at $i = 1$, thereby verifying φ^\star over $t_{2[1:6]}$ as per Lemma 2. Thanks to Theorem 4, φ_M^\star can be verified on $t_{2[1:6]}$ with the reversed automaton $\overleftarrow{A}_{\varphi_M^\star}$, depicted in the top-left corner of Table 2.

We remark that in this way the pure past formulae of sep.aut.sets can be verified by parsing sub-traces from the beginning of the trace till the activator event.

4 Algorithm

Algorithm 1 shows the pseudo-code of our on-line technique to compute the interestingness degree $\zeta(\Psi, t)$ of a RCon Ψ with respect to a trace t . Its fundamental data structure is a set of pairs, each associating an automaton to its current state, as the replay of the trace proceeds. We call it *Janus state* and denote it as \mathcal{J} .

More specifically, only future automata of sep.aut.sets are considered. The naïve approach would indeed parse trace t , and apply the check of Theorem 3 whenever Ψ is activated. This is an impractical solution, because it requires the replay of prefix $t_{[1:i]}$ and suffix $t_{[i:n]}$ on the respective automata at every instant of activation i . We save computation time by keeping track of the past valuation state, so that at each activation it is already known, thus improving on the sub-valuation of pure past formulae φ^\star . We rely on the fact that automata preserve the history of replays in the reached state: if at instant i the current state of A^\star is s_i , then at $i + 1$ it is known that $s_{i+1} = \delta(s_i, t(i + 1))$. To this extent, the algorithm requires all past automata A^\star to be already reversed as per Theorem 4.

The runtime of the algorithm will be explained considering as input:

- $t = t_1 = \langle d, f, a, f, c, a, f, b, a, f \rangle$ from Example 1,

Algorithm 1: Computing the interestingness degree of an RCon $\Psi = \alpha \square \rightarrow \varphi$ on trace t , given the sep.aut.set $\mathcal{A}^{\star\star}$ of φ

```

1  $\mathcal{O} \leftarrow$  empty bag ;
2 while event  $t(i) \in t$  do
3   foreach  $A^{\star} \in \mathcal{A}^{\star\star}$  do perform transition  $t(i)$  on  $A^{\star}$ ;
4   if  $t(i) = \alpha$  then // Activation triggered
5      $\mathcal{J} \leftarrow$  empty set of pairs ; //  $\mathcal{J}$  stores the replay-state of future
      automata for one activation
6     foreach  $(A^{\star}, A^{\star}, A^{\star}) \in \mathcal{A}^{\star\star}$  do
7       if  $A^{\star}$  is in an accepting state and  $A^{\star}$  accepts  $\langle t(i) \rangle$  then
8         Take  $A^{\star} = (\Sigma, S^{\star}, \delta^{\star}, s_0^{\star}, S_F^{\star})$  and add  $(s_0^{\star}, A^{\star})$  to  $\mathcal{J}$ 
9       add  $\mathcal{J}$  to  $\mathcal{O}$ ; //  $\mathcal{O}$  collects replay-states for all activations
10    foreach  $\mathcal{J} \in \mathcal{O}$  do
11      foreach  $(s^{\star}, A^{\star}) \in \mathcal{J}$  do  $s^{\star} \leftarrow \delta^{\star}(s^{\star}, t(i))$  // Perform  $t(i)$  on  $A^{\star}$  and save
        state ;
12 if  $|\mathcal{O}| > 0$  then return  $\frac{|\{\mathcal{J} \in \mathcal{O} : \text{at least a } (s^{\star}, A^{\star}) \in \mathcal{J} \text{ is s.t. } s^{\star} \in S_F^{\star}\}|}{|\mathcal{O}|}$  else return 0 ;

```

- $\Psi = a \square \rightarrow (\ominus b \vee \diamond c)$, and
- the sep.aut.set $\left\{ (A_{\ominus b}^{\star}, A_{True}^{\star}, A_{True}^{\star}), (A_{True}^{\star}, A_{True}^{\star}, A_{\diamond c}^{\star}) \right\}$ of φ , with past automaton $A_{\ominus b}^{\star}$ already reversed as depicted in Table 2.

Let i denote the current instant in the following.

At lines 1–2, a bag of Janus states \mathcal{O} is initialised, to store the states of current replays for the verification of pure future formulae. Every replay is triggered by the occurrence of an activation, thus every Janus state refers to an activator. Thereupon, trace t is parsed one event at a time starting with the left-first one, e.g., $t(1) = d$. We remark that no knowledge is assumed on the subsequent events of the trace, as per the on-line setting. At line 3, past automata replay t performing each transition $t(i)$ as it is read, i.e., not waiting for the activation to occur. By contrast, future automata will begin with independent replays at each occurrence of an activator. Therefore every A^{\star} automaton starts a replay from $i = 1$. At line 4, the activator is captured, as, e.g., when i is equal to 3, 6, and 9, i.e., when $t(i) = a$. Consequently at line 5 a new Janus state \mathcal{J} is initialised to store information on the new replay. At line 7 it is checked that, for every triple in the sep.aut.set, (i) A^{\star} is in an accepting state, and (ii) A^{\star} accepts the trace made of the current event. If this is the case the replay on the future automaton of the triple, A^{\star} , can start. At line 8 the pair consisting of A^{\star} and its initial state s_0^{\star} is added to \mathcal{J} . In the example, the triple $(A_{\ominus b}^{\star}, A_{True}^{\star}, A_{True}^{\star})$ at $i = 3$ and $i = 6$ has $A_{\ominus b}^{\star}$ not accepting the prefix $t_{[1:i]}$ because $t(i-1) \neq b$. On the contrary, the replay of A^{\star} can start at $i = 9$. For what the triple $(A_{True}^{\star}, A_{True}^{\star}, A_{\diamond c}^{\star})$ is concerned, $A_{\diamond c}^{\star}$ always starts a new replay because A_{True}^{\star} and A_{True}^{\star} accept any trace.

We remark that for A^{\star} no state is retained because the scope of a pure present formula is limited to a single event. Notice that because every triple $(A^{\star}, A^{\star}, A^{\star}) \in \mathcal{A}^{\star\star}$ represents a conjunctive formula, if A^{\star} is not in an accepting state then the activation leads the entire triple to violation, regardless of the replay on A^{\star} .

The new Janus state \mathcal{J} is added to \mathcal{O} at line 9. In the example, at $i = 9$, $\mathcal{O} = \left\{ \{(s_1, A_{\diamond c}^{\star})\}, \{(s_0, A_{\diamond c}^{\star})\}, \{(s_0, A_{\diamond c}^{\star}), (s_0, A_{True}^{\star})\} \right\}$. At lines 10–11 all states in every $\mathcal{J} \in$

\mathcal{O} are updated by executing the current transition on their respective future automata. For example, upon the reading of $t(5) = c$, $\{(s_0, A_{\diamond c}^*)\}$ is updated to $\{(s_1, A_{\diamond c}^*)\}$.

An event in a trace interestingly fulfils an RCon if, upon activation, at least a triple (A^*, A^*, A^*) is such that A^* , A^* , and A^* all accept the respective sub-traces, as per Theorem 3. By construction, this holds true if at least a future automaton in \mathcal{J} is in its accepting state at the end of the replay. To measure the interestingness degree of the RCon, we thus compute the ratio between the number of all such Janus states and the cardinality of \mathcal{O} at line 12. For instance, at $i = 10$, $\mathcal{O} = \{\mathcal{J}_1, \mathcal{J}_2, \mathcal{J}_3\} = \{\{(s_1, A_{\diamond c}^*)\}, \{(s_0, A_{\diamond c}^*)\}, \{(s_0, A_{\diamond c}^*), (s_0, A_{true}^*)\}\}$ and $|\mathcal{O}| = 3$. Only \mathcal{J}_1 and \mathcal{J}_3 contain automata in accepting states, therefore $\zeta(\Psi, t) = \frac{2}{3} = 0.667$.

4.1 Computational Cost.

To compute the asymptotic computational cost, we consider the worst case scenario, occurring when at each instant (i) Ψ is activated, and (ii) all past and present automata are in an accepting state. In such a case, given an event log L , a trace t of length t , an RCon $\Psi = \alpha \square \rightarrow \varphi$ for which the sep.aut.set \mathcal{A}^{***} of φ is generated with separation degree D , the cost of verifying Ψ on t is:

$$\sum_{j=1}^n (D + (n - j)D) = nD + D \sum_{j=1}^n (n - j) = Dn \left(1 + \frac{(n - 1)}{2}\right).$$

The cost is linear in the number of activations, as each one requires a single replay of the trace for every automaton. For each activation only trace suffixes are replayed, owing to our optimisation over past formulae, hence the $1/2$ factor. For $D \ll n$, the upper-bound is $O(n^2)$, which is comparable to state-of-the-art techniques [15, 13]. Because every trace of L is parsed singularly, the cost is $O(|L|)$. Finally, denoting as m the maximum amount of parameters of a template, the cost is $O(|\Sigma|^m)$ because constraints are verified for every permutation of symbols in alphabet Σ . For standard Declare, e.g., $m = 2$.

5 Concluding remarks and research outlook

The discussion has been focused on events occurrence to make clearer the technique. However, it is equivalently applicable on the other attributes associated the event itself because the structure of any RCon $\Psi = \alpha \square \rightarrow \varphi$. Indeed the formulas are extensible to cover the multiple perspectives as seen in Deliverable D5.1.

In the present work, we leverage the process ontology developed in Deliverable D5.2, to be able to read a stream of linked event data as a process traces. Then the checking is computed outside of the linked world, to be relinked afterwards only once retrieved the result. The advantage of *SHACLare*, proposed in D5.2, was to be natively employing semantic technologies. It is indeed our next step to integrate this new streaming feature into it to have this advantage back.

6 Summary

In this deliverable, we presented a technique to monitor online linked data event streams with respect to declarative constraints. We extended the logic behind MP-Declare, proposed in D5.1, to deal with the streaming setting, exploiting the reactive nature of constraints. We derived the ability to interpret streams of linked data from D5.2 results.

Deliverable 5.3 – v1.0

Each constraint activation is independent from another one, even in within the same process case. This led to a fine granular measurement system more resistant to noise, both for the singles traces and the entire observed stream. We exploited automata theory to perform efficiently the verification of reactive constrains online, as soon as new events are recorded into the stream. We detailed the implementation of the technique with an algorithm.

References

- [1] A. Cecconi, C. Di Ciccio, G. De Giacomo, and J. Mendling, “Interestingness of Traces in Declarative Process Mining: The Janus LTLpf Approach,” in *Business Process Management*. Springer International Publishing, 2018, vol. 11080, pp. 121–138.
- [2] I. M. Hodkinson and M. Reynolds, “Separation-Past, Present, and Future.” in *We Will Show Them!(2)*, 2005, pp. 117–142.
- [3] M. Reynolds, “The complexity of temporal logic over the reals,” *Annals of Pure and Applied Logic*, vol. 161, no. 8, pp. 1063–1096, May 2010.
- [4] D. Gabbay, “The declarative past and imperative future,” in *Temporal logic in specification*. Springer, 1989, pp. 409–448.
- [5] M. Y. Vardi and P. Wolper, “An Automata-Theoretic Approach to Automatic Program Verification,” in *Proceedings of the First Symposium on Logic in Computer Science*. Cambridge: IEEE Computer Society, 1986, pp. 322–331.
- [6] Y. S. Ramakrishna, L. E. Moser, L. K. Dillon, P. M. Melliar-Smith, and G. Kutty, “An automata-theoretic decision procedure for propositional temporal logic with Since and Until,” *Fundam. Inform.*, vol. 17, no. 3, pp. 271–282, 1992.
- [7] P. Gastin and D. Oddoux, “LTL with past and two-way very-weak alternating automata,” in *MFCS*, vol. 3. Springer, 2003, pp. 439–448.
- [8] G. De Giacomo, R. De Masellis, and M. Montali, “Reasoning on LTL on Finite Traces: Insensitivity to Infiniteness.” in *AAAI*, 2014, pp. 1027–1033.
- [9] G. De Giacomo and M. Y. Vardi, “Linear temporal logic and linear dynamic logic on finite traces,” in *IJCAI’13 Proceedings of the Twenty-Third international joint conference on Artificial Intelligence*. Association for Computing Machinery, 2013, pp. 854–860.
- [10] C. Di Ciccio, F. M. Maggi, M. Montali, and J. Mendling, “Resolving inconsistencies and redundancies in declarative process models,” *Information Systems*, vol. 64, pp. 425–446, March 2017.
- [11] J. Adamo, *Data mining for association rules and sequential patterns - sequential and parallel algorithms*, J. Adamo, Ed. Springer New York, 2001.
- [12] F. M. Maggi, R. J. C. Bose, and W. M. van der Aalst, “Efficient discovery of understandable declarative process models from event logs,” in *International Conference on Advanced Information Systems Engineering*. Springer, 2012, pp. 270–285.
- [13] F. M. Maggi, C. Di Ciccio, C. Di Francescomarino, and T. Kala, “Parallel algorithms for the automated discovery of declarative process models,” *Information Systems*, 2018.
- [14] J. E. Hopcroft, R. Motwani, and J. D. Ullman, *Introduction to automata theory, languages, and computation*, 3rd ed. Boston: Pearson/Addison Wesley, 2007, oCLC: ocm69013079.
- [15] C. Di Ciccio and M. Mecella, “On the discovery of declarative control flows for artful processes,” *ACM Transactions on Management Information Systems (TMIS)*, vol. 5, no. 4, p. 24, 2015.