Complete and Interpretable Conformance Checking of Business Processes

Luciano García-Bañuelos, Nick R.T.P. van Beest, Marlon Dumas, Marcello La Rosa and Willem Mertens

Abstract—This article presents a method for checking the conformance between an event log capturing the actual execution of a business process, and a model capturing its expected or normative execution. Given a process model and an event log, the method returns a set of statements in natural language describing the behavior allowed by the model but not observed in the log and vice versa. The method relies on a unified representation of process models and event logs based on a well-known model of concurrency, namely event structures. Specifically, the problem of conformance checking is approached by converting the event log into an event structure, converting the process model into another event structure, and aligning the two event structures via an error-correcting synchronized product. Each difference detected in the synchronized product is then verbalized as a natural language statement. An empirical evaluation shows that the proposed method can handle real datasets and produces more concise and higher-level difference descriptions than state-of-the-art conformance checking methods. In a survey designed according to the technology acceptance model, practitioners showed a preference towards the proposed method with respect to a state-of-the-art baseline.

Index Terms—process mining, conformance checking, process model, event log, event structure.

1 Introduction

Process mining [1] is a family of methods concerned with the analysis of event logs produced by software systems that support the execution of business processes. Process mining methods allow analysts to understand how a business process is actually executed on top of the software system, and to detect and analyze deviations with respect to performance objectives or normative executions.

The main input of a process mining method is an *event log* of a business process. An event log is a set of *traces*, each consisting of the sequence of *event records* produced by one execution of the process (a.k.a. a *case*). An event denotes the start, end, abortion or other relevant state change of a task. As a minimum, an event record contains a timestamp, an identifier of the case to which the event refers, and an *event class*, that is, a reference to a task in the business process.

This article is concerned with a recurrent process mining operation, namely *business process conformance checking* [1]. Given an event log recording the actual executions of a business process, and given a process model capturing its intended or normative executions, the goal of conformance checking is to pinpoint and to describe differences between the behavior observed in the event log and the behavior captured in the process model.

Business process conformance checking is used in a variety of settings, including compliance auditing, model maintenance and automated process model discovery. In the context of compliance auditing, a typical task is to detect deviations in the process execution with respect to a normative model, that is, discovering behavior observed in the log that does not fit with what the model stipulates (i.e. unfitting behavior). In the context of process model maintenance, conformance checking allows one to identify both behavior observed in the log that does not fit with the model (unfitting behavior) as well as behavior observed in the model but not in the log (additional behavior). The former situation indicates that the model may need to be extended to capture the unfitting behavior, while the latter situation suggests that some execution paths in the process model have become spurious, meaning that these paths are no longer used and may be pruned if they are found to have lost relevance. Finally, conformance checking is used in the context of iterative automated discovery of process models. In this setting, an initial process model is discovered from the event logs using one of various existing automated process discovery algorithms. Conformance checking is then applied to assess the extent to which this process model captures all the behavior in the log as well as the amount of additional behavior allowed by the model but not observed in the log. Given the output of conformance checking, the discovered process model may be adjusted by restricting its behavior in order to remove paths in the model that are never observed (to avoid over-generalization) or conversely, by extending the model in order to capture behavior observed in the log but not captured in the model [2], [3]. In summary, conformance checking allows analysts to monitor the use of a software system by process workers in order to detect undesired deviations, to align the model so as to reflect reality, or to reverse-engineer an accurate model of the process that is being executed on top of the system. These tasks in turn support the continuous alignment of the intended process and the supporting software system [4].

Previous approaches to business process conformance checking are designed to identify the number and the location of the differences between the model and the traces

- *L. García-Bañuelos and M. Dumas are with the University of Tartu, Estonia.*
  E-mail: {luciano.garcia, marlon.dumas}@ut.ee
- *N.R.T.P. van Beest is with Data61, CSIRO, Brisbane, Australia.*
  E-mail: nick.vanbeest@data61.csiro.au
- *M. La Rosa and W. Mertens are with the Queensland University of Technology, Australia.*
  E-mail: {m.larosa, w.mertens}@qut.edu.au

in the log, rather than providing a diagnosis that would allow analysts to understand these differences. For example, these approaches can identify that there is a state in the process model where the model has additional behavior not observed in the log, but without describing this additional behavior. Similarly, these approaches can find points in a trace where the behavior observed in the trace deviates from the model, but without explaining what behavior the trace has that the model does not.

This article addresses these limitations by proposing a method for business process conformance checking that: (i) identifies all differences between the behavior in the model and the behavior in the log; and (ii) describes each difference via a natural language statement capturing task occurrences, behavioral relations or repeated behavior captured in the model but not observed in the log, or vice-versa.

The proposed method, namely *behavioral alignment*, is built on the idea of representing process models and event logs using a unified model of concurrent behavior, specifically *event structures* [5]. In other words, both the input process model and the event log are transformed into event structures. The two resulting event structures are then aligned via an error-correcting synchronized product that identifies all the behavioral differences between the process model and the event log. This synchronized product is used as a basis for enumerating the differences and verbalizing them as natural language statements. The choice of event structures is driven by the fact that they allow us to characterize the detected differences in terms of behavioral relations corresponding to well-accepted elementary workflow patterns [6], namely causality (sequence pattern), concurrency (parallel split & syncronization patterns), conflict (exclusive choice & merge patterns) and repetition (cycles).

As a running example, Fig 1 presents a model of a loan application process using the Business Process Model and Notation (BPMN). The process starts with the receipt of a loan application. Two tasks are performed in parallel – "Check credit history" and "Check income sources". Next the application is assessed, leading to two possible branches. In one branch, a credit offer is made to the customer and the process ends. In the other, a negative decision is communicated to the customer. In some cases, the customer is asked to provide additional information. Once the customer provides this information, the application is assessed again.

Consider now a log {ABCDEH, ACBDEH, ABDEH, ABCDFH, ACBDFH, ABDFH} where, for convenience, traces are represented as words over the single-letter labels A-H shown in the top-right corner of each model element in Fig 1. Given this log, the proposed method identifies the following differences: (i) task C is optional in the log; (ii) the cycle including IGDF is not observed in the log. The first statement characterizes the behavior observed in the log but not in the model, while the second characterizes the behavior captured in the model but not observed in the log.

The rest of the article is structured as follows. Section 2 discusses previous work on conformance checking. Section 3 introduces event structures so as to make the article self-contained. Next, Section 4 gives an overview of the proposed conformance checking method, which relies on the construction of event structures from process models and from event logs (Section 5), a partially synchronized

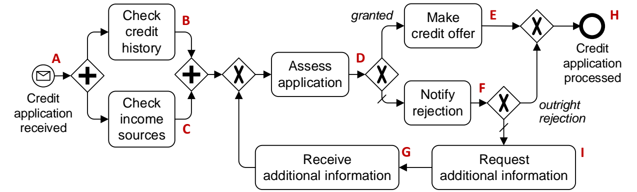

Fig. 1: Example: loan application process model

product of event structures (Section 6) and a method for extracting and verbalizing differences from the partially synchronized product (Section 7). Section 8 presents an empirical evaluation using both synthetic and real logs as well as a usability survey with process management researchers and professionals. Finally, Section 9 summarizes the contributions and outlines directions for future work.

2 Background and Related Work

The purpose of conformance checking is to identify two types of discrepancies:

1) Unfitting log behavior: behavior observed in the log that is not allowed by the model.
2) Additional model behavior: behavior allowed in the model but never observed in the log.

The identification of unfitting log behavior has been approached using two types of methods: *replay* and *trace alignment*.

Replay methods take as input one trace at a time and determine the maximal prefix of the trace (if any) that can be parsed by the model. When it is found that a prefix can no longer be parsed by the model, this "parsing error" is corrected either by skipping the event and continuing with the next one, or by changing the state of execution of the process model to one where the event in question can be replayed. A representative replay method is *token fitness* [7], which replays each trace in the log against the process model (represented as a Petri net) and identifies two types of errors: (i) *missing tokens*: how many times a token needs to be added to a place in the Petri net in order to correct a replay (parsing) error; and (ii) *remaining tokens*: how many tokens remain in the Petri net once a trace has been fully replayed. De Medeiros [8] proposed two extensions of this technique: *continuous parsing* and *improved continuous semantics fitness* (ICS). These extensions rely on the same principle, but sacrifice completeness of the output in order to gain in performance. Another extended replay method has been proposed by vanden Broucke et al. [9]. This method starts by decomposing the process model into single-entry single-exit (SESE) regions, so that the replay can be done on each region separately. This decomposition allows the replay to be performed independently in each region, thus making it more scalable and suitable for real-time conformance analysis. It also produces more localized feedback, since each replay error can be traced back to a specific region. An analysis of different approaches to decompose the model into regions for the purpose of measuring unfitting log behavior is provided in [10].

A general limitation of replay methods is that error recovery is performed locally each time that an error is

encountered. Hence, these methods might not identify the minimum number of errors that can explain the unfitting log behavior. This limitation is addressed by *trace alignment fitness* [11], [12]. This latter method identifies, for each trace in the log, the closest corresponding trace parsed by the model and computes an alignment showing the points of divergence between these two traces. The output is a set of pairs of "aligned traces". Each pair shows a trace in the log that does not match exactly a trace in the model, together with the corresponding closest trace(s) produced by the model. For example given the model shown in Fig. 1 and log {ABCDEH, ACBDEH, ABDEH, ABCDFH, ACBDFH, ABDFH}, trace alignment produces two aligned traces: one between trace ABDEH of the log and trace ABCDEH of the model; and another between trace ABDFH of the log and trace ABCDFH of the model. From this, the user can infer that the unfitting log behavior is that task C is optional in the log but always executed in the model.

The number of aligned traces produced by the above method is often too large to be explored exhaustively. Visualizations are proposed to cope with large sets of aligned traces. Fundamentally though, the limitation of trace alignment fitness (also shared with replay methods) is that it identifies differences at the level of individual traces rather than at the level of behavioral relations observed in the log but not captured in the model. This observation is one of the starting points for the method proposed in this article. Instead of aligning traces produced by the model with traces produced by the log, our idea is to compute an optimal alignment between an event structure representing the entire behavior of the model and an event structure representing all the behavior observed in the log.

Methods to identify additional behavior include those based on *negative events* and those based on *prefix automata*. An exemplar of the former class is *negative event precision* [13]. This method works by inserting inexistent (so-called *negative*) events to enhance the traces in the log. A negative event is inserted after a given prefix of a trace if this event is never observed preceded by that prefix anywhere in the log. For example, if event c is never observed after prefix ab, then c can be inserted as a negative event after ab. The traces extended with negative events are then replayed on the model. If the model can parse some of the negative events, it means that the model has additional behavior. This approach to detect additional behavior is however heuristic: it does not guarantee that all additional behavior is identified. An extension of this method [14] addresses its scalability limitations and can also better deal with noisy logs, but again it does not guarantee that all additional behavior is identified.

A method to detect the presence of additional model behavior based on prefix automata is outlined in [15]. The first step in this method is to generate a prefix automaton that fully represents the entire log. Each state in this automaton corresponds to a unique trace prefix. For each state S_a in the automaton, the corresponding trace prefix is replayed in the model in order to identify a matching state S_m in the model. The set of tasks enabled in S_m is then determined. If there is a task enabled in state S_m in the model but not in state S_a in the automaton, this is marked as additional model behavior by adding a so-called "escaping edge" to state S_a of the automaton. This edge is labelled with the task in question and considered as a sink state in the automaton. The edge represents the fact that in state S_a, there is additional behavior in the model that is not observed in the log. This basic method suffers from two limitations: (i) it cannot handle tasks with duplicate labels in the model nor tasks without labels (so-called *invisible tasks*), which are needed to capture decisions based on the evaluation of data conditions; and (ii) it assumes that all traces in the log fit the model. These limitations are addressed in [16]. The idea of this latter method is to first calculate an alignment between the traces in the log and traces in the model, using the trace alignment technique mentioned above. This leads to a log with aligned traces, which include invisible tasks. These aligned traces and any prefix thereof can always be replayed by the model. The prefix automaton is then computed from the model-projection of the aligned traces rather than from the original traces. The automaton is then used to detect "escaping edges" in the same way as described above.

The methods described in [15] and [16] are able to pinpoint states in the model where behavior is allowed that is not present in the log. However, they cannot characterize the additional allowed behavior, beyond stating that the additional behavior starts with the execution of a given task. For example, given the model in Fig. 1 and log {ABCDEH, ACBDEH, ABDEH, ABCDFH, ACBDFH, ABDFH}, these methods identify an escaping edge after a prefix in the automaton that finishes with "Notify Rejection". However, they do not detect that there is repetitive behavior in the model whereas there is no such repetitive behavior in the log (e.g. in the model task "Assess application" can be repeated whereas this repetition is not observed in the log).

To recap, existing conformance checking methods can be characterized along the following dimensions:

- *Scope*: whether the method detects unfitting (*Unfitting*) behavior, additional (*Additional*) behavior or both (*Both*).
- *Completeness (Comp.)*: whether the method detects a complete (*Complete*) or a partial (*Partial*) set of differences within its scope.
- *Unit of feedback*: elementary unit of output of the method; in other words, how does the method characterize one difference between the model and the log? The unit of feedback may be a missing or additional token in a place (*Tokens*), a pair of (aligned) traces (*2 × Traces*), a negative event occurrence (*Neg. events*), an escaping edge (*Escap. edges*) or a behavioral relation between a pair of tasks (*Behavioral rel.*).
- *Modality of feedback*: whether the method produces visual feedback over the model (*VisualM*), visual feedback over the traces (*VisualT*), or textual feedback (*Text*).

Table 1 characterizes the overviewed methods (including the behavioral alignment method herein proposed) in terms of the above dimensions. We observe that the method herein proposed is the only one that covers both unfitting and additional behavior (under the same conceptual umbrella). It provides complete feedback and it is also the only one to provide feedback in terms of behavioral relations. One could combine a (complete) technique for unfitting behavior

and another one for additional behavior to obtain a hybrid method that is comparable to ours except for the feedback unit and modality. Accordingly, in the experimental evaluation, we compare our method with a combination of the trace alignment and prefix automata methods. We opt for trace alignment (as opposed to token fitness) because it returns more detailed feedback. Indeed, whereas token fitness only tells us the places in the model where there is a parsing error for a given unfitting trace, trace alignment gives us for each unfitting trace in the log: (i) one or all of the closest corresponding traces in the model; and (ii) the locations of the differences between the trace in the log and the closest corresponding trace(s) in the model. On the other hand, we opt for prefix automata because of its higher accuracy compared to negative event precision.

Method \ Dimension	Scope	Comp.	Feedback unit	Feedback mod.
Token fitness [7]	Unfitting	Complete	Tokens	VisualM
Continuous parsing [8]	Unfitting	Partial	Tokens	VisualM
ICS [8]	Unfitting	Partial	Tokens	VisualM
SESE replay [10]	Unfitting	Complete	Tokens	VisualM
Alignments [11]	Unfitting	Complete[1]	2×Traces	VisualT VisualM
Neg. event prec. [13]	Additional	Partial	Neg. events	-
Prefix automata [16]	Additional	Partial	Escap. edges	VisualM
Behavioral alignment	Both	Complete	Behavioral rel.	Text

TABLE 1: Comparison of conformance checking methods

The method proposed in this paper relies on the construction of an event structure from the model, an event structure from the log, and the computation of a synchronized product between these two event structures, from which a set of differences are extracted and verbalized. Some of these ideas are inspired from our previous work on model-to-model comparison [17] and model-to-log comparison [18]. In our work on model-to-model comparison, we developed an unfolding algorithm to create an event structure from a process model and an algorithm to calculate a synchronized product from two event structures representing the behavior of two process models. This latter algorithm however differs fundamentally from the algorithm introduced in the present article because the algorithm in [17] is not designed to produce complete feedback; specifically, it cannot enumerate all differences when the model contains repeating behavior, instead it is only able to assert that a given task can occur multiple times in a model and at most once in the other. In particular, the technique in [17] cannot identify the start and the end of the repeated behavior, while the one presented in this paper can. Meanwhile, in [18] we presented an algorithm to calculate a synchronized product of two event structures, which is complete (it detects all differences) but it operates over event structures representing finite behavior (that is extracted from a log), while in the present article the event structure extracted from the model can represent infinite (cyclic) behavior. In summary, the synchronized product proposed in this article is both complete and can handle cyclic behavior in the event structure produced from the process model. Additionally, the difference extraction and verbalization algorithm is finer-grained than the ones presented in [17] and [18]. In particular, the algorithm presented here detects situations where a task in the model is observed in the log but in a different state relative to the model, as well as situations where a task in the model is substituted by a task with a different label in the log.

In a previous short paper [19], we briefly sketched the idea of using a synchronized product of event structures for conformance checking. The present article elaborates this idea by providing formal definitions and algorithms both for the computation of the product and for the extraction and verbalization of differences. The article also adds an empirical evaluation and a comparison against existing conformance checking techniques.

3 Event Structures

A Prime Event Structure (PES) [5] is a graph of events, where an event e represents the occurrence of an action (e.g. a task) in the modeled system (e.g. a business process). If a task occurs multiple times in an execution, each occurrence is represented by a different event. The order of occurrence of events is defined via three binary relations: i) *Causality* ($e < e'$) indicates that event e is a prerequisite for e'; ii) *Conflict* ($e \# e'$) implies that e and e' cannot occur in the same execution; iii) *Concurrency* ($e \parallel e'$) indicates that no order can be established between e and e'. Formally:

Definition 1. A *Labeled Prime Event Structure* is the tuple $\mathcal{E} = (E, \leq, \#, \lambda)$ where E is the set of events, $\leq \subseteq E \times E$ is a partial order, referred to as *causality*, $\# \subseteq E \times E$ is an irreflexive, symmetric relation, referred to as *conflict*, and $\lambda : E \rightarrow \mathcal{L} \cup \{\tau\}$ is a labeling function.
The irreflexive version of causality is denoted as $<$. The *concurrency relation*, in turn, can be derived from causality and conflict relations, i.e. $\parallel \triangleq E \times E \setminus (< \cup <^{-1} \cup \#)$. Moreover, conflict is "inherited" via the causality relation, i.e. $e \# e' \wedge e' \leq e'' \Rightarrow e \# e''$ for $e, e', e'' \in E$. The relations $<$, $\leq$, $\#$ and $\parallel$ are together referred to as *behavioral relations*.

Fig. 2 presents a variant of the process model introduced in Fig. 1[2] and its corresponding PES $\mathcal{E}^1$. In the PES, nodes are labelled by an event identifier and a task label, e.g. "e_2:C" tells us that event e_2 represents an occurrence of task "C". For brevity, we will often omit the event label. Causal dependencies are drawn as solid arcs, whereas conflict relations as dotted edges. In order to simplify the graphical representation of an event structure, transitive causal and hereditary conflict relations are not drawn. Every two events that appear neither directly nor transitively connected are considered to be concurrent.

An execution context (i.e. a "state") in an event structure is described in terms of sets of events that can occur together in an execution of the underlying system. Such a set of events is called a *configuration*. Formally, we say that a set of events $C \subseteq E$ is a configuration iff (i) C is causally closed: for each event $e \in C$, the configuration C also contains all causal predecessors of e, i.e.

1. There are two versions of the trace alignment technique: "one alignment" and "all optimal alignments". The former is more scalable but produces only partial results (it may miss some differences) while the latter produces complete results.

2. In this variant of the process, task C can be skipped – e.g. "Check income sources" may not be required for existing customers – and applicants cannot request for reviewing a rejected application.

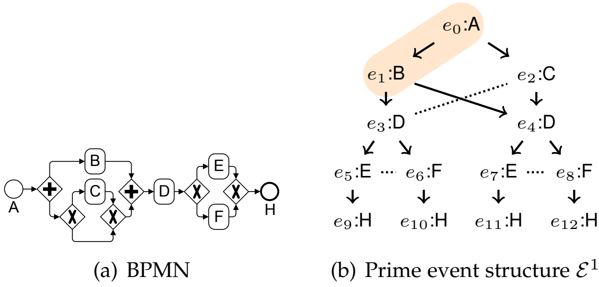

(a) BPMN (b) Prime event structure $\mathcal{E}^1$

Fig. 2: Sample process model and event structure

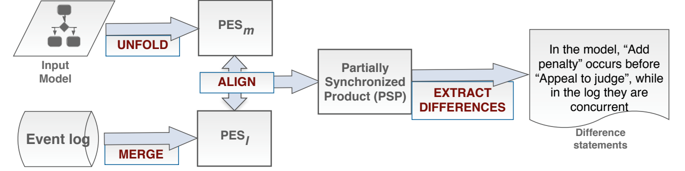

Fig. 3: Overall view of the behavioral alignment method

$\forall e' \in E, e \in C : e' \leq e \Rightarrow e' \in C$, and (ii) C is conflict free: C does not contain any pair of events in mutual conflict, i.e. $\forall e, e' \in C \Rightarrow \neg(e \# e')$. An event e is an extension of a configuration C, denoted $C \oplus e$, if and only if $C \cup \{e\}$ is also a configuration. We distinguish between observable and silent (or τ) events. In the following, we write $C|_\lambda$ to denote the restriction of configuration C to its subset of observable events, i.e. $C|_\lambda \triangleq \{e \in C \mid \lambda(e) \neq \tau\}$. We denote by $\mathcal{F}(\mathcal{E})$ the set of all the configurations of $\mathcal{E}$ and by $\mathcal{F}_{\mathrm{m}}(\mathcal{E})$ the set of configurations that are maximal with respect to set inclusion. Moreover, we define the concept of *set of possible extensions of a configuration* C as $\mathrm{PE}(C) \triangleq \{e \mid C \oplus e\}$.

For example, let C_1 be the set of events $\{e_0, e_1\}$ – highlighted in Fig. 2(b). Intuitively, the configuration C_1 of $\mathcal{E}^1$ represents the state of computation in which tasks "A" and "B" have occurred. Moreover, given the configuration C_1 we say that the computation can evolve by executing an event from the set $\{e_2, e_3\}$, given that this set of events corresponds to the possible extensions of C_1, that is $\mathrm{PE}(C_1)$. Now, if we consider the occurrence event e_3, we would have to consider a new configuration, say $C_2 = \{e_0, e_1, e_3\}$, which we can also denote as $C_1 \oplus e_3$. Note that in the context of configuration C_2 the occurrence of e_2 is no longer possible because the event e_3 is in conflict with e_2. Finally, in this same example the set of maximal configurations is $\mathcal{F}_{\mathrm{m}}(\mathcal{E}^1) = \{\{e_0, e_1, e_3, e_5, e_9\}, \{e_0, e_1, e_3, e_6, e_{10}\}, \{e_0, e_1, e_2, e_4, e_7, e_{11}\}, \{e_0, e_1, e_2, e_4, e_8, e_{12}\}\}$.

We use the term *local configuration* of an event e to refer to $\lceil e \rceil \triangleq \{e' \mid e' \leq e\}$, and the term *strict causes* of an event to refer to $\lceil e) \triangleq \lceil e \rceil \setminus \{e\}$. Finally, we say that events e_1 and e_2 are in *immediate conflict*, denoted $e_1 \#_\mu e_2$, if and only if $e_1 \# e_2$ and they are both possible extensions of the same configurations. Formally, the latter property can be verified by checking if $\lceil e_1) \cup \lceil e_2 \rceil$ and $\lceil e_1 \rceil \cup \lceil e_2)$ are both configurations or not.

4 Overview of Behavioral Alignment

The proposed behavioral alignment method takes as input a process model captured in the standard BPMN language and an event log (cf. Fig. 3). In order to leverage Petri net-based techniques for constructing event structures, the input process model is first converted into a Petri net using the transformation proposed in [20]. The resulting Petri net is then unfolded into a prime event structure (cf. PES$_m$ in Fig. 3) using Petri net unfolding techniques [21]. Each event in the resulting PES corresponds to an occurrence of a task in the process model. The procedure for constructing an event structure from a Petri net is outlined in Section 5.1.

Meanwhile, the input event log is transformed into another prime event structure (cf. PES$_l$ in Fig. 3) by first transforming the set of traces in the log into a set *partially-ordered runs* and then "prefix-merging" the resulting set of runs. The procedure for constructing an event structure from a log is elaborated upon in Section 5.2.

Given the prime event structures PES$_m$ and PES$_l$ obtained from the model and the log respectively, we compute a so-called Partially Synchronized Product (PSP) of the two event structures. In a nutshell, a PSP is a representation of an error-correcting synchronized traversal of two input PESs, such that when a discrepancy between the PESs is detected, it is explicitly recorded and the traversal resumes from a"suitable" configuration in each of the two PESs. The procedure for calculating the PSP is presented in Section 6.

If two event structures PES$_m$ and PES$_l$ have a behavioral difference of type unfitting log behavior, this difference will be captured in a node of the PSP. Thus, we can enumerate all unfitting log behavior by traversing their PSP. To expose additional model behavior, we define a notion of coverage of a PES extracted from a model by a PES extracted from a log. The parts of PES$_m$ not covered by PES$_l$ can then be isolated and enumerated.

Differences between two event structures can be of several types. For example, one type of difference is that a task t is always executed according the model, but it is skipped in some traces in the log. Another type of difference occurs when two tasks t_1 and t_2 are causally related in the model (e.g. t_1 occurs always before t_2), but the corresponding events appear in any order in the log (i.e. sometimes t_1 occurs before t_2, sometimes the other way around). In order to generate an interpretable difference diagnosis from the PSP, we define a set of disjoint and complete mismatch patterns, as well as rules to verbalize each mismatch pattern as a natural language statement. The patterns and their verbalization are presented in Section 7.

5 Construction of Event Structures

This section shows how event structures are derived from a Petri net and from an event log.

5.1 From Petri nets to PES

A Petri net is a bipartite graph, consisting of transitions (rectangles), places (hollow circles), tokens (filled circles) and arcs. A transition represents a system action (e.g. a task). Each transition has a set of input places and a set of output places. At a given point in the execution of a Petri net, a place can hold a number of tokens. The distribution of tokens across places on the net is called a *net marking*. A transition is enabled and can "fire" when all its input

places have at least one token. When a transition fires, one token is removed from each input place and one token is put into each output place. A Petri net is called *safe* iff in every possible marking each place holds at most one token.

Definition 2. A tuple (P, T, F, λ) is a *labeled Petri net*, where P is a set of *places*, T is a set of *transitions*, with $P \cap T = \emptyset$, $F \subseteq (P \times T) \cup (T \times P)$ is a set of arcs, and $\lambda : P \cup T \to \mathcal{L} \cup \{\tau\}$ a labeling function. A *net marking* $M : P \to \mathbb{N}_0$ is a function that associates a place $p \in P$ with a natural number (viz., place tokens). A *net system* $N = (P, T, F, M_0)$ is a Petri net (P, T, F) together with an *initial marking* M_0.

Places and transitions are conjointly referred to as *nodes*. We write $\bullet y = \{x \in P \cup T \mid (x, y) \in F\}$ and $y\bullet = \{z \in P \cup T \mid (y, z) \in F\}$ to denote the *preset* and *postset* of node y, respectively. F^+ and F^* denote the irreflexive and reflexive transitive closure of F, respectively.

The dynamics of a net system can be expressed in terms of markings. A marking M *enables* a transition t if $\forall p \in \bullet t : M(p) > 0$. Moreover, the firing of t leads to a new marking M', with $M'(p) = M(p) - 1$ if $p \in \bullet t \setminus t\bullet$, $M'(p) = M(p) + 1$ if $p \in t\bullet \setminus \bullet t$, and $M'(p) = M(p)$ otherwise. We also use $M \xrightarrow{t} M'$ to denote the firing of t. The marking M_n is said to be reachable from M if there exists a sequence of transition firings $\sigma = t_1 t_2 \dots t_n$ such that $M \xrightarrow{t_1} M_1 \xrightarrow{t_2} \dots \xrightarrow{t_n} M_n$. A marking M of a net is *n-safe* if $M(p) \leq n$ for every place p. A net system N is said *n-safe* if all its reachable markings are *n-safe*. In the following we restrict ourselves to *1-safe* net systems. Hence, we identify the marking M with the set $\{p \in P \mid M(p) = 1\}$.

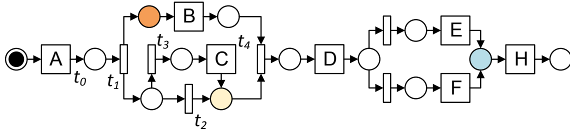

Fig. 4: Petri net for BPMN model in 2(a)

Fig. 4 presents a Petri net corresponding to the BPMN process model in Fig. 2(a). Each task, start and end event is mapped into a transition, which carries the same label as the corresponding BPMN construct. The Petri net additionally contains some unlabeled transitions. These transitions correspond to parallel gateways in the BPMN process model as well as branches stemming out of decision gateways. The materialization of gateways and decision branches as unlabelled (a.k.a. silent or τ) transitions is an artifact of the transformation from BPMN to Petri nets [20]. These unlabeled transitions will be eliminated during the construction of the event structure as discussed later.

An alternative approach to represent the dynamics of a net system is by means of another Petri net that explicitly represents all partially-ordered runs of the original net system. A run of a system is a partially-ordered set of events that can occur in one execution thereof. All the partially-ordered runs can be accommodated in a single tree-like structure, called *branching process* [5]. Fig. 5 presents the branching process of the net system in Fig. 4.

Branching processes are intimately related with prime event structures because they explicitly represent the same set of behavioral relations. This fact is formally captured in the following definition.

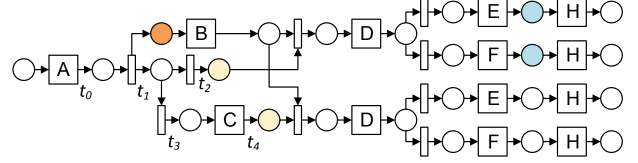

Fig. 5: Branching process of the marked net in Fig. 4

Definition 3. Let $N = (P, T, F)$ be a net and $x, y \in P \cup T$ two nodes in N.

- x and y are *causal*, written $x <_N y$, iff $(x, y) \in F^+$,
- x and y are in *conflict*, denoted $x \#_N y$, iff $\exists\, t, t' \in T : t \neq t' \wedge \bullet t \cap \bullet t' \neq \emptyset \wedge (t, x), (t', y) \in F^*$,
- x and y are *concurrent*, denoted $x \parallel_N y$, iff neither $x <_N y$, nor $y <_N x$, nor $x \#_N y$.

Armed with the above, we can now provide a formal definition of a branching process.

Definition 4. Let $N = (P, T, F, M_0)$ be a net system. The *branching process* $\beta = (B, E, G, \rho)$ of N is the net (B, E, G) defined by the inductive rules in Figure 6. The rules also define the function $\rho : B \cup E \to P \cup T$ that maps each node in β to a node in N. Given the set $X \subseteq B \cup E$, $\varrho(X)$ is a shorthand for $\{\rho(x) \mid x \in X\}$.

$$\frac{p \in M_0}{b = \langle \emptyset, p \rangle \in B \qquad \rho(b) = p}$$

$$\frac{t \in T \quad B' \subseteq B \quad B'^2 \subseteq \parallel_\beta \quad \varrho(B') = \bullet t}{e = \langle B', t \rangle \in E \qquad \rho(e) = t}$$

$$\frac{e = \langle B', t \rangle \in E \quad t\bullet = \{p_1, \dots, p_n\}}{b_i = \langle t', p_i \rangle \in B \qquad \rho(b_i) = p_i}$$

Fig. 6: Inductive construction of a branching process

The elements of B and E in a branching process β are respectively called *conditions* and *events*. G in turn denotes the flow relation of the branching process. $Min(\beta)$ denotes the set of minimal elements of $B \cup E$ with respect to the transitive closure of G and, hence, $Min(\beta)$ corresponds to the set of places in the initial marking of N, i.e., $\varrho(Min(\beta)) = M_0$. Labels on a net N can be carried over to its branching process β by composing λ and ρ, i.e., $\lambda_\beta \triangleq \lambda_N \circ \rho$.

Clearly, the behavioral relations derived from the branching process generate a prime event structure [5]. Events in the branching process correspond to events in the event structure. Consequently, the notion of configuration can be extrapolated from prime event structures to branching processes. Specifically, given a branching process $\beta = (B_\beta, E_\beta, G_\beta, \lambda_\beta)$ of a marked net N, the event structure $\mathcal{E}$ of N is defined as $\mathcal{E}(N) \triangleq (E_\beta, \leq_\beta \cap E_\beta^2, \#_\beta \cap E_\beta^2, \lambda_\beta|_{E_\beta})$.[3] The latter definition maps both observable and silent transitions to events in the event structure. In [17], the authors

3. We use A^2 to denote the cartesian product of a set A.

proved that silent events can be abstracted away (i.e. removed) in a behavior-preserving manner, under a well-known notion of behavioral equivalence, namely visible-pomset equivalence.[4] The PES presented in Fig. 2(b) is the one that corresponds to the branching process in Fig. 5 after removing all silent events. Accordingly, in the rest of the paper we assume, without loss of generality, that the event structures we manipulate do not have silent transitions.

The branching process of a Petri net with cycles may be infinite. In [21], McMillan showed that for safe nets a prefix of a branching process fully encodes the behavior of the original net. Such prefix of a branching process is referred to as the *complete prefix unfolding* of a net. We will use β_m to denote a maximal branching process and β_f to denote a complete prefix unfolding of β_m.

Definition 5. Let $\beta_m = (B_m, E_m, G_m, \rho_m)$ be the maximal, possibly infinite branching process of the net system N.

- A *local configuration* $\lceil e \rceil$ of an event e in a branching process is the set of events that causally precede e, i.e. $\lceil e \rceil = \{e' \in E_m \mid (e', e) \in G_m^*\}$.
- The *reachable marking* of a local configuration, denoted $Mark(\lceil e \rceil)$, is the set of places in N that get marked after all the transitions in $\varrho(\lceil e \rceil)$ fire.
- An *adequate order* $\lhd$ is a strict well-founded partial order on local configurations, such that $\lceil e \rceil \subset \lceil e' \rceil$ implies $\lceil e \rceil \lhd \lceil e' \rceil$[5].
- An event e of a branching process is a *cutoff event* if there exists a *corresponding event* e', such that $Mark(\lceil e \rceil) = Mark(\lceil e' \rceil)$ and $\lceil e' \rceil \lhd \lceil e \rceil$. The pair cutoff/corresponding events (e, e') is referred to as a *cc-pair*.
- Let $E_f \subseteq E_m$ be the set of events of β_m such that $e \in E_f$ iff no event $e' <_{\beta_m} e$ is a cut-off event. A *complete prefix unfolding* β_f is the subnet of β_m having E_f as set of events.

The complete prefix unfolding for the net system in Fig. 4 (and its branching process in Fig. 5) is given in Fig. 7.

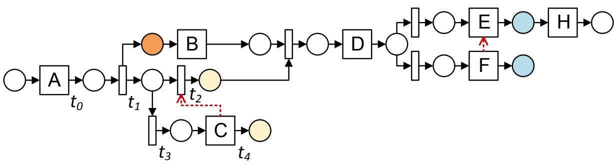

Fig. 7: Complete prefix unfolding of the marked net in Fig. 4

To illustrate the intuition behind a complete prefix unfolding, let us consider the local configurations $\lceil t_2 \rceil = \{t_0{:}\mathsf{A}, t_1{:}\tau, t_2{:}\tau\}$ and $\lceil t_4 \rceil = \{t_0{:}\mathsf{A}, t_1{:}\tau, t_3{:}\tau, t_4{:}\mathsf{C}\}$. To ease the tracking of events in the aforementioned local configuration from the unfolding in Fig. 7 back to events in the branching process and to transitions in the net system in Figures 4 and 5, we added labels t_0, t_1, etc., next to the corresponding graphical element (transition or event). Clearly, the "future" of event t_2, denoted $\lceil t_2 \rceil\Uparrow$, is isomorphic to that of t_4. Indeed, the firing of the transitions that correspond with the events in $\lceil t_2 \rceil$ would lead to a marking where places colored orange and yellow in Fig. 4 would hold a token each, which would be the same making that would produce the firing of the transitions that correspond with set of events in $\lceil t_4 \rceil$. Therefore, we can safely stop unfolding the branching process once we reach $t_4{:}\mathsf{C}$ provided that we continue unfolding from $t_2{:}\tau$ and onwards. Following Def. 5, the pair (t_4, t_2) is called a *cc-pair*. Moreover, the isomorphism on the future of the events in a cc-pair (e, f), that is $\lceil e \rceil\Uparrow$ and $\lceil f \rceil\Uparrow$, will be denoted as $\mathcal{I}_{(\lceil e \rceil, \lceil f \rceil)}$. In the graphical representation of unfoldings and PESs (as introduced later), a cc-pair is indicated via a dashed red arrow from the cutoff event to the corresponding event (cf. for example the dashed red arc between t_4 and t_2 in Fig. 7).

To represent the behavior specified by a BPMN process model, we will use the prime event structure derived from the complete prefix unfolding of the model's Petri net. The latter is herein called the *PES prefix unfolding* of a model and is formally defined as follows:

Definition 6. Let $\beta_f = (B_f, E_f, G_f, \rho)$ be the complete prefix unfolding of the net system N, with labelling function λ_β. Let $\overline{E_f} \subseteq E_f$ be the set of events of β_f such that $e \in \overline{E_f}$ iff e is labelled (i.e. $\lambda_\beta(e) \neq \tau$), or e is cutoff or corresponding event. The *PES prefix unfolding* of N, denoted $\overline{\mathcal{E}}(N)$, is defined as:

$$\overline{\mathcal{E}}(N) \triangleq (\overline{E_f}, \leq_\beta \cap \overline{E_f}^2, \#_\beta \cap \overline{E_f}^2, \lambda_\beta|_{\overline{E_f}})$$

From the previous definition we can see that the computation of a PES prefix unfolding is the same as that for a regular prime event structure except for the following: (i) we keep track of cc-pairs, and (ii) for convenience, we do not abstract away a silent event when such event is either a cutoff or a corresponding event. For example, Fig. 8 presents the PES prefix unfolding corresponding to the marked net in Fig. 4 and, hence, with the process model in Fig. 2(a).

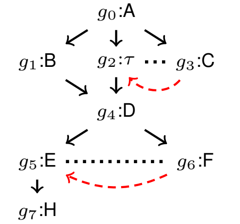

Fig. 8: PES prefix unfolding $\overline{\mathcal{E}}_2$ of the net in Fig. 4

Reasoning about possible executions of a PES prefix unfolding is not convenient because some configurations are not explicitly represented. To make it more convenient to explore the configurations of a PES prefix unfolding, we adapt to our setting the "shift" operation on net unfoldings introduced in [23]. Intuitively, given a cc-pair (e, f), since the futures of $\lceil e \rceil$ and $\lceil f \rceil$ are isomorphic, we can "shift" from one configuration to the other. The shift operation is thus a "step" function that allows us to move from one configuration to another. This intuition is captured below.

Definition 7. Let (e, f) be a cc-pair of the PES prefix $\overline{\mathcal{E}}$ and $\mathcal{I}_{(\lceil e \rceil, \lceil f \rceil)}$ the isomorphism from $\lceil e \rceil\Uparrow$ to $\lceil f \rceil\Uparrow$. Moreover, let C be a configuration of $\overline{\mathcal{E}}$. The (e, f)-shift of C, denoted $\mathcal{S}_{(e,f)}(C)$, is defined as follows:

$$\mathcal{S}_{(e,f)}(C) = \lceil f \rceil \cup \mathcal{I}_{(\lceil e \rceil, \lceil f \rceil)}(C \setminus \lceil e \rceil)$$

We say that $\mathcal{S}_{(e,f)}(C)$ is a *backward shift* iff $\lceil f \rceil \subset \lceil e \rceil$, that is, the corresponding event f is included in the local

4. This result holds on condition that every sink event in the event structure is a labeled (non-silent) event – something that we can easily ensure by adding, when needed, a "fake" labelled final event to the Petri net from which the event structure is generated.

5. Several definitions of *adequate order* exist; we use the one defined in [22], because it has been shown to generate compact unfoldings.

configuration of the cutoff event e, otherwise $\mathcal{S}_{(e,f)}(C)$ is called a *forward shift*. Moreover, an event e is said a *backward cutoff event* iff it entails backward shift. Intuitively, a backward shift "moves back" to a configuration that has already be observed in the past of the run.

With abuse of notation, we will use the following variant:

$$\overline{\mathcal{S}}_e(C) = \begin{cases} C & \text{if } e \text{ is not cutoff event} \\ \mathcal{S}_{(e,\mathrm{corr}(e))}(C) & \text{otherwise} \end{cases}$$

Consider for example configuration $C_1 = \{g_0\text{:A}, g_1\text{:B}, g_3\text{:C}\}$ of event structure $\overline{\mathcal{E}}_2$ in Fig. 8. C_1 contains the cutoff event g_3, which is associated with the cc-pair (g_3, g_2). Given that $\mathcal{S}_{(g_3,g_2)}(C_1) = \{g_0\text{:A}, g_1\text{:B}, g_2\text{:}\tau\}$, we infer that g_2:D is a possible extension of C_1.

Esparza [23] shows that any property that holds over a branching process (and thus on a maximal PES) also holds on its prefix unfolding by applying a sequence of shift operators. In other words, if we wish to compare a maximal PES with a PES prefix unfolding, we can apply shift operations on the PES prefix in order to materialize behavior that is not explicitly represented. This latter observation is used later when simultaneously traversing a PES prefix derived from a process model and a maximal PES derived from a log.

The extraction of a complete prefix unfolding from a Petri net (and the size of the prefix unfolding itself) is exponential on the size of the net [23]. This entails in turn that the derivation of the event structure from an input BPMN model is worst-case exponential.

5.2 From log to PES

As stated in Section 1, an event log consists of a set of traces such that each trace records one execution of a process. A trace is a totally ordered sequence of events. Each event corresponds to a transition in the lifecycle of a task (e.g. a task became enabled, the execution of a task started or completed). Event logs are formally defined as follows:

Definition 8. Let $\mathcal{L}$ be a set of task labels and let E be a universe of possible event occurrences, each denoting a transition in the lifecycle of some task. Let $\lambda : E \to \mathcal{L}$ be a labeling function that maps every event occurrence to a task label. A *trace* σ of length n is a function that maps each $i \in [0, n-1]$ to an event (occurrence) in $E_\sigma \subseteq E$. For convenience, we will refer to a trace and its elements as follows: $\sigma = \langle \lambda(e_0), \lambda(e_1), \ldots, \lambda(e_{n-1}) \rangle$. An event log L is a set of traces, i.e. $L \in \mathcal{P}(E^*)$.

In previous work [18], we presented a method to generate a PES from an event log. The method consists of two steps. First the event log, seen as a set of traces, is transformed into a set of partially-ordered runs by invoking a *concurrency oracle*. A concurrency oracle is a function that given a log, returns a set of pairs of event labels that are in a concurrency relation. Given a concurrency oracle, each trace is turned into a run by relaxing the total order induced by the trace into a partial order such that two events are not causally related if the concurrency oracle has determined that they occur concurrently. Several approaches have been proposed to extract concurrency relations between pairs of events from an event log [24], [25], [26]. Here we use the $\alpha+$ concurrency oracle [26] which is a refinement of the well-known α concurrency oracle [25], but other approaches can be similarly used. Two event labels A and B appearing in an event log are $\alpha-$concurrent if A is sometimes observed immediately after B and vice-versa. The latter intuition is formalized as follows [25]:

Definition 9. Let L be an event log over the set of event labels $\mathcal{L}$ and $\sigma \in L$ be a log trace. A pair of tasks with labels $A, B \in \mathcal{L}$ are said to be in *alpha directly precedes relation*, written $A \prec_{\alpha(L)} B$, if there exists a trace $\sigma = \langle \lambda(e_0), \lambda(e_1), \ldots, \lambda(e_{n-1}) \rangle$ in L, s.t. $A = \lambda(e_i)$ and $B = \lambda(e_{i+1})$. A pair of tasks $A, B \in \mathcal{L}$ are *alpha concurrent*, written $A \parallel_{\alpha(L)} B$, if $A \prec_{\alpha(L)} B \wedge B \prec_{\alpha(L)} A$.

The notion of $\alpha-$concurrency is inaccurate when the process model that generated the log contains loops involving one or two tasks. For instance, an event trace $\langle A, B, C, B, C, B, D \rangle$ might be generated by a process model that contains a loop involving tasks B and C. An α oracle would assert that $B \parallel_{\alpha(L)} C$, which is false. The $\alpha+$ concurrency approach extends the basic α oracle such that (1) short loops involving one task are identified and removed from the event log in a preprocessing step; and (2) concurrency detected on pairs of tasks that are involved in short loops are rectified in a post-processing step.

Once the concurrency relation is computed, the set of runs are merged into an event structure in a lossless manner, meaning that the set of maximal configurations of the resulting event structure is exactly equal to the set of runs. In this way and modulo the accuracy of the concurrency oracle, we ensure that the resulting event structure is a lossless representation of the input log.

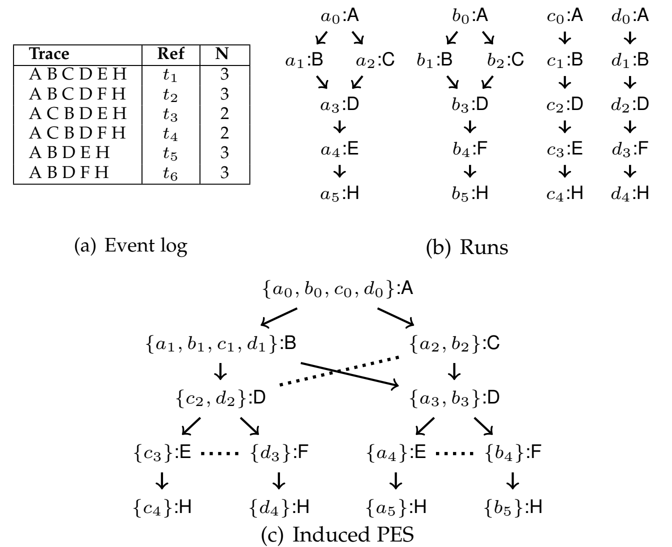

Trace	Ref	N
A B C D E H	t_1	3
A B C D F H	t_2	3
A C B D E H	t_3	2
A C B D F H	t_4	2
A B D E H	t_5	3
A B D F H	t_6	3

(a) Event log (b) Runs (c) Induced PES

Fig. 9: Example of construction of a PES from a set of traces

For example, consider the log given in Fig. 9(a), where the last column lists the number of instances of each distinct trace. Using the α oracle we conclude that event classes B and C are concurrent. Thus, we construct the set of runs in Fig. 9(b), where the notation e:A indicates that event e represents an occurrence of event class A in the original log. By merging events with the same label and the same history (i.e. same prefix), we obtain the PES in Fig. 9(c), where the

notation $\{e_1, e_2 \dots e_i\}$:A indicates that events $\{e_1, e_2 \dots e_i\}$ represent occurrences of event class A in different runs.

The algorithms to transform traces into runs (given a concurrency oracle) and to merge runs into event structures, are illustrated in [18]. This latter paper also shows that the complexity of this transformation is $O(|\sigma_m|^3)$, where $|\sigma_m|$ is the length of the longest trace in the event log.

6 Partially Synchronized Product

The *Partially Synchronized Product* (PSP) of two event structures [17] is a state machine in which the states correspond to pairs of configurations visited during an error-correcting synchronized traversal of the two input event structures, starting from their empty configurations and ending with all pairs of maximal configurations of the two event structures. A technique for constructing a PSP of two acyclic PESs (without cc-pairs) has been proposed in [17]. In this section, we extend the notion of PSP and the PSP construction technique proposed in [17] in order to handle the case where one of the input event structures is the PES prefix of a process model (and thus contains cc-pairs), and the other is a PES derived from an event log as discussed in Section 5.2.

To illustrate the notion of PSP, consider the pair of event structures shown in Fig. 10. The synchronized product starts with the empty configurations. In this initial state, events a_0 from $\mathcal{E}^a$ and b_0 and b_1 from $\mathcal{E}^b$ are enabled. Since a_0 and b_0 carry the same label (i.e. A), an event match is asserted in the PSP via a so-called "match" operation. This operation leads to a state corresponding to the pair of configurations containing the occurrences of a_0 and b_0. In this state of the PSP, events a_1 from $\mathcal{E}^a$ and b_1 from $\mathcal{E}^b$ are enabled. Although events a_1 and b_1 carry the same label (i.e. B), they cannot be matched because of the discrepancy in the causal relation of the PESs: in $\mathcal{E}^a$ it holds that $a_0 < a_1$ whereas in $\mathcal{E}^b$ it holds that $b_0 \parallel b_1$. In other words, there is an error in the synchronized product of the two event structures. To recover from this error, events a_1 and b_1 are declared as "hidden" in the PSP and the synchronized simulation can proceed. This example illustrates two requirements for two events to be matched in the PSP, namely that an event matching must be label preserving (i.e. both events must have the same label) and order-preserving (i.e. event matchings in the PSP are consistent with the causal relation of the input PESs).

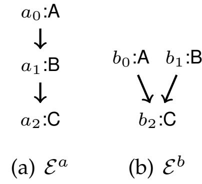

Fig. 10: Sample PESs

Fig. 11 presents a fragment of the PSP of event structures $\mathcal{E}^a$ and $\mathcal{E}^b$ shown in Fig. 10. In the general case, the PSP of a pair of event structures is not commutative. Therefore, we fix the following convention: the left-hand side PES is the one that is derived from an event log, while the PES derived from a process model is always at the right hand side. Correspondingly, we use "lhide" to denote the hiding of an event from the PES at the left-hand side and "rhide" when the hidden event comes from the other PES.[6]

6. In Fig. 11 the operations are equipped with parameters indicating the affected event. However, this is shown for understandability purposes only, as the event is implied by the difference in the configurations in the state and as such it is not included in Algorithm 1.

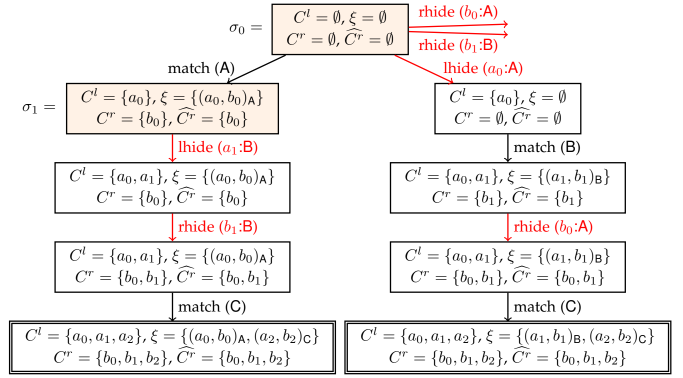

Fig. 11: Fragment of the PSP of $\mathcal{E}^a$ and $\mathcal{E}^b$

Note that the fragment of the PSP shown in Fig. 11 has two branches. The leaf state in each branch corresponds to states that can no longer be extended, because their state refers to maximal configurations. Informally, the left-hand side branch states that, without considering the occurrence of events a_1 and b_1, in both PESs we observe the execution of a task A followed by the execution C. If we consider the other branch, we would have a somehow symmetric conclusion. We note also that we could also find in the PSP a sequence of PSP operations hiding all the events from both event structures. However, such sequence would not be informative. Instead, we are interested only in sequences of PSP operations that maximize the number of event matches or, symmetrically minimize the number of hide operations, which we will call optimal event matchings.

From a conceptual point of view, a PSP is a directed acyclic graph. A node in a PSP represents a state in the synchronized simulation. An arc in a PSP, on the other hand, represents a transition between states in the simulation and is labeled by the type of operation that was used in the transition. Formally, a PSP is denoted as a tuple (Σ, OP, A), where Σ is a set of states; OP is the set of operations in a PSP, i.e. $OP \triangleq \{\text{match}, \text{lhide}, \text{rhide}\}$; and $A \subseteq \Sigma \times OP \times \Sigma$ is the set of directed labeled arcs. A state in the PSP stores the configurations of the input event structures, as a way to keep track of the moves in the synchronized simulation. Given a pair of event structures $\mathcal{E}^l$ and $\overline{\mathcal{E}^r}$, a state $\sigma \in \Sigma$ of a PSP is defined as a tuple $(C^l, \xi, C^r, \widehat{C^r})$. There, C^l and $\widehat{C^r}$ represent configurations from $\mathcal{F}(\mathcal{E}^l)$ and $\mathcal{F}(\overline{\mathcal{E}^r})$, respectively. Since $\overline{\mathcal{E}^r}$ is a PES prefix, $\widehat{C^r}$ may be the result of a shift operation. In the tuple, C^r is a multiset of events from $\overline{\mathcal{E}^r}$, which records not only the set of events but also the number of event occurrences during the synchronized simulation. Finally, $\xi \subseteq E(\mathcal{E}^l) \times E(\overline{\mathcal{E}^r})$ holds a set of event pairs, representing the event matches that have been observed in a path from the root state of the PSP to state σ.

We approached the problem of computing the set of optimal event matchings in the PSP with Algorithm 1, which is based on the well-known A* heuristic search [27]. The problem at hand can be naturally mapped into a multi-objective heuristic search. However, as for other domains, the memory requirements of a multi-objective approach are high. As a result, Algorithm 1 is designed as a single-objective heuristic search, which matches one maximal con-

figuration from the event log at a time. The result of this algorithm is later combined to produce the entire PSP. With abuse of notation, we use $C \xrightarrow{e} C'$ to denote the extension of the configuration C with event e, such that $C' = C \oplus e$.

Algorithm 1 Partially synchronized product

1: **function** BUILDPSP($\mathcal{E}^l$, $\overline{\mathcal{E}^r}$, C^l_{m})
2: OPEN $\leftarrow \{(\emptyset, \emptyset, \emptyset, \emptyset)\}$
3: Initialize PSP
4: **while** OPEN $\neq \emptyset$ **do**
5: Choose $\sigma = (C^l, \xi, C^r, \widehat{C^r}) \in$ OPEN, with min. $\varphi(\sigma, C^l_{\mathrm{m}})$
6: OPEN $\leftarrow$ OPEN $\setminus \{\sigma\}$
7: **return** (PSP, σ) **if** $C^l \in \mathcal{F}_{\mathrm{m}}(\mathcal{E}^l) \wedge \widehat{C^r} \in \mathcal{F}_{\mathrm{m}}(\overline{\mathcal{E}^r})$
8: **for each** $\left(\begin{array}{l} C^l \xrightarrow{e} C^{l\prime}, \widehat{C^r} \xrightarrow{f} \widehat{C^{r\prime}}, \text{s.t.} \\ e \in C^l_m \wedge \lambda^l(e) = \lambda^r(f) \wedge \\ \text{FINDCAUSALINC}(\text{PSP}, \sigma, e, f) = \bot \end{array}\right)$ **do**
9: $\sigma' \leftarrow (C^{l\prime}, \xi \cup \{(e, f)\}, C^r \cup \{f\}, \overline{\mathcal{S}}_f(\widehat{C^{r\prime}}))$
10: ADDARC(PSP, (σ, "match", σ'))
11: PUSH(OPEN, σ')
12: **end for**
13: **for each** $C^l \xrightarrow{e} C^{l\prime}$, s.t. $e \in C^l_m$ **do**
14: $\sigma' \leftarrow (C^{l\prime}, \xi, C^r, \widehat{C^r})$
15: ADDARC(PSP, (σ, "lhide", σ'))
16: PUSH(OPEN, σ')
17: **end for**
18: **for each** $\widehat{C^r} \xrightarrow{f} \widehat{C^{r\prime}}$ **do**
19: $\sigma' \leftarrow (C^l, \xi, C^r \cup \{f\}, \overline{\mathcal{S}}_f(\widehat{C^{r\prime}}))$
20: ADDARC(PSP, (σ, "rhide", σ'))
21: PUSH(OPEN, σ')
22: **end for**
23: **end while**
24: **end function**

Given two event structures $\mathcal{E}^l$ and $\overline{\mathcal{E}^r}$, and a maximal configuration C^l_{m} of $\mathcal{E}^l$, Algorithm 1 proceeds as follows. It starts by considering the root state (all configurations are set to empty set) in line 2 and enters a while loop in line 4, which will be repeated as long as there is an unprocessed state in OPEN. In line 5, a state σ is taken from OPEN such that σ has minimum cost φ. The cost function φ will be discussed later. In lines 8-12, the algorithm identifies the set of event matches. To this end, given the configurations C^l and $\widehat{C^r}$ in the state σ, we iterate over the set of possible extensions for both configurations. When a pair of events is found to be label-preserving and order-preserving, a new state σ' is instantiated and an arc $(\sigma, \text{match}, \sigma')$ is added to the PSP in line 10. Label preservation is straightforwardly checked by comparing the event labels, i.e. $\lambda^l(e) = \lambda^r(f)$. Order preservation, on the other hand, is checked by calling the function FINDCAUSALINC: if this function returns $\bot$ a pair of events is found order-preserving in state σ. The function FINDCAUSALINC will be further discussed later in this Section. Then, the new state σ' is added to OPEN. Note that the configuration $\widehat{C^r}'$ is updated accordingly (i.e. it is shifted) when the event f is a cutoff event. Then, a hide operation is processed, i.e. a new state and an arc are added to the PSP, for each possible extension found in the configuration C^l (lines 13-17) and C^r (lines 18-22). Note that we restrict the processing of events from $\mathcal{E}^l$ to those that are part of C^l_m. This is done by checking $f \in C^l_{\mathrm{m}}$ in lines 8 and 13. In this way, the search is explicitly directed to find an optimal matching for C^l_{m}. The algorithm stops in line 7 when it first reaches a state where C^l_{m} has been matched.

Let us now focus on the problem of checking if a candidate event match is order-preserving. A key challenge is to consider the shift operations that have occurred in the path that leads to a given state in the PSP. Algorithm 2 provides a solution to this problem.

Algorithm 2 Find causally inconsistent events in the PSP w.r.t. a given pair of events

1: **function** FINDCAUSALINC(PSP, σ, e, f)
2: cutoffs $\leftarrow \emptyset$
3: **while** $\exists(\sigma_{\text{pred}},\ op,\ \sigma) \in A(\text{PSP})$ **do**
4: $(e', f') \leftarrow$ GETDELTAEVENTS(σ, σ_{pred})
5: **if** f' is cutoff event **then**
6: cutoffs $\leftarrow$ cutoff $\circ f'$
7: $f \leftarrow \mathcal{I}^{-1}_{f'}[f]$
8: **end if**
9: **if** op is "match" **then**
10: **return** $(e, e', f, f', \text{cutoffs})$ **if** $\neg(e' < e \Leftrightarrow f' < f)$
11: **end if**
12: $\sigma \leftarrow \sigma_{\text{pred}}$
13: **end while**
14: **return** $\bot$
15: **end function**

The algorithm backward-traverses the PSP from a given state to the root state, checking if the input pair of events are causally consistent with the matched events as recorded in a path of the PSP. If a pair of events is found that is causally inconsistent, the algorithm returns a tuple containing the input events (possibly updated to compensate the effect of shift operations), the causally inconsistent events, and the sequence of cutoff events that are traversed during the search. This algorithm is used later in the characterization of behavior mismatches as explained in Section 7.

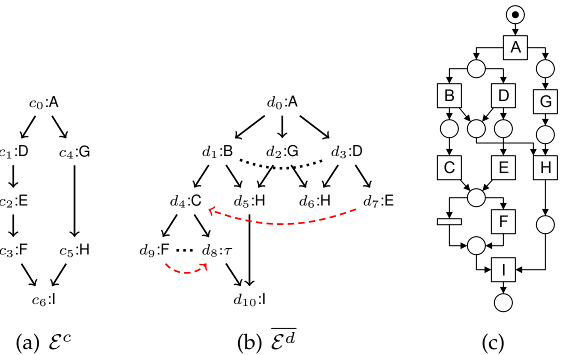

(a) $\mathcal{E}^c$ (b) $\overline{\mathcal{E}^d}$ (c)

Fig. 12: Sample PESs ($\overline{\mathcal{E}^d}$ is the PES prefix of (c))

To illustrate Algorithm 2, we will use the pair of PESs in Fig. 12. Here, $\overline{\mathcal{E}^d}$ is the PES prefix of the Petri net shown in Fig. 12(c). Assuming that the PSP, shown in Fig. 13, has been computed up to state s_2, events c_2 of $\mathcal{E}^c$ and d_7 of $\overline{\mathcal{E}^d}$ are enabled. As a result, function FINDCAUSALINC is called by Algorithm 1 (line 8). In the call, the input parameter PSP refers to the excerpt of the PSP shown of Fig. 13 comprising states s_0, s_1 and s_2; σ refers to state s_2; e refers to c_2 and f to d_7. In line 2, variable cutoffs – which stores the sequence of cutoff events found during the traversal – is initialized with an empty sequence. The while

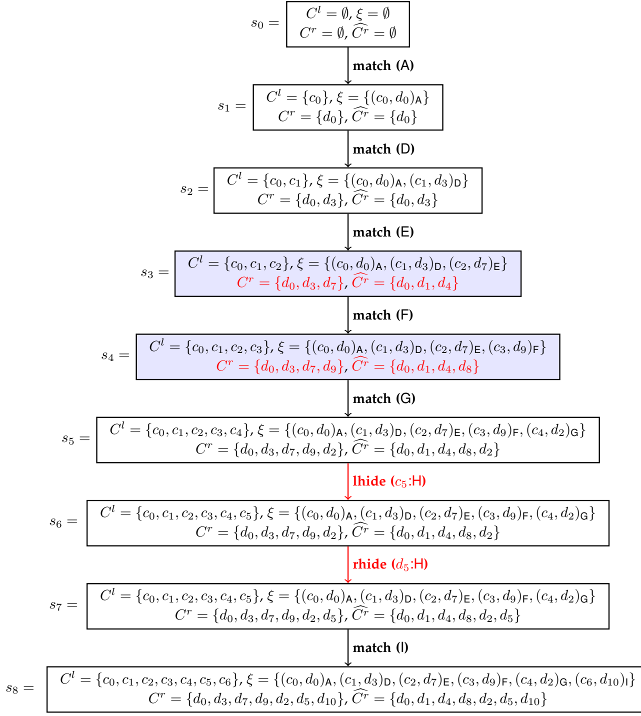

Fig. 13: Excerpt of the PSP of $\mathcal{E}^c$ and $\overline{\mathcal{E}^d}$

loop in lines 3-13 traverses the PSP backwards, processing one arc from the input PSP at a time. In the first iteration, the algorithm analyzes arc $(s_1, \text{"match"}, s_2)$. Next, in line 4, function GETDELTAEVENTS is called to determine the set of events involved in the operation associated with arc $(s_1, \text{"match"}, s_2)$. This function GETDELTAEVENTS can be straightforwardly implemented by computing the difference of C^l and C^r in the states s_2 and s_1. In this case, the function returns and set $e' = c_1$ and $f' = d_3$, respectively. Then, the block in lines 9-11 is executed because *op* is a "match" operation. Let us consider the expression in line 10, i.e. $\neg(e' < e \Leftrightarrow f' < f)$, which in this case is false, because it holds that $c_1 < c_2$ and $d_3 < d_7$. Clearly, the latter condition checks the consistency of the causal relations of input events (e.g. a pair of events that form candidate event match to be added to the PSP) with all the event matches recorded in the PSP that precede the state that activated the input events. In addition to validating the requirement of order preservation, as explained before, this function also returns the pair of events for which causal consistency does not hold if any pair is found. The first iteration concludes by updating σ with the value s_1 in line 12. In the second iteration, the algorithm processes arc $(s_0, \text{"match"}, s_1)$ in a similar way as in the first iteration. Since $c_0 < c_2$ and $d_0 < d_7$ hold true, order preservation is also decided true. This is the last iteration of the loop and the function returns $\perp$, indicating the absence of causal inconsistencies.

Let us now assume that the PSP has been computed up to state s_5 and function FINDCAUSALINC is called with events $e = c_5$ and $f = d_5$. In this case, this function processes the arcs associated with operations "match (G)" and "match (F)" in a similar way as described before. We further analyze the iteration where FINDCAUSALINC processes the arc associated with operation "match (E)". In this iteration, GETDELTAEVENTS returns and set $e' = c_2$ and $f' = d_7$, respectively. In line 6, event d_7 is appended to cutoff that, at this point, is set to $[d_9, d_7]$. Next, variable f is updated from d_5 to d_6 in line 7 (i.e. $\mathcal{I}_{d_2}^{-1}[d_5] = d_6$). It is because of this update that, in the following iteration of the while loop, the function finds that $c_1 \parallel c_5$ and $d_3 < d_6$, and hence concludes that the matching of c_5 and d_5 is not order-preserving.

We now turn our attention to defining the cost function φ, used in Algorithm 1. As for any conventional A*-based algorithm, the cost function is used as a criterion for selecting the next state to expand while constructing the PSP. As usual, the function φ relies into two other functions. The first function, called g, accounts for the cost of the PSP operations (i.e. match and hide operations) incurred in building the PSP from the root state up to a given state. The second function, referred to as h, corresponds to an optimistic approximation to the cost of PSP operations that will be required to reach the target state (the one that matches with a minimal cost the run that comes from the event log with one of the runs stemming from the model PES). Although all possible runs can be dynamically recomputed from the model PES prefix by applying the shift operation, we need a finite representation that could be used for formulating the function h. To this end, we will define a way to compute representative acyclic and cyclic runs in the form of an acyclic graph, where nodes are configurations and edges are configuration extensions. Inspired from the notion of elementary paths of a graph, we call such representation *elementary acyclic/cyclic pomsets*.[7] Specifically, we compute the set of all acyclic runs and that of cyclic runs, with the restriction that cyclic runs correspond to the unrolling of cyclic behavior once.

Let us first define acyclic runs, that is elementary acyclic pomsets. Let $\mathcal{X}_A(C)$ be the set of elementary acyclic pomsets that extend configuration C. Specifically, we say that $X_A \in \mathcal{X}_A(C)$ is an elementary acyclic pomset iff there exists an order $i \in [0 \ldots n-1]$ and a set of events $X_A = \{e_0 \ldots e_{n-1}\}$ with $|X_A| = n$ such that $\overline{\mathcal{S}}_{e_{n-1}}(\ldots(\overline{\mathcal{S}}_{e_0}(C))\ldots) \in \mathcal{F}_{\text{m}}$.

For defining cyclic runs, we need first to compute the unrolling of cyclic behavior until one full iteration is completed. Let $\mathcal{X}_R(C)$ be the set of elementary cyclic pomsets that extend configuration C. Specifically, we say that $X_R \in \mathcal{X}_R(C)$ is an elementary cyclic pomsets iff there exists an order $i \in [0 \ldots n-1]$ and a set of events $X_R = \{e_0 \ldots e_{n-1}\}$ with $|X| = n$ such that e_{n-1} is cutoff event, $\overline{\mathcal{S}}_{e_{n-1}}(\ldots(\overline{\mathcal{S}}_{e_0}(C))\ldots) \in \mathcal{F}$ and $\text{corr}(e_{n-1}) \in C \cup X_R$. Moreover, we will say that $\text{corr}(e_{n-1})$ is the entry event to the elementary cyclic pomset and will be denoted by $entry(X_R(C))$. Note that a repetitive suffix never reaches a maximal event because it finishes in the event that reenters a (possibly nested) loop, entailing the repetitive behavior. Finally, we will denote by $\mathcal{X}_W(C)$ the set of maximal elementary pomsets that extend configuration C. Intuitively, a maximal suffix of a configuration represents a complete run. Clearly, $\mathcal{X}_W(C)$ includes all the acyclic suffixes of configuration C, i.e. $\mathcal{X}_A(C) \subseteq \mathcal{X}_W(C)$. In the case of repetitive suffixes, we need to complement them with acyclic suffixes to complete sets of possible runs. The idea is that, once

7. *Pomset* (standing for *partially ordered multiset*) is a widely known formalism in the literature of concurrency theory [28], which corresponds with the intuition described here.

repetitive behavior has been executed one or several times, the computation should follow the sequence captured by an acyclic suffix. The set of possible acyclic suffixes to consider depends on the entry event of the repetitive suffix. We now can formally define $\mathcal{X}_W(C)$ as follows:

$$\mathcal{X}_W(C) = \mathcal{X}_A(C) \cup \left\{ X_A \cup X_R \;\middle|\; \begin{array}{l} X_R \in \mathcal{X}_R(C) \wedge X_A \in \mathcal{X}_A(C) \\ \wedge\ \text{entry}(X_R(C)) \in X_A \cup C \end{array} \right\}$$

The notion of elementary pomsets along with an algorithm to compute them all (e.g. as a preprocessing step) are further discussed later in Subsection 7.3.

Armed with the above, we define function φ as follows:

Definition 10. Let $\sigma = (C^l, \xi, C^r, \widehat{C^r})$ be a state in PSP and $C_\text{m} \in \mathcal{F}_\text{m}(\mathcal{E}^l)$ be a maximal configuration of $\mathcal{E}^l$ (the PES of the event log), such that $C^l \subseteq C_\text{m}$. With abuse of notation, given a configuration C we will use $\lambda(C)$ to denote the set of labels of the events in C, i.e. $\lambda(C) = \{\lambda(e) | e \in C\}$. The function φ is defined as

$$\varphi(\sigma, C_\text{m}) = g(\sigma) + h(\sigma, C_\text{m})$$

where

$$g(\sigma) = |C^l| + |C^r|_{\lambda^r}| - |\xi| * 2$$

and

$$h(\sigma, C_\text{m}) = \min_{X_W \in \mathcal{X}_W(C^r)} \left| \lambda^l(C_\text{m} \setminus C^l) \setminus \lambda^r(X_W) \right|$$

Intuitively, the cost function g corresponds to the number of hide operations incurred in the path starting from the root state in the PSP and leading to a given state. Clearly, we would like to find a sequence with only match operations, if such a sequence exists, or the path that includes the minimum number of hide operations. Note that in the case of C^r, we restrict our attention to the set of observable events, denoted $C^r|_{\lambda^r}$. In fact, we are interested in characterizing differences in terms of visible events, but we have to keep some of the invisible events in the PES prefix to maintain the information about cc-pairs in the complete prefix unfolding. Thus, in our definition of g we search to not penalizing operations that involve invisible events recorded in C^r.

As per the conventional A*-based algorithm, function h corresponds to an optimistic approximation to the "future cost": the cost of a path starting from a given state up to a goal state. In our context, the goal state corresponds to the optimal state in which a maximal configuration C_m (coming from the PES of the event log) is matched. To this end, we consider the possible futures for both configurations. In the case of C^l, we consider only the set of events in $C_\text{m} \setminus C^l$, because Algorithm 1 looks at finding an optimal matching for C_m. In the case of C^r, we consider the set of elementary acyclic/cyclic pomsets that can extend the configuration C^r. Note that we take into account all the possible elementary pomsets, and select the minimal cost of matching C^l with each elementary pomset as the value for h. The idea is that only the elementary pomset with minimal cost will eventually be used in the future, if that leads to the optimal matching. Since we are matching the set of event labels, the repetition of cycles is abstracted away. That is, the cost will be the same regardless of the fact that repetitive behavior occurs once or more times, provided that the matching of the full behavior is found to use the same number of hide operations. Since h is an optimistic cost function, it follows that Algorithm 1 is admissible [27] and hence it returns an optimal solution.

When the input PESs have concurrent behavior, the PSP may contain paths with redundant information. This redundancy stems from the interleaved enablement of concurrent events. Fig. 14 gives an example of this situation.

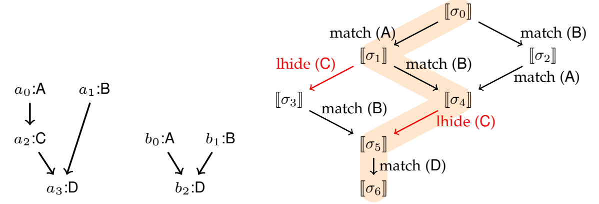

Fig. 14: PSP with redundant information due to concurrency

From this figure it can be easily checked that the PSP contains three different paths, all capturing the same information. This is a well known problem in areas such as Model Checking, where techniques have been developed to discard some paths when exploring the underlying state space [29]. Therefore, we can leverage results from that field to reduce redundant information at the time of the construction of the PSP. In that context, it is well known that the exploration of the state space can be reduced by analyzing the commutativity of the transitions in the state space, which translates to our setting to the notion of commutativity of operations. Intuitively, we say that a pair of operations can be commuted if their underlying events are enabled concurrently. For instance, the event a_0 and a_1 are concurrently enabled by the empty configuration and so are the events b_0 and b_1 in the example on Fig. 14. As shown in the PSP, the operations "match (A)" and "match (B)" appear in two distinct orders, in paths that start in state σ_0 and finishing in σ_4. We will say that operations "match (A)" and "match (B)" are commutative. A similar situation happens at state σ_1 with operations "lhide (C)" and "match (B)", because the events a_1 and a_2 are concurrent.

Commutativity of operations is closely inspired by the notion of commutativity of transitions in model checking. This property is at the heart of a large number of partial order reduction techniques used in model checking [29]. What is more, we can recover their theoretical results to claim that one path encodes the same information as the other paths, reducing significantly the size of PSP. In this respect, Algorithm 3 presents the modifications to apply on Algorithm 1 to achieve the partial order reductions.

This modified algorithm would build a reduced PSP for the example in Fig. 14 as follows. When the algorithm analyzes the root state in the PSP, that is σ_0 in Fig. 14, it processes the operations "match (A)" and "match (B)". Observe that we assume that configuration extensions are ordered according to the event labels. This heuristic allows us to implement the partial order reduction in a more straightforward way. Therefore, "match (A)" will be appended first to the PSP and the variable $\mathbb{E}$, standing for *enablements*, will be set to $\{(a_0, b_0)\}$. In the next iteration

Algorithm 3 Partially synchronized product with partial order reductions

▷ Replace lines 8-12 with the following block

1: $\mathbb{E} \leftarrow \emptyset$
▷ *Consider configuration extensions according to the lexicographical order of the event labels*
2: **for each** $\left(\begin{array}{l} C^l \xrightarrow{e} C^{l\prime}, \widehat{C^r} \xrightarrow{f} \widehat{C^{r\prime}}, \text{s.t.} \\ e \in C^l_m \wedge \lambda^l(e) = \lambda^r(f) \wedge \\ \text{FINDCAUSALINC}(\text{PSP}, \sigma, e, f) = \bot \end{array} \right)$ **do**
3: **if** $\exists (e', f') \in \mathbb{E}$, s.t. $e \parallel e' \vee f \parallel f'$ **then**
4: **continue**
5: **else**
6: $\mathbb{E} \leftarrow \mathbb{E} \cup \{(e, f)\}$
7: **end if**
8: ▷ Keep lines 9-11 as they are
9: **end for**

▷ Replace line 13 with the following one

for each $C^l \xrightarrow{e} C^{l\prime}$, s.t. $e \in C^l_m \wedge \nexists (e', f') \in \mathbb{E} : e \parallel e'$ **do**

▷ Replace line 18 with the following one

for each $\widehat{C^r} \xrightarrow{f} \widehat{C^{r\prime}}$, s.t. $\nexists (e', f') \in \mathbb{E} : f \parallel f'$ **do**

of the for loop, the algorithm will no longer consider the operation "match (B)" as a successor of σ_0, as if the events a_1 and b_1 were not enabled in such a state. For the same reason, the hide operations associated with events a_1 and b_1 will not be appended to σ_0 either. In the next iteration, the algorithm will consider the operations "match (B)", which was warranted because of the commutability of the operations "match (A)" and "match (B)". Once the operation"match (B)" is appended, the algorithm proceeds as usual (following the path highlighted in Fig. 14), because all the remaining operations are not commutative.

We observe that in Algorithm 3 the testing of commutativity has been slightly modified. It is by means of the expression in line 3, namely $e_1 \parallel e_1' \vee e_2 \parallel e_2'$, that the algorithm checks commutativity of operations. In fact, line 3 will discard any match operation that is found commutative with one that has been previously appended during a previous iteration of the for loop. Let us further analyze the condition in line 3. First, if both terms in $e_1 \parallel e_1' \vee e_2 \parallel e_2'$ hold true, then it is because we have two potential match operations that are clearly commutative. Fig. 15 presents an example that illustrates the second case. Assume that events a_1 and b_1 have already been processed (i.e. $(a_1, b_1) \in \mathbb{E}$). When we want to process the match operation associated with the events a_2 and b_2, we will have that the first term in the expression $e_1 \parallel e_1' \vee e_2 \parallel e_2'$ hold true whereas the second term does not, because $b_1 \# b_2$. Fig. 15 shows e_1, e_2, etc. in red font to ease the mapping of the example to the expression. Due to the presence of conflict, the operation "match (C)" will not be appended to the PSP in the path that follows the operation "match (B)" and, as a result, we will see a "lhide (a_1)" and this operation is commutative with "match (B)" The remaining case, that is when the first term in $e_1 \parallel e_1' \vee e_2 \parallel e_2'$ is false and the second term is true, is the symmetric of the second case.

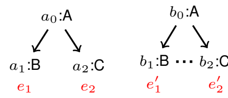

Fig. 15: Example of PESs

Complexity analysis. The complexity of the PSP computation depends on the size of the state space explored by the A* search. This state space is in $O(3^{|\mathcal{F}(E_1)| \cdot |\mathcal{F}(E_2)|})$ where $\mathcal{F}(E)$ is the set of configurations of E. Indeed, each configuration in E_1 is associated with a configuration from E_2 via three possible operations (i.e. $match$, $lhide$ and $rhide$). This worst case complexity may be reached when the event structures are completely different. Conversely, when the event structures are identical, the heuristic search converges in linear time. In practice, we expect a high behavior overlap between a process model and the corresponding log, and hence a complexity far below the worst case.

7 Difference Extraction and Verbalization

In the previous section, we showed that a PSP contains a minimal set of hide operations required to capture all behavioral discrepancies between the PES of an event log and the PES of a process model. In this section, we show how to traverse the PSP in order to extract a set of difference statements that characterize the behavior observed in the log and not captured in the model and vice-versa.

In order to extract such difference statements, we rely on a collection of nine *mismatch patterns* classified into the following disjoint categories:

- *Unfitting behavior patterns*, capture behavior observed in the log but not allowed by the model. Unfitting behavior patterns are further classified into two subcategories:
  - *Relation mismatch patterns*, capture cases where two events in the PES of the log are related via a given behavioral relation (immediate causality, direct conflict or concurrency) but they are related via a different relation in the PES prefix of the model – e.g. they are related via immediate causality in the log while they are related via direct conflict in the PES of the model.
  - *Event mismatch patterns*, capture all other cases of unfitting behavior. These patterns capture events in the PES of the log that cannot be directly matched to an event in the PES prefix of the model.
- *Additional behavior patterns*, cover all cases where behavior is allowed in the model but not observed in the log.

In the following we describe each pattern.

7.1 Relation mismatch patterns

The first category of mismatch patterns corresponds to situations where a pair of events – one from the log PES and one from the model PES prefix – have the same label but they are not matched in the PSP because they have different behavioral relations with at least one other event. In other words, there is a pair of events in the PES of the log linked via a given relation, which could be matched to a corresponding pair of events in the PES prefix of the model, if it was not for the fact that these pairs are related via different behavioral relations. Since there are only three behavior relations, there are also three possible (symmetric)

relation mismatches: *Immediate Causality vs. Concurrency*, *Immediate Causality vs. Direct Conflict*, and *Concurrency vs. Direct Conflict*.[8] As we will see later, the last two types of mismatches (those involving conflict) have very similar manifestations in the PSP and hence we will treat them as one single pattern. Hence below we introduce two patterns, namely *Causality/Concurrency* and *Conflict*.

Like all other patterns introduced later, relation mismatch patterns occur in a given *context*, characterized by a pair of configurations (one configuration from each PES). Consider the example shown in Fig. 16. We note that $a_1 \parallel a_2$ whereas $b_1 < b_2$ and these two pairs of events have matching labels. Thus, there is a Causality/Concurrency mismatch. This mismatch is observed in the state of the PSP associated with the pair of matching configurations $\{a_0, a_1\}$ and $\{b_0, b_1\}$ (the mismatch context). We also note that one of the events that is hidden in the PSP (μb_2) is the target of an immediate causality relation stemming from an event in the configuration $\{b_0, b_1\}$ (specifically note that $b_1 <_\mu b_2$).

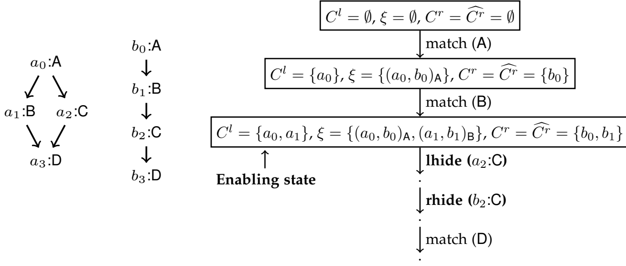

Fig. 16: Causality/Concurrency mismatch

We also observe that a relation mismatch pattern manifests itself in the PSP in the form of two "hide" operations. In the example shown in Fig. 16, the pair of "hide" operations are contiguous in one path of the PSP. However, the hide operations do not necessarily happen always contiguously in the PSP as shown in the example of Fig. 17. The reason why the "hide" operations are not contiguous in the PSP of this second example stems from the fact that "match (C)" and "lhide (a_1:B)" are commutative. Hence, the identification of this type of mismatches requires one to take into account the commutativity of operations.

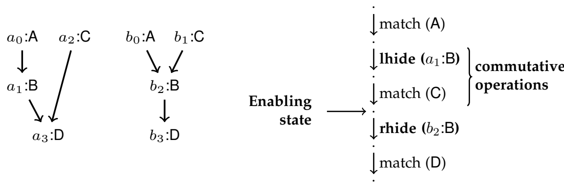

Fig. 17: Another Causality/Concurrency mismatch

The proposed approach to identify Causality/Concurrency relation mismatches (later referred to as CAUSCONC mismatches) is formalized in Algorithm 4. This algorithm analyzes one path from the PSP at a time (line 2). In line 3, it selects three arcs each corresponding to an "lhide", an "rhide" and a "match" operation respectively. Next, in lines 4-6 it determines the set of events that are involved in the three operations. In line 7, it maps cutoff with its corresponding event, if required. Note that in lines 8-10, the algorithm discards the operations being analyzed (i.e. the loop is forced to proceed with the next iteration in line 9), if the same pattern can be built with another hide operation that is causally predecessor of one of the hide operations being processed. This situation is checked by function CHECKPREDS. Finally, in line 11 the algorithm checks the conditions defining an elementary Causality/-Concurrency mismatch: the pair of hidden events carry the same label and one pair of events are in immediate causal relation whereas the other pair is concurrent.

Algorithm 4 Finding Causality/Concurrency mismatches

1: **procedure** FINDCAUSALCONCMISMATCHES(PSP)
2: **for each** PATH in PSP **do**
3: **for each** $(\sigma_1, \text{lhide}, \sigma_1'), (\sigma_2, \text{rhide}, \sigma_2'), (\sigma_3, \text{match}, \sigma_3') \in$ PATH **do**
4: $(\bot, f) \leftarrow$ GETDELTAEVENTS(σ_1', σ_1)
5: $(e, \bot) \leftarrow$ GETDELTAEVENTS(σ_2', σ_2)
6: $(e', _f') \leftarrow$ GETDELTAEVENTS(σ_3', σ_3)
7: $f' \leftarrow$ ISCUTOFF$(_f')$? corr$(_f')$: $_f'$
8: **if** $\left(\begin{array}{l} \text{CHECKPREDS(PSP}, \sigma_3', \sigma_1, \text{“lhide”}, e, \lambda_l(e)) \vee \\ \text{CHECKPREDS(PSP}, \sigma_2', \sigma_1, \text{“rhide”}, f, \lambda_r(f)) \end{array} \right)$ **then**
9: **continue**
10: **end if**
11: **if** $\lambda_l(e) = \lambda_r(f) \wedge (e' <_\mu e \vee f' <_\mu f) \wedge (e' \parallel e \vee f' \parallel f)$ **then**
12: **assert**(CAUSCONC$(\sigma_3, e, f, e', _f', f')$)
13: **end if**
14: **end for**
15: **end for**
16: **end procedure**
17: **function** CHECKPREDS(PSP, σ_0, σ_2, op, e, $label$)
18: ▷ This function returns $true$ if there does not exist an arc $(\sigma, (op), \sigma') \in A(\text{PSP})$ that involves an event that carries the label $label$ and that causally precedes event e, and $false$ otherwise
19: **end function**

Let us now turn our attention to the two other cases, namely *Causality/Conflict* and *Concurrency/Conflict*. Again, we observe that these mismatches show up in the PSP in the form of two "hide" operations, but these two "hide" operations appear in different branches. The latter holds because configurations are conflict-free (i.e. two conflicting events cannot occur in the same computation).

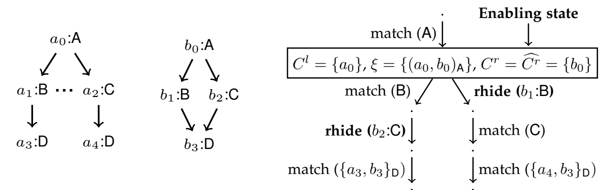

Fig. 18: Concurrency/Conflict mismatch

The example shown in Fig. 18 corresponds to a case of *Concurrency/Conflict* mismatch. There, it can be seen that the hide operations occur in different branches of the model on the right. As the left model allows either B or C (whereas they are concurrent in the right model), either one of the two events needs to be hidden in the model on the right. Due to the presence of concurrency, the hide operation and the conflicting match operation can appear in either order. Fig. 19 presents a more complex situation. In this example,

8. We only report mismatches involving immediate causality (not transitive causality) and direct conflict (not transitive conflict), because we are only interested in reporting the origin of a difference.

a *Causality/Conflict* mismatch is intertwined with a pair of concurrent events. This leads to the hide operation and the conflicting match operation being separated. Therefore, an approach to identify these mismatches must take into account the commutativity of operations, in the same way as for the *Causality/Concurrency* mismatch pattern.

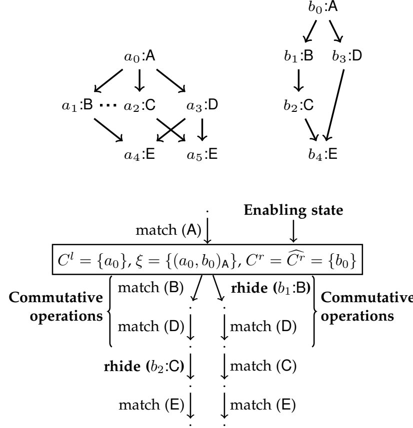

Fig. 19: Causality/Conflict mismatch

The above observations on the characteristics of mismatch patterns involving conflict are formalized in Algorithm 5, which can identify both *Causality/Conflict* and *Concurrency/Conflict* mismatches (herein referred to as CONFLICT mismatches). Algorithm 5 is similar way to Algorithm 4, except for two key points. First, Algorithm 5 processes three arcs at a time, but the arcs are not required to come from the same path. Recall that the hide operations associated with a Causality/Conflict or a Concurrency/-Conflict mismatch pattern will be located in two different branches. Second, the conditions in line 10 are the ones that define a conflict-related mismatch: the hidden events carry the same label and one pair of events are in immediate conflict relation whereas the other pair are in either causal or in concurrency relation. Note that the symmetric condition holds, i.e. $\neg(e'\#e) \wedge (f'\#f)$.

Algorithm 5 Finding Conflict mismatches

1: **procedure** FINDCAUSALCONCMISMATCHES(PSP)
2: **for each** $(\sigma_1, \text{lhide}, \sigma_1'), (\sigma_2, \text{rhide}, \sigma_2'), (\sigma_3, \text{match}, \sigma_3') \in A(\text{PSP})$ **do**
3: $(\perp, f) \leftarrow$ GETDELTAEVENTS(σ_1', σ_1)
4: $(e, \perp) \leftarrow$ GETDELTAEVENTS(σ_2', σ_2)
5: $(e', _f') \leftarrow$ GETDELTAEVENTS(σ_3', σ_3)
6: $f' \leftarrow$ ISCUTOFF$(_f')$? corr$(_f')$: $_f'$
7: **if** (CHECKPREDS$(\sigma_3', \sigma_1, \text{"lhide"}, \lambda_l(e)) \vee$ CHECKPREDS$(\sigma_2', \sigma_1, \text{"rhide"}, \lambda_r(f))$) **then**
8: **continue**
9: **end if**
10: **if** $\lambda_l(e) = \lambda_r(f) \wedge (e'\#_\mu e \vee f'\#_\mu f) \wedge \neg(e'\#e \wedge f'\#f)$ **then**
11: **assert**(CONFLICT$(\sigma_3, e, f, e', _f', f')$)
12: **end if**
13: **end for**
14: **end procedure**

Complexity analysis. Algorithms 4 and 5 are O(n^3), where n is the number of arcs in the PSP. However, with some technical optimizations, the identification of all relation mismatch patterns can be implemented using a single depth-first search traversal of the PSP – hence with an O(n) complexity. Details of how the two algorithms can be combined into a single depth-first search traversal are omitted as they are purely technical optimizations.

7.2 Event mismatch patterns

In this category of patterns, we group together all cases of unfitting behavior that cannot be characterized via a relation mismatch. A naive way of characterizing such cases would be to simply state that there are some events in the PES of the log that are not matched to any event in the PES prefix of the model. However, such an approach to diagnose differences is too low level and would lead to a high number of difference statements. Instead, we introduce four mismatch patterns that capture possible reasons for the presence of an unmatched event at a higher level of abstraction, namely *task skipping*, *unmatched repetition*, *task substitution* and *task relocation*. When a given unfitting behavior cannot be characterized using any of these four patterns, we use a fifth "catch all" pattern (namely *Task absence*), which essentially states that there is an event that can occur in the PES of the log in a given configuration but not in the corresponding configuration in the PES prefix of the model. Below we present these five patterns in turn.

Task skipping (TASKSKIP). This pattern is illustrated in our running example and, for discussion purposes, in the PSP fragment presented in Fig. 20. The way this pattern shows up in the PSP bears some similarity with how the *Causality/Conflict* mismatch pattern shows up, in the sense that it requires us to combine information coming from two branches of the PSP stemming at a given state. What makes this pattern different from the *Causality/Conflict* mismatch is that the operation "match $((a_2, b_2)_{\mathsf{C}})$" (which is interfering with the event a_1 (i.e. $a_1 \#_\mu a_2$) has a counterpart match operation in the other branch, namely "match $((a_3, b_2)_{\mathsf{C}})$" and both match operations involve the event b_2.

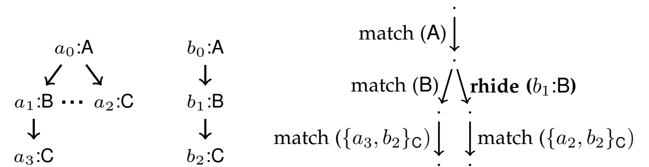

Fig. 20: Task skipping

The identification of this pattern requires a second traversal of the PSP. In the first traversal, a *Causality/Conflict* mismatch is identified and the information about the state where this mismatch is enabled is also gathered. In the second traversal, we analyze the sibling branches to look for the counterpart match operation. If the latter is found, we assert an occurrence of a *Task skipping* pattern instead of asserting an occurrence of a *Causality/Conflict* pattern. We do not provide a separate algorithm to detect this pattern as it would be largely redundant with Algorithm 5.

Unmatched repetition (UNMREPETITION). A second scenario where one hide operation cannot be matched is when the event log is capturing repetitive behavior that is not specified in the process model. In that context, every occurrence in the log of the same label will be mapped to a

different event. Every time an event is repeated in the log that cannot be matched a hide operation will be appended to the PSP. Fig. 21 presents a simple example of this pattern.

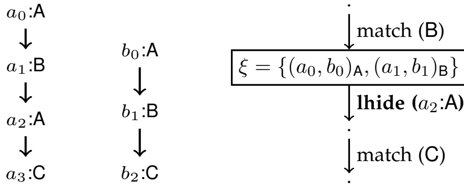

Fig. 21: Unmatched repetition

This pattern can be straightforwardly detected in the PSP by analyzing the set of matchings ξ, which is stored along with every state in the PSP. In our example, we observe that ξ contains the match $(a_0, b_0)_A$ that carries the same label as event a_2. We can therefore conclude that there is an activity with label A that occurs twice in the same trace, but cannot be matched to the behavior specified in the model. This test can be piggybacked in Algorithm 4 in line 12 and not adding the hide operation to n_chs if it has been found to be an unmatched repetition.

Note that the symmetric case (where repetitive behavior specified in the process model cannot be matched with behavior observed in the event log) cannot be processed in the same way. This is because the PES corresponding to the process model explicitly represents repetitive behavior by means of cc-pairs and shift operations. The problem of identifying repetitive behavior captured in the model but not observed in the log is discussed later (cf. additional behavior patterns).

Task substitution (TASKSUB). In some cases, an event cannot be matched because its counterpart has been substituted by a task with a different label. Fig. 22 presents an example of this pattern.

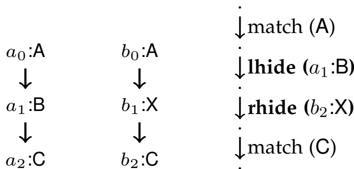

Fig. 22: Task substitution

Our assumption is that a task substitution must happen in the same execution context. Concretely, we require that the candidate events are enabled immediately after the events involved in a "match" operation. Given this requirement, the identification of this mismatch matching can be done with a variant of Algorithm 4. The changes to that Algorithm are basically to modify the condition of line 11 to eliminate the requirement about the equality of the event labels, checking that the events are immediately activated after a match operation $(e' <_\mu e \wedge f' <_\mu f)$ and, additionally, that the hide operations are causally consistent with all the operations that precede the event match. To make this analysis deterministic, we order the hide operations in alphabetic order of the event labels.

Task relocation (TASKRELOC). Let us now consider the case where a pair of events carrying the same labels cannot be matched because they appear in different places in the same path of a PSP. The main difference with respect to the relation mismatch patterns is that the events are not enabled after the same event. One simple case is the one where the order of a pair of events in two PESs is inverted. For instance, let us assume that in one PES it holds a_1:A $< a_2$:B whereas in the other PES it holds b_1:B $< b_2$:A, and there exists one state in the PSP where a_1 and b_1 are both enabled. Evidently, the PSP would have two different branches, each one with two hide and one match operation, respectively. This situation can be generalized to the case the events are not contiguous, as illustrated in Fig. 23.

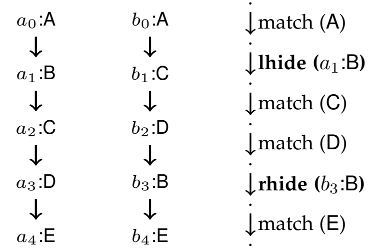

Fig. 23: Task relocation

Occurrences of the relocation pattern can be identified by keeping track of the events that were not found to be part of a relation mismatch pattern, and then checking equality of the labels associated with the hidden events.

Task absence/insertion (TASKABS). Any hide operation in the log of the PES that is not involved in any occurrence of one of the previous patterns is treated as an occurrence of a *Task absence* pattern, meaning that a task is observed in the log but missing in the model.[9] In other words, *Task absence* is a "catch all" pattern for all remaining cases of unfitting log behavior, thus ensuring that the set of patterns is complete. In the simplest case, an occurrence of the *Task absence* pattern corresponds to the situation where a task label is observed in the event log despite the fact no task with such label is specified in the process model. However, this pattern also captures the case where there exists at least a pair of events (one from each PES) with the same label, which are enabled in different states. The example shown in Fig. 24 illustrates the latter situation.

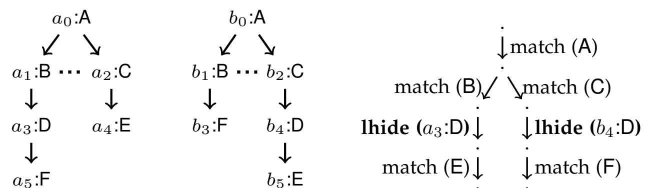

Fig. 24: Task absence/insertion

Complexity analysis. As said before, Task skipping requires traversing the PSP twice, making its complexity linear on

9. Symmetrically, we can state that a task has been inserted in the log.

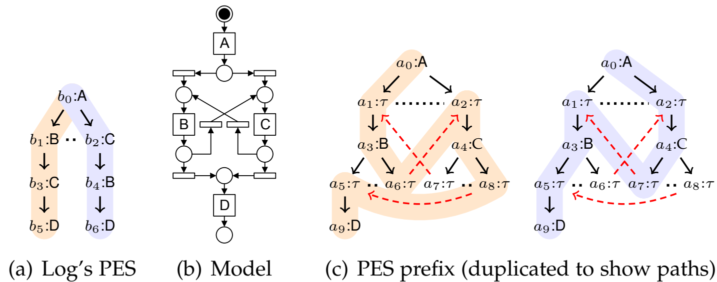

(a) Log's PES (b) Model (c) PES prefix (duplicated to show paths)

Fig. 25: Additional (cyclic) model behavior

the size of the PSP. All the other patterns in this section can be computed as part of one of the traversals.

7.3 Patterns of additional model behavior

The seven patterns presented above characterize behavior observed in the event log but not allowed in the process model. We now seek to characterize behavior allowed in the model but not observed in the log. Such additional behavior is captured by two patterns:

- *Unobserved acyclic interval* (UNOBSACYCLICINTER) – an acyclic fragment of a process model not observed in the log. Each such fragment is characterized by an initial task and a final task and is thus called an *interval*.
- *Unobserved cyclic interval* (UNOBSCYCLICINTER) – a cyclic fragment (*interval*) of a process model not observed in the log.

An example of additional model behavior is depicted in Fig. 25. Fig. 25(a) denotes a PES constructed from a log. Next, Fig. 25(b) shows a process model (in the form of a Petri net), while Fig. 25(c) shows the PES prefix derived from the model. The PES prefix appears in two copies, in order to show how each of the two paths in the PES of the log is also found in the PES of the model. In other words, there is no unfitting behavior in this example. On the other hand, there is additional behavior: the PES of the log does not contain any repetitive behavior, while the process model has a loop with two entry points and two exit points. Yet, if we constructed the PSP, we would find that it contains no hide operations and it covers all events and causal relations in both PESs. This is because the PSP is constructed with the goal of finding optimal matchings for every maximal configuration of the log's PES, and does not try to achieve full coverage of the model's PES prefix. In other words, a PSP with no hide operations only means that all behavior observed in the event log is fully captured in the process model, but not vice-versa.

Hence, in order to characterize additional model behavior (both acyclic and cyclic), we need to define a notion of *coverage* of the PES prefix of a process model. In other words, we need to answer the question: What does it mean that all the behavior captured in a process model is "covered by" (i.e. observed in) an event log? To answer this question, we use the notions of elementary paths and elementary cycles from the field of graph theory [30], [31]. Intuitively, we will say that an event log "covers" the behavior of a process model, if every elementary path and elementary cycle in the PES prefix of the model is represented by a path in the PSP, and thus represented by a maximal configuration in the PES of the log after hide operations have been applied to account for unfitting log behavior.

We recall some basic definitions from graph theory. A directed graph is a set of vertices and a set of directed edges. A path is a sequence of vertices connected by edges. A path is said elementary if no vertex is contained twice. A cycle is a path where the initial and the final vertex are the same. A cycle is said to be elementary if, after removing the last vertex in the sequence, the resulting path is elementary.

The above concepts provide a straightforward approach to define a notion of coverage of a graph by a set of traces. However, we cannot directly apply the above concepts to characterize the possible executions a PES prefix. Indeed, a path (along the direct causality relation) in a PES prefix does not characterize a possible execution, because an execution may contain concurrent events, and a single path in the PES prefix necessarily misses some of these events. Instead, we characterize the executions of a PES prefix by means of the set of elementary paths on a graph where the vertices are the configurations explicitly represented in the PES prefix, and the edges are the possible configuration extensions (i.e. direct transitions from one configuration to the next), including possible extensions induced by a cc-pair in the PES prefix.

In order to reason over this graph of configurations and configuration extensions, we rely on the notion of *pomset* from the literature of concurrency theory [28], which we introduced in Section 6. A pomset is a Directed Acyclic Graph (DAG) where the nodes are configurations, and the edges represent direct causality relations between configurations. An edge is labeled by an event. Unlike an event structure, a pomset does not have any conflict relation, since a pomset represents one possible execution. The behavior of a PES can be characterized by the set of pomsets it induces.[10]

In the case of a PES prefix, the set of induced pomsets is infinite when the PES prefix captures cyclic behavior via cc-pairs. Hence in general we cannot enumerate all pomsets of a PES prefix in order to check if each of them is observed in the PES of the log. However, we can extract a set of elementary pomsets (inspired by the notion of elementary paths), which collectively cover all the possible pomsets induced by a PES prefix without unfolding the cyclic behavior infinitely. Intuitively, this corresponds to unfolding every cycle so that it is traversed once only.

This intuition is formalized by Algorithm 6, which computes the set of *elementary pomsets* of a PES prefix. Fig. 26 illustrates the execution of this algorithm taking as input the PES shown in Fig. 25. Function FINDEPOMSETS builds an expanded prefix by successively applying configuration extensions and shift operations (cf. Section 5.1) on the PES prefix. Specifically, function FINDEPOMSETS adds one path (or branch) to the expanded prefix every time an elementary pomset is found (in lines 10 and 17). The result is a directed acyclic graph reflecting the configuration extension relation. For illustration purposes, Fig. 25 presents the expanded prefix as a PES prefix, with the corresponding label in the right-hand side of each "event" in the expanded prefix.

10. This is exactly the definition of a notion of equivalence known as *visible pomset equivalence*: two PESs are equivalent iff they induce the same set of pomsets.

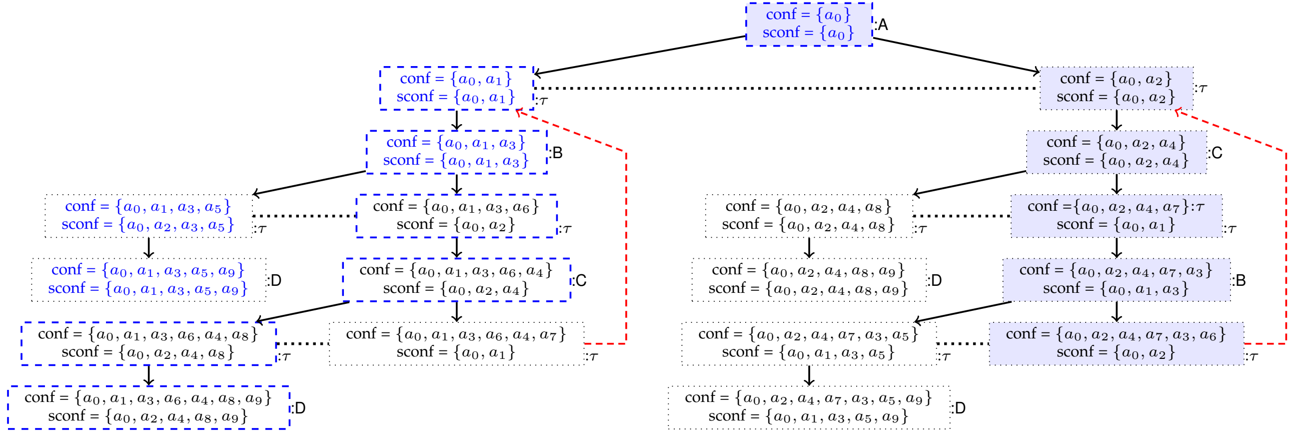

Fig. 26: Expanded prefix with elementary pomsets of PES in Fig. 25(c)

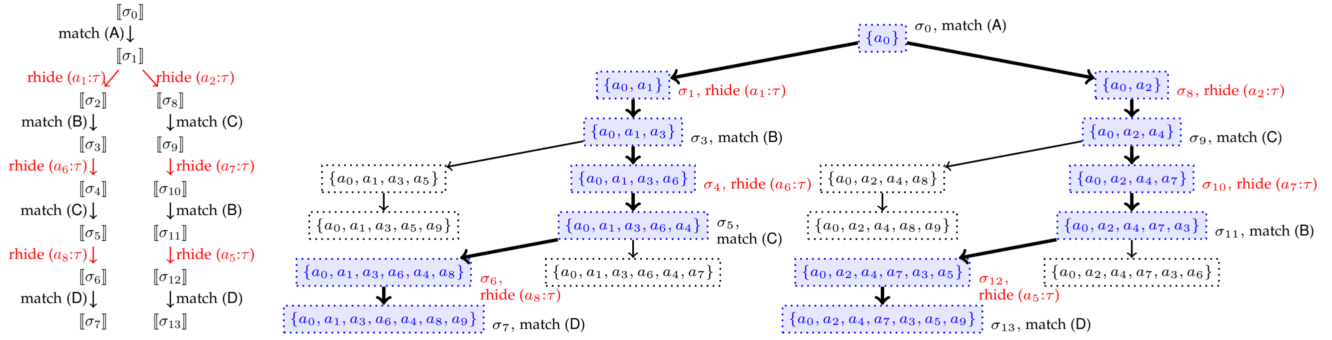

Fig. 27: Using PSP and expanded prefix for identifying additional model behavior in the example shown in Fig. 25

Algorithm 6 Identification of elementary pomsets

1: **procedure** FINDEPOMSETS(conf, sconf, visited, **var** cycles, **var** runs)
2: APPEND(visited, (conf, sconf))
3: **for** (sconf $\xrightarrow{e}$ n_sconf) **do**
4: n_conf $\leftarrow$ conf $\cup$ $\{e\}$
5: **if** e is cutoff event **then**
6: n_sconf $\leftarrow$ $\mathcal{S}$(n_sconf)
7: **end if**
8: **if** $\exists$(entryConf, n_sconf) $\in$ visited $\wedge$ n_sconf $\cap$ $\|[e]$ $= \emptyset$ **then**
9: cycles $\leftarrow$ cycles $\cup$ $\{$(n_conf $\setminus$ entryConf, entryConf)$\}$
10: ADDBRANCHTOEXPPREFIX(visited)
11: **else**
12: FINDEPOMSETS(n_conf, n_sconf, visited, cycles, runs)
13: **end if**
14: **end for**
15: **if** sconf *is a maximal configuration* **then**
16: runs $\leftarrow$ runs $\cup$ $\{$conf$\}$
17: ADDBRANCHTOEXPPREFIX(visited)
18: **end if**
19: REMOVELAST(visited)
20: **end procedure**

A key observation is that the value of conf can be used as a unique identifier for each event in the expanded prefix. The uniqueness of conf stems from the following facts: 1) for elementary acyclic pomsets, conf is finitely extended by a different event until a complete configuration is found, 2) for elementary cyclic pomsets, conf would be finitely extended up to the point where a duplicate event occurs. It is by means of the shifted version of conf, i.e. the variable sconf, that cycles can be identified.

The proofs of completeness and correctness of Algorithm 6 follow directly from the proofs for the algorithms to identify elementary cycles [30], [31]. With the aim of explaining how Algorithm 6 works and to sketch the proofs, we describe three cases.

Case 1: *Identification of elementary acyclic pomsets that include no cutoff event.* This case is illustrated with the sequence of "events" that are shown in blue font in the expanded prefix shown in Fig. 26. The function FINDEPOMSETS is first called with conf and sconf set to empty set. We note that in this case, conf and sconf are updated in such a way that they both hold the same value. The function FINDEPOMSETS extends the configuration conf by one event in line 4 and recursively calls itself in line 12. The recursive call will eventually stop, when sconf contains a maximal configuration, because a maximal configuration has no further possible extension. Moreover, a maximal configuration is warrantied to be found because we consider only PESs coming from sound systems and consisting of a finite number of events. The elementary acyclic pomset that has been found is added to the set runs in line 16, just before the recursive call returns.

Case 2: *Identification of elementary acyclic pomsets that contain at least one cutoff event.* This case is illustrated with the sequence of events in the blue dashed box in the expanded prefix in Fig. 26. As for the previous case, conf and sconf contain the same value at every recursive call of FINDEPOMSETS, as long as no cutoff event is found. When a (forward) cutoff event is found, such an event is added to conf and,

afterwards, n_sconf is assigned with the shifted configuration $\mathcal{S}(\text{sconf} \oplus \text{e})$, in line 6. In line 12, FINDEPOMSETS is recursively called with the extended configuration n_conf and also with the shifted configuration n_sconf. Note that sconf is used for testing if the function has found a maximal configuration in line 15 and also to compute the set of possible extensions in line 3. Conceptually, conf keeps track of the events that are "executed" by the underlying run, dynamically unrolling towards a larger prefix, whereas sconf maps the execution back to a configuration in the original PES prefix. Note that the recursive call to the function can only find a finite number of forward cutoff events. Thus, if no cycle is found as the recursive call to FINDEPOMSETS proceeds, the function will extend the configuration until a maximal configuration is found.

Case 3: *Identification of elementary cyclic pomsets.* We observe that, when the input PES has repetitive behavior, the function FINDEPOMSETS can only be recursively called a finite number of times before it finds an elementary cyclic pomset, because the PES has only a finite number of events. One case corresponds with finding a cutoff event that induces a backward shift. For simplicity, let us assume that the input PES has only one cutoff. Let e be a cutoff event. By definition, we know that e induces a backward shift if and only if $\lceil \text{corr}(e) \rceil \subset \lceil e \rceil$. Therefore, there exists a sequence of calls that finds first $\lceil \text{corr}(e) \rceil$, storing the pair $(\lceil \text{corr}(e) \rceil, \lceil \text{corr}(e) \rceil)$ in visited. Since the input PES has a finite number of events, FINDEPOMSETS will eventually process $\lceil e \rceil$, which will induce a backward shift operation in line 6. Since visited contains a pair associated with such a configuration, FINDEPOMSETS will record a new elementary cyclic pomset in line 9. Note that the cycle is stored with a pair of configurations, $(\lceil e \rceil \setminus \lceil \text{corr}(e) \rceil, \lceil \text{corr}(e) \rceil)$ in our example, where the first set corresponds with the set of events in the body of the cycle, herein called the *cyclic interval*, and the configuration characterizing the entry point to the cyclic behavior.

However, an elementary cyclic pomset does not always involve a backward shift as it is the case for the PES prefix shown in Fig. 25. In fact, the PES prefix has two elementary cycles but no backward cutoff. The elementary cycles can still be detected, because we store the shifted configuration (along with the unshifted configuration) at every recursive call of the function FINDEPOMSETS. This information can then be used to check if the current shifted configuration has been previously visited, in which case an elementary cyclic pomset is reported. One example of this case is illustrated by the sequence of events shown in the filled blue boxes in the expanded prefix in Fig. 26.

When an elementary cyclic pomset is embedded in a block of concurrency, FINDEPOMSETS will find the cycle multiple times. For instance, for the PES prefix shown in Fig. 28, FINDEPOMSETS will find an elementary cycle comprising the set of events $\{a_1, a_4\}$, when the function processes the configuration $\{a_0, a_1\}$ and the cutoff event a_4. The same cycle will be found when the function processes the configuration $\{a_0, a_1, a_2\}$ and later when it processes the configuration $\{a_0, a_1, a_2, a_3\}$. In order to prevent the recording of multiple copies of the same elementary cyclic pomset, line 8 checks if n_sconf contains the cycle and no other concurrent event. Only when that condition holds, the elementary cyclic pomset is retained.

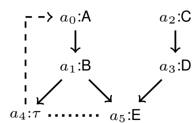

Fig. 28: Cycle within a block of concurrency

The set of elementary pomsets along with the expanded prefix are then used for identifying all the additional model behavior as follows. Traverse the PSP in depth-first search order. At every step, when operation associated with an arc involves an event coming from the model, mark the "event" in the expanded prefix associated with the configuration C^r, with a reference to the operation and the state in the PSP that is reached as an outcome to the operation at hand. Fig. 27 illustrates the result of the previous stage on the example PES presented in Fig. 25. The "events" in the expanded prefix that have been marked are shown with a blue background. Additionally, the information about the state in the PSP and the corresponding operation is shown to the right-hand side of each "event" in the expanded prefix.

In a second stage, we iterate over the set of elementary pomsets to identify those that were not marked. When an elementary pomset is not marked, we find it in the expanded prefix and traverse bottom-up the prefix to find the closest "event" marked with a match operation in the PSP. The state in the PSP associated with this "event" serves to give context to an occurrence of one of the two additional behavior patterns (UNOBSACYCLICINTER and UNOBSCYCLICINTER). Let us consider again the example shown in Fig. 27. If we proceed from left to right, the first elementary pomset is associated with the "event" labeled as $\{a_0, a_1, a_3, a_5, a_9\}$ and its closest match operation occurs at state σ_3 in the PSP. This case corresponds to an elementary acyclic pomset. Part of the pomset has been observed, i.e. the sequence of tasks A and B, and only task D was not observed. The interval of tasks that are not observed can be computed with a set difference. Thus, in this example we will assert a mismatch with the constructor UNOBSACYCLICINTER$(\sigma_3, \{a_5, a_9\})$. Given that a_5 corresponds to an invisible task, one can remove that event from the difference diagnosis. In this case, we can report that in the event log the interval of tasks between a_5 and a_9 is not observed.

The second elementary pomset in the example is associated with the "event" labeled as $\{a_0, a_1, a_3, a_6, a_4, a_7\}$ and its closest match operation occurs at state σ_5 in the PSP. This case corresponds to an unobserved elementary cyclic pomset. The diagnostic will be asserted with the constructor UNOBSCYCLICINTER$(\sigma_5, \{a_3, a_6, a_4, a_7\})$, meaning that the loop formed by tasks B and C is not observed in the log. Note that in this example, there is a loop with two entries and two exits comprising tasks B and C. When expanded, this loop leads to two elementary cyclic pomsets not covered by the PES of the event log and thus two occurrences of the UNOBSCYCLICINTER mismatch pattern

In summary, for the example in Fig. 25 we will assert four mismatches:

Two unobserved elementary acyclic pomsets

- UNOBSACYCLICINTER$(\sigma_3, \{a_5, a_9\})$
- UNOBSACYCLICINTER$(\sigma_9, \{a_8, a_9\})$

Two unobserved elementary cyclic pomsets

- UNOBSCYCLICINTER$(\sigma_5, \{a_3, a_6, a_4, a_7\})$
- UNOBSCYCLICINTER$(\sigma_{11}, \{a_4, a_7, a_3, a_6\})$

Complexity analysis. The complexity of Algorithm 6 is exponential on the size of the PES prefix, due to the fact that so is the number of elementary cycles in a directed graph [31]. In practice, however, given the typical topology of process models (e.g. small number of loops and low edge density) elementary pomsets can be identified quite efficiently. Detecting additional model behavior requires a single traversal of the expanded prefix. Thus, the complexity of this latter step is linear on the size of the expanded prefix.

7.4 Verbalization

The last step in the method is to turn occurrences of mismatch patterns identified using the PSP, into plain natural language statements that can be interpreted by users. Table 2 shows the statements corresponding to each of the nine mismatch patterns defined in Section 7. For some of the mismatch patterns, we identify multiple sub-cases based on conditions on the events of the log and/or model, leading to more than one statement type per mismatch pattern. This depends on whether the events in question are causal or concurrent, or if the event in the log is defined. As a result, the nine constructors give rise to 16 different types of difference statements.

In each statement, the states of the PSP (σ, and when required, also σ') are used to precisely localize where the difference occurs in the log and/or in the model. The text "after σ" means that a difference is observed *immediately* after the occurrence of that state.

8 EVALUATION

We implemented the proposed behavioral alignment method in a standalone Java tool called ProConformance,[11] as well as an OSGi plugin called Compare for the Apromore [32] online process model repository.[12] The tool takes as input a process model in BPMN format and a log in MXML or XES format. Its output is a set of difference statements. The tool allows users to customize the output by switching on/off PSP states, and selecting which elements of a state to show, e.g. only the last matched event.

Using this tool, we conducted a two-pronged (qualitative and quantitative) evaluation of the proposed method, complemented by a user evaluation. First, we performed a qualitative evaluation of the output produced by the method on a real-life event log and a corresponding process model. Next, we performed a quantitative evaluation of time performance and number of produced difference statements, based on large collections of real-life process models. Finally, we measured the perceived ease of use, usefulness and likelihood of using our method, by administering an online survey to process modeling experts. In all three evaluations, we compared our method with *trace alignment*, which, as discussed in Section 2, is a state-of-the-art method in business process conformance checking.

8.1 Qualitative evaluation

For the qualitative evaluation, we used a publicly available log extracted from an information system for managing road traffic fines in Italy [33], and a normative process model, which we derived from the description of this business process in [12]. The normative model in Petri nets is shown in Fig. 29. This traffic fines management process starts when a fine is created. The fine can be paid by the offender right away, after a notification is sent to the offender by the police, or when the offender receives the notification. The payment itself can be done in one or more instalments, depending on the amount of the fine. The case is closed as soon as the payment for the full amount has been done. If the fine is not paid within 180 days, a penalty is charged on top of the fine and if after further 180 days the fine is still due, a credit collection organization will take over the handling of the case. At any time after receiving the notification, the offender can appeal against the fine through a judge or a prefecture. In case of a successful appeal, the case is dismissed and the process ends. If the appeal is unsuccessful, the fine is still to be paid. An appeal can be made more than once, depending on the circumstances (e.g. when escalating the appeal to a higher court).

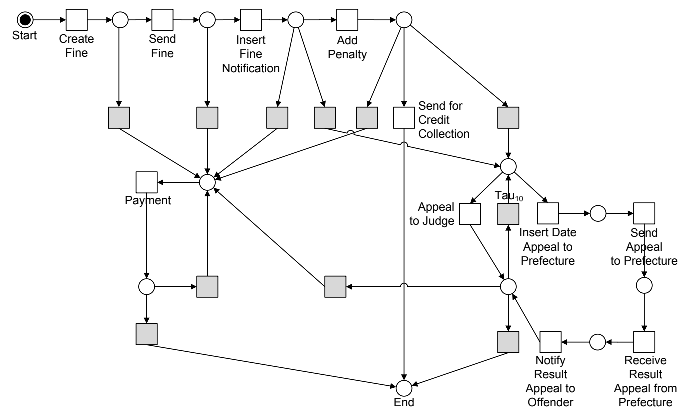

Fig. 29: Traffic fines management process model.

The log covers fines recorded in the period 2000–2013. It contains 150,370 traces comprising 231 distinct traces and a total of 561,470 events.

We assessed the conformance of this log with the Petri net of Fig. 29. Our method produced 15 distinct statements capturing all the differences between the log and the model. As an example, the following statements were retrieved (states are indicated through the last matched event):

1) *In the log, "Send for credit collection" occurs after "Payment" and before the end state*
2) *In the model, after "Insert fine notification", "Add penalty" occurs before "Appeal to judge", while in the log they are concurrent*
3) *In the log, after "Add penalty", "Receive results appeal from prefecture" is substituted by "Appeal to judge"*
4) *In the log, the cycle involving "Insert date appeal to prefecture, Send appeal to prefecture, Receive result appeal from prefecture, Notify result appeal to offender" does not occur after "Insert fine notification".*

Statement 1 (an example of task insertion) denotes a potential compliance issue: credit collection should never occur if the payment has been done, though there are cases in the log where this happens. Similarly, Statement 2 (an example of causality/concurrency mismatch) indicates that there are cases in the log where the penalty is charged

11. Available at http://apromore.org/platform/tools
12. Available at http://apromore.qut.edu.au

Constructor	Condition	Statement type
CausConc(σ, e, f, e', f',coff)	if $e' < e$	*In the log, after σ, $\lambda(e')$ occurs before $\lambda(e)$, while in the model they are concurrent*
	else	*In the model, after σ, $\lambda(f')$ occurs before $\lambda(f)$, while in the log they are concurrent*
Conflict(σ, e, f, e', f',coff)	if $e' \parallel e$	*In the log, after σ, $\lambda(e')$ and $\lambda(e)$ are concurrent, while in the model they are mutually exclusive*
	else if $f' \parallel f$	*In the model, after σ, $\lambda(f')$ and $\lambda(f)$ are concurrent, while in the log they are mutually exclusive*
	else if $e' < e$	*In the log, after σ, $\lambda(e')$ occurs before task $\lambda(e)$, while in the model they are mutually exclusive after σ*
	else	*In the model, after σ, $\lambda(f')$ occurs before $\lambda(f)$, while in the log they are mutually exclusive*
TaskSkip(σ, e, f, e', f',coff)	if $e \neq \perp$	*In the log, after σ, $\lambda(e)$ is optional*
	else	*In the model, after σ, $\lambda(f)$ is optional*
TaskSub(σ, e, f, e', f',coff)		*In the log, after σ, $\lambda(f)$ is substituted by $\lambda(e)$*
UnmRepetition(σ, e, f, e', f',coff)		*In the log, $\lambda(e)$ is repeated after σ*
TaskReloc($\sigma, e, f, \sigma', e', f'$)	if $e \neq \perp$	*In the log, $\lambda(e)$ occurs after σ instead of σ'*
	else	*In the model, $\lambda(f)$ occurs after σ instead of σ'*
TaskAbs(σ, σ', e, f)	if $e \neq \perp$	*In the log, $\lambda(e)$ occurs after σ and before σ'*
	else	*In the model, $\lambda(f)$ occurs after σ and before σ'*
UnobsAcyclicInter(σ, inter)		*In the log,* inter *do(es) not occur after σ*
UnobsCyclicInter(σ, inter)		*In the log, the cycle involving* inter *does not occur after σ*

TABLE 2: Verbalization of mismatch patterns

even after the appeal, while this should be done only if the appeal is unsuccessful. Given that these two events have been observed in any order in the log, they are identified as concurrent. These compliance issues may be related to recording errors in the system (e.g. a payment not being recorded or being recorded for a lower amount).

Statement 3 (an example of task substitution) pinpoints that in the log there are traces where after "Add penalty", event "Receive results appeal from prefecture" is observed. In the PSP, this event in the log is substituted by "Appeal to judge" in the model, after which we know the process can complete. This means that tasks "Insert date appeal to prefecture", "Send appeal to prefecture" and "Notify result appeal to offender", which are in the same path as "Receive results appeal from prefecture" in the model, are not observed in the log. The method substitutes "Receive results appeal from prefecture" with "Appeal to judge" because this minimizes the number of mismatches, as opposed to skipping the three tasks above.[13] This statement suggests that in some cases, the results of an appeal to the prefecture are received by the police, without the appeal having actually been lodged by the offender. This might be due to a mistake at the prefecture (e.g. fines being swapped), which explains why the police does not notify the offender (event "Notify result appeal to offender" is not observed after "Receive results appeal from prefecture").

Finally, Statement 4 (an example of unobserved elementary cycle) indicates that while in principle an offender can appeal to the prefecture multiple times, this has not being observed in the log. Given that the log covers over 10 years of behavior, this may suggest that our model perhaps generalizes the behavior in the log, or that subsequent appeals are never recorded in the system.

For trace alignment, we used the plugins "Replay a Log on Petri Net for All Optimal Alignments"[14] and "Replay a Log on Petri Net for Conformance Analysis"[15] for the ProM 6.5.1 environment. Both plugins report on conformance issues related to fitness, by computing several visual diagnostics as well as a fitness metric (a value from 0 to 1). The former plugin finds *all optimal alignments* for each distinct trace of the log,[16] while the latter provides a good approximation of this result by computing only *one optimal alignment* per distinct trace.

The main diagnostic consists in projecting the results of alignment onto the log, which results in a list of individual trace alignments (a small except of this view for our example is shown in Figure 30). Besides statistics on fitness, this diagnostic shows a great deal of information for each distinct trace, including the exact order in which synchronous moves, (silent) moves on model and moves on log occur, and for each move, the label of the involved event.

The "Replay a Log on Petri Net for Conformance Analysis" plugin, while providing a sub-optimal solution, offers a range of additional diagnostics. For example, one can visualize all trace alignments in a single tabular view and apply various filters on top of it. More interestingly, one can also project the results of alignment onto the normative Petri net (see Fig. 31). This diagnostic can be used to show which model tasks are often skipped (those with a red border), and when tasks that should not be performed according to the model are actually performed according to the log (the darker the color of a path, the more frequent the path is executed in the log). Further, a colored bar at the bottom of a task box shows the ratio between the number of times the task is executed synchronously in the log and in the model (called synchronous move) and the number of times the task is only executed in the model (called move on model).

Although the model view pinpoints, to a certain extent, differences in executions, the exact differences have to be obtained by inspecting the individual misalignments, i.e. those

13. The same holds for task "Send for credit collection" though the substitution considers the lexicographical order of task labels.

14. Parameters used: graph-based state space replay to obtain all optimal alignments with maximum explored states equal to 10,001,000.

15. Parameters used: A^* cost-based fitness express with ILP with maximum explored states equal to 10,001,000.

16. This plugin, with the graph-based state space replay, does not effectively retrieve the complete list of all optimal alignments, but only the *representative* ones as only one possible serialization of each run is returned, and assumes that the model does not have loops made of silent transitions only.

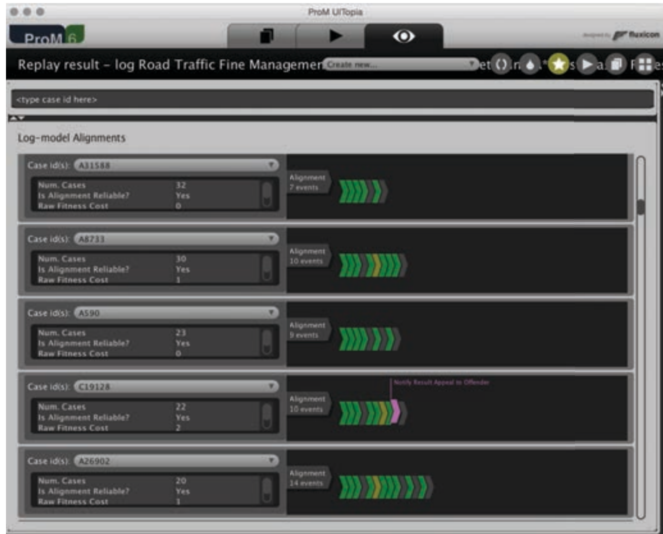

Fig. 30: Excerpt of trace alignments projected on log for the traffic fine management process (green = synchronous move, purple = move on model, grey = silent move on model, yellow = move on log).

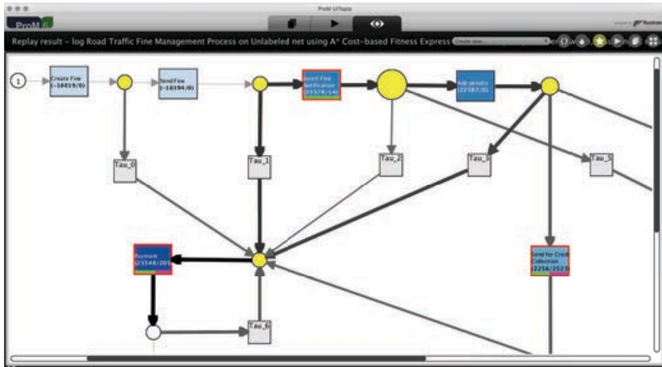

Fig. 31: Trace alignments projected on model for the traffic fine management process.

trace alignments that have at least one move on model or on log.[17] This requires additional analysis. In our example, we need to examine 205 misalignments out of 231 alignments when using one-optimal alignment, and 406 misalignments out of 412 alignments when using all-optimal alignments. Still, differences related to additional model behavior, such as that captured by Statement 4 with our method, cannot be distilled from the misalignments as these only focus on fitness. For this, the underlying technique of the plugin "Check Precision based on Align-ETConformance" could be used, which relies on the prefix automaton built from trace alignments to identify the escaping edges from which the additional model behavior starts, as discussed in Section 2.[18] In our example, however, the escaping edge being reported would be the invisible task Tau_{10}, because this is the last event before the tasks in the interval referred to by Statement 4 can be repeated. From this, by looking at the model, one may infer that there are tasks in the model that can be repeated after Tau_{10}, which are not observed in the log. Similarly, Statement 2 refers to two events being concurrent in the log and causal in the model. This difference cannot be detected by examining the misalignments, because in trace alignment diagnostics are provided at the level of *individual* traces, while concurrency is a behavioral relation that can only be observed *across* traces.

17. Silent moves on model are excluded as they do not capture observable differences.

18. This plugin only provides statistics such as a precision metric. However, one could extract the escaping edges from the code.

8.2 Quantitative evaluation

In order to test the scalability of our approach to increasing model and log complexity, we used two collections of process models: the IBM Business Integration Technology (BIT) library, a publicly-available collection of process models in financial services, telecommunication and other domains, gathered from IBM's consultancy practice [34],[19] and the SAP R/3 collection, the reference model used by SAP to customize their R/3 ERP product, documented in [35].

The BIT collection contains 735 models, while the R/3 collection contains 604 models. We extracted 348 and 494 models, respectively, from these collections, by removing models that were not single-entry single-exit (i.e. models that were not Workflow nets) and that were behaviorally incorrect (i.e. unsound). Next, for each model, we generated an event log using the ProM plugin "Generate Event Log from Petri Net" documented in [36]. This plugin generates a distinct log trace for each possible execution sequence in the model.[20] The tool was only able to parse 274 models from the BIT collection, and 438 models from the R/3 collection, running into out-of-memory exceptions for the remaining models. As such, our quantitative evaluation is based on the logs generated from 712 sound Workflow nets. The statistics on these models are provided in Table 3. The models range from simple ones, with a minimum of 7 nodes and a small number of XOR and AND splits, to very large and complex models with up to 177 nodes and a large number of XOR and AND splits with many outgoing arcs.

Collection		Min	Max	Mean	σ
BIT	Size	7	177	38.1	30.08
	# XOR splits	0	5	0.61	1
	Outdegree XOR	2	10	2.42	1.18
	# AND splits	0	33	6.49	5.72
	Outdegree AND	2	7	2.08	0.42
R/3	Size	7	85	27.62	17.78
	# XOR splits	0	4	0.59	0.82
	Outdegree XOR	2	8	2.48	0.94
	# AND splits	0	4	0.83	0.86
	Outdegree AND	2	5	2.29	0.61

TABLE 3: Statistics on model complexity.

Next, in order to create random differences between each log and its corresponding model, we injected noise in each original log. We achieved this by repeatedly adding or removing a random event that already existed in the original log in a random position of a randomly selected trace, until the number of added and removed events equals a percentage of the total number of events in the original log.

19. The BIT collection is available at http://apromore.qut.edu.au

20. Parameters used: simulation method: complete generation; min./max. traces to add for each generated sequence: 1; max. times marking seen: 2; only include traces that reach end state; only include traces without remaining tokens.

We applied four noise levels, corresponding to 5%, 10%, 15% and 20% of the total number of events, thus obtaining four "noisy" variants for each original log. The noise injection procedure is inspired by the technique documented in [37]. Before performing this operation, we duplicated each distinct trace in every original log, so that each distinct execution sequence in the corresponding model is represented twice in the log. We did so in order not to increase the total number of traces in the log when injecting noise.

Table 4 provides statistics on the complexity of the logs for both collections, divided by noise level. The logs range from 6 to 1,433 total events, with a maximum of 38 traces (having 29 distinct traces on average) in the case of the BIT collection, and from 6 to 9,462 total events, with a maximum of 840 traces (38 distinct traces on average) for the R/3 collection. In the remainder, with *total log size* we refer to the total number of events, which corresponds to the sum of the lengths of the traces.

Collection	Noise	Min	Max	Mean	σ
BIT	None	6	1,432	58	120
	5%	6	1,427	58	120
	10%	7	1,433	58	120
	15%	7	1,428	57	120
	20%	7	1,428	57	120
R/3	None	6	9,408	515	1,377
	5%	6	9,410	514	1,377
	10%	7	9,432	515	1,379
	15%	7	9,449	515	1,382
	20%	7	9,462	515	1,384

TABLE 4: Total log size in terms of number of events.

Using these two logsets, we measured the execution time and counted the number of statements provided by our method for each model-log pair for all noise levels. We performed the tests on a machine with a dual core Intel Core i7-4710HQ 2.5GHz (4 cores), 16GB RAM, running Windows 8.1 x64 and Java 1.8.0_31 with 14GB of allocated memory. To eliminate load time from the measures, we executed each test five times and recorded average times of three executions, removing the fastest and the slowest executions.

The execution times against the total log size for each noise level are plotted in Figures 32 (BIT) and 33 (R/3), where the measuring points are color-coded depending on the level of noise. For the BIT case (Figure 32), we have omitted the five outlier logs with more than 1,400 events, in order to improve clarity. Summary statistics of the time performance of each operation (log to PES, model to PES, PSP computation and verbalization), as well as of the total time taken by our method are shown in Table 5.

Without noise, the execution time is linear on the log size (R^2 = 0.83 for BIT and 0.95 for R/3), reaching a peak of 67ms for a log of 1,432 events (BIT) and 183ms for a log of 8,352 events (R/3). The plots show that the discrepancies between the model and the log result in higher execution times than the case without noise. This is due to the complexity of building the PSP. Still, execution times are always under 10sec for BIT (6.6sec max at 20% noise on a log of 1,428 events and a model of 177 nodes, with a peak memory heap space of 894MB) and under 2min for R/3 (92sec max at 20% noise on a log of 8,352 events and a model of 68 nodes, with a peak memory heap space of 1GB).[21]

21. The memory tests were performed using YourKit Java Profiler version 2016.02-b40.

From Table 5 we can observe that the increase of noise (and, hence, the increase in the number of differences) results in an increased execution time for generating the PSP and for verbalizing the differences. In particular, higher PSP computation times correspond to logs with a larger number of distinct traces. However, the execution time required to convert the log to PES and the model to PES only depends on the size of the log and the model, respectively, and are thus not affected by the increase of noise.

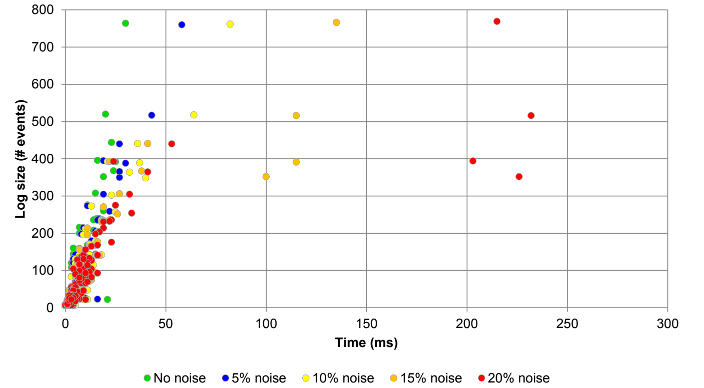

Fig. 32: Effect of log size and noise on the time performance of our method (BIT).

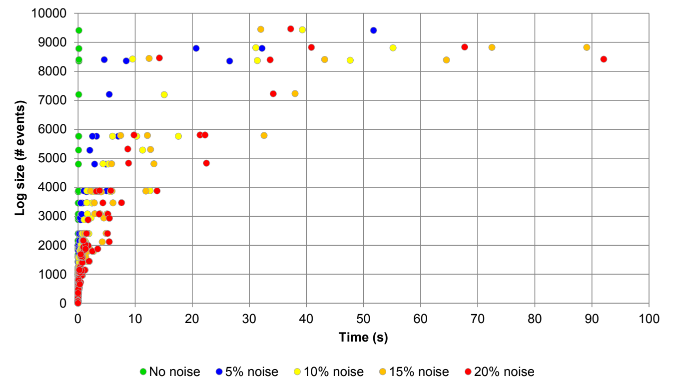

Fig. 33: Effect of log size and noise on the time performance of our method (R/3).

Table 6 reports the number of statements produced by our method for each logset. As expected, the noiseless logs all produce zero statements. In the extreme case of a log with 20% noise, 104 statements were required to describe all the differences for the BIT collection (with a log of 1,428 events and model of 177 nodes – this is the pair that took 6.6sec to be compared), and 593 statements for the R/3 collection (with a log of 8,352 events and model of 68 nodes).

Comparing execution times with number of statements, we can observe a relatively sharp increase in average execution time between 0% and 5% noise in the R/3 dataset (11ms vs 467ms). This, however, coincides with a similarly sharp increase in the amount of produced statements (0 vs 18 on average and 7,469 in total). For example, the number of additional statements required for the 10% noise level compared to the 5% noise level is smaller, resulting in a similarly smaller increase in required execution time.

Collection	Noise	Operation	Min	Max	Mean	95%	σ
BIT	None	Log to PES	0	24	1	4	2
		Model to PES	0	34	2	8	3
		PSP	0	9	0	0	1
		Verbalization	0	4	0	0	0
		Total	**0**	**67**	**4**	**14**	**6**
	5%	Log to PES	0	28	1	4	2
		Model to PES	0	34	2	8	3
		PSP	0	78	1	2	5
		Verbalization	0	62	0	1	4
		Total	**0**	**141**	**5**	**15**	**10**
	10%	Log to PES	0	29	1	4	2
		Model to PES	0	34	2	8	3
		PSP	0	856	4	3	52
		Verbalization	0	88	1	2	5
		Total	**0**	**919**	**8**	**17**	**56**
	15%	Log to PES	0	31	1	4	2
		Model to PES	0	34	2	8	3
		PSP	0	2,547	11	4	154
		Verbalization	0	99	1	2	6
		Total	**0**	**2,615**	**15**	**17**	**158**
	20%	Log to PES	0	31	1	4	2
		Model to PES	0	34	2	8	3
		PSP	0	6,592	28	7	399
		Verbalization	0	108	1	2	7
		Total	**0**	**6,659**	**32**	**20**	**403**
SAP	None	Log to PES	0	160	9	50	25
		Model to PES	0	36	2	8	4
		PSP	0	8	0	0	1
		Verbalization	0	2	0	0	0
		Total	**0**	**183**	**11**	**60**	**28**
	5%	Log to PES	0	192	9	55	27
		Model to PES	0	23	2	7	3
		PSP	0	51,562	455	596	3,431
		Verbalization	0	1,194	20	66	111
		Total	**0**	**51,768**	**467**	**656**	**3,452**
	10%	Log to PES	0	197	10	58	28
		Model to PES	0	26	2	7	3
		PSP	0	54,985	811	2,031	4,779
		Verbalization	0	2,217	53	199	260
		Total	**0**	**55,196**	**824**	**2,072**	**4,806**
	15%	Log to PES	0	204	10	62	29
		Model to PES	0	23	2	7	3
		PSP	0	88,889	1,224	3,687	7,383
		Verbalization	0	2,990	75	282	349
		Total	**0**	**89,124**	**1,238**	**3,759**	**7,410**
	20%	Log to PES	0	225	10	63	30
		Model to PES	0	23	2	7	3
		PSP	0	91,874	1,213	3,716	6,886
		Verbalization	0	3,490	89	341	410
		Total	**0**	**92,102**	**1,226**	**3,798**	**6,915**

TABLE 5: Execution time (ms) of each operation for the two logsets, using our method.

Collection	Noise	Min	Max	Mean	σ	Total
BIT	None	0	0	0	0	0
	5%	0	42	2	4	480
	10%	0	65	4	7	943
	15%	0	89	5	9	1,409
	20%	0	104	7	11	17,76
R/3	None	0	0	0	0	0
	5%	0	370	18	43	7,469
	10%	0	408	28	62	11,720
	15%	0	544	37	79	15,532
	20%	0	593	43	88	18,048

TABLE 6: Statements produced for each logset.

Next, we carried out the same tests using trace alignment with all optimal alignments (using the plugins "Replay a Log on Petri Net for All Optimal Alignments" to obtain the individual trace alignments, and "Check Precision based on Align-ETConformance"[22]) to obtain the escaping edges. The execution times against log size and noise level are reported in Figures 34 (BIT) and 35 (R/3). Again, we have omitted the five outlier logs with more than 1,400 events for the BIT case (Figure 34), in order to improve clarity. Table 7 provides the summary statistics. Similar to our method, performances grow linearly on the log size in the case of no noise (R^2=0.87 for BIT and 0.93 for R/3). Comparatively, trace alignment is faster than our method, reaching a peak of just 0.5sec for BIT and 1.9sec for R/3 in the case of 20% noise, compared to 6.6sec and 92sec, respectively, with our method. In particular, the difference in performance is more evident for larger logs with many model-log discrepancies, where the complexity of the PSP is exposed. Conversely, on simpler logs our method tends to be slightly faster than trace alignment. This is the case for the model-log pairs with up to 5% noise levels for the BIT collection, and the model-log pairs with no noise for the R/3 collection.

22. Parameters used: ordered representation; all optimal alignments.

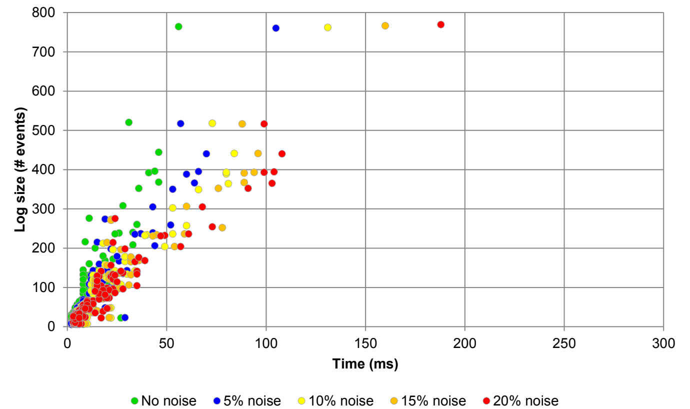

Fig. 34: Effect of log size and noise on the time performance of all optimal alignments (BIT).

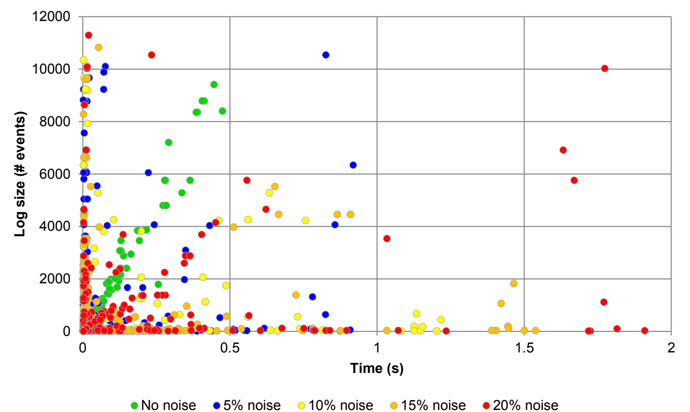

Fig. 35: Effect of log size and noise on the time performance of all optimal alignments (R/3).

Tables 8 and 9 report the number of misalignments and escaping edges. From these we can observe that the number of diagnostics provided by trace alignment is significantly higher than that reported by our method, with a total of 6,968 misalignments + escaping edges for the BIT collection and 153,698 misalignments + escaping edges for the R/3 collection (summing up across all noise levels), compared to a total of 4,608 statements and 52,769 statements, respectively, with our method.

Trace alignment reports escaping edges also in the case of logs with no noise (316 for BIT and 2,495 for R/3). This is due to the fact that these edges are detected whenever there

is repetitive behavior (i.e. infinite behavior) in the model, since the log records finite behavior. For example, if a loop in the model is only observed twice in the log, an escaping edge will be reported on the state enabling the third iteration of this loop in the model.

Collection	Noise	Min	Max	Mean	95%	σ
BIT	None	2	102	7	25	10
	5%	2	221	10	31	18
	10%	3	317	11	36	24
	15%	3	414	12	38	30
	20%	3	522	13	42	37
R/3	None	2	475	29	148	70
	5%	2	919	55	290	143
	10%	2	1,218	72	406	193
	15%	2	1,539	86	486	239
	20%	2	1,910	98	553	282

TABLE 7: Execution time (ms) for each logset using all optimal alignments.

Collection	Noise	Min	Max	Mean	σ	Total
BIT	None	0	0	0	0	0
	5%	0	36	3	4	759
	10%	1	61	4	6	1,194
	15%	1	69	6	8	1,567
	20%	1	94	7	10	1,864
R/3	None	0	0	0	0	0
	5%	0	826	35	96	14,287
	10%	1	1,689	65	180	26,285
	15%	1	2,554	93	267	37,271
	20%	1	3,711	121	358	48,725

TABLE 8: Misalignments for each logset using all optimal alignments.

Collection	Noise	Min	Max	Mean	σ	Total
BIT	None	0	8	1	1	316
	5%	0	8	1	1	315
	10%	0	7	1	1	317
	15%	0	7	1	1	316
	20%	0	7	1	1	320
R/3	None	0	494	6	32	2,495
	5%	0	1,246	14	78	5,748
	10%	0	1,533	16	92	6,266
	15%	0	1,795	16	102	6,405
	20%	0	1,821	16	102	6,216

TABLE 9: Escaping edges for each logset using all optimal alignments.

8.3 User evaluation

As a final step, we performed a user evaluation of the proposed method to improve the ecological validity of our findings. We have shown that our method is more accurate and concise than trace alignment, but its potential can only materialize if people want to use it. Thus, we designed an online survey instrument that presented a simple process model and assumed that this model was accompanied by a log recording 53 executions of the corresponding business process. The model, a Petri net with 31 nodes (10 visible transitions), was created from a real-life claims handling process model, whose labels were anonymized to avoid domain bias in the respondents. We showed a Petri net rather than a BPMN model to be consistent with the visualization produced by trace alignment.

Next, we showed the output of the trace alignment method in the form of a set of 32 misalignments and a Petri net with alignment information overlaid, as well as the output of our method in the form of a list of nine statements. For trace alignment, we used one-optimal instead of all-optimal alignment to reduce the number of misalignments showed as output (from over 120 alignments to 32), so as to facilitate the analysis of this output by the respondents in the context of a short survey. The complete instrument is reported in the Appendix.

Using this survey, we asked the respondents to compare both methods across the main constructs of the Technology Acceptance Model [39]. The Technology Acceptance Model is a widely used model in information systems research to evaluate the likelihood that people will use a certain information technology. It asserts that people who find a given technology easy to use will find it more useful, and in turn if they find it both easy to use and useful they will be more likely to use the technology. Accordingly, we measured each of those constructs by asking respondents to indicate which of the methods they perceived as easier to use, more useful, and more likely to be used for i) checking the conformance of event logs to process models, and ii) identifying differences between process models and event logs. This led to the following six questions:

- “What is the easiest approach for checking the conformance of an event log to a process model?”
- “What is the easiest approach for identifying the differences between a process model and an event log?”
- “What is the most useful approach for checking the conformance of an event log to a process model?”
- “What is the most useful approach for identifying the differences between a process model and an event log?”
- “Which approach would you likely use for checking the conformance of an event log to a process model?”
- “Which approach would you likely use for identifying the differences between a process model and an event log?”.

For each question, we used a seven point Likert-type response scale ranging from “Strongly prefer Trace Alignment” to “Strongly prefer Behavioral Alignment”.

The type of comparison we chose in our experiment may have favored trace alignment since visual representation is more appealing and informative than plain text. In the absence of comparable feedback, we decided to favor trace alignment, rather than our method, for which we were trying to find support. This minimizes the risk for a Type I error in statistics, i.e. that of erroneously rejecting the null hypothesis, which in our case was that both methods are equally good for conformance checking.

Respondents were also asked to share their general occupation (academic vs. professional); their experience in business process modelling, including questions about the frequency of creating and analyzing process models in the past twelve months, as well as the complexity of those models; their familiarity with and confidence and competence in working with Petri nets (evaluated on 7-point Likert-type response scales ranging from “strongly disagree” to “strongly agree”); and the amount of training and self-education in process modeling they had engaged in over the past twelve months (in number of days).

We hypothesized that on average:

1) respondents would have a preference for behavioral alignment,

	Full sample (n = 71)			Academics (n = 38)			Professionals (n = 33)			p-value difference academics vs. professionals
	Mean	StDev	Min-Max	Mean	StDev	Min-Max	Mean	StDev	Min-Max	
Experience (yrs)	8.11	6.65	0–30	7.97	5.35	0–25	8.27	7.80	0–30	.57
Familiarity Petri nets[a]	4.97	1.98	1–7	5.89	1.35	1–7	3.91	2.07	1–7	.00***
Confidence Petri nets[a]	5.15	1.75	1–7	6.03	1.13	3–7	4.15	1.82	1–7	.00***
Competence Petri nets[a]	4.61	2.13	1–7	5.68	1.60	1–7	3.36	2.00	1–7	.00***
Models analyzed (nr.)	86.11	107.28	0–500	119.00	125.05	4–500	48.24	65.84	0–300	.00***
Models created (nr.)	25.17	34.91	0–250	33.47	44.31	4–250	15.61	14.83	0–70	.01*
Activities in models (nr.)	19.58	12.37	0–70	20.50	12.52	6–70	18.52	12.30	0–50	.77
Training (days)	1.90	4.16	0–30	0.92	1.84	0–7	3.03	5.61	0–30	.02*
Self-education (days)	31.25	69.14	0–365	43.63	91.38	0–365	17.00	20.26	0–100	.26

[a] Questions that were rated on a 7-point Likert-type response scale with high scores representing high familiarity, confidence and competence, respectively.
* $p < .05$; *** $p < .001$ For these variables, the null hypothesis that distributions for academics and professionals were equal was rejected using a Mann-Whitney U test; academics scored significantly higher on each of these, except training, where professionals scored significantly higher than academics.

TABLE 10: Summary statistics for the full sample and for academics and professionals separately.

	Median			Mann-Whitney U comparing Academics and Professionals		
	Full sample (n = 71)	Academics (n = 38)	Professionals (n = 33)	U	p	r
Ease of use for checking differences	0.00	-0.50	2.00	414.00	.01	-0.30
Ease of use for checking conformity	0.00	-1.00	2.00	376.50	.00	-0.35
Useful for checking differences	0.00	-1.00	1.00	427.00	.02	-0.28
Useful for checking conformity	0.00	-1.00	1.00	415.50	.01	-0.29
Likely to use for checking differences	0.00	-1.00	1.00	396.50	.01	-0.32
Likely to use for checking conformity	0.00	-1.00	1.00	416.00	.01	-0.29

Note. Variables were measured on a 7-point Likert-type response scale ranging from (-3) "Strongly prefer Trace Alignment" to (0) "Neutral" to (3) "Strongly prefer Behavioral Alignment". U denotes the Mann-Whitney U test result; r refers to an effect size estimation that is considered medium at .3 and high above .5 [38].

TABLE 11: Medians and Mann-Whitney U test results for perceived ease of use, usefulness and likelihood of use.

	Ease of use for checking		Useful for checking		Likely to use for checking	
	differences	conformance	differences	conformance	differences	conformance
Experience (years)	-0.13	-0.08	-0.01	-0.06	-0.1	-0.02
Familiarity Petri nets	-0.24*	-0.22	-0.2	-0.27*	-0.13	-0.30*
Confidence Petri nets	-0.29*	-0.27*	-0.26*	-0.35**	-0.22	-0.34**
Competence Petri nets	-0.25*	-0.23	-0.22	-0.28*	-0.19	-0.27*
Models analyzed (nr.)	-0.07	-0.07	0	-0.01	0.08	-0.11
Models created (nr.)	0.13	0.11	0.16	0.1	0.16	0.11
Activities in models (nr.)	0.15	0.26*	0.18	0.18	0.32**	0.2
Training (days)	0.26*	0.234*	0.26*	0.27*	0.32**	0.24*
Self-education (days)	0.19	0.23*	0.21	0.17	0.23	0.27*

* $p < .05$; ** $p < .01$
Note. Low values in the dependent variables represent a preference for alignments, while high values represent a preference for behavioral alignment.

TABLE 12: Spearman Correlations between dependent variables and indicators of experience and expertise.

2) respondents with less experience, familiarity, confidence and competence in the use of Petri nets would have a stronger preference for behavioral alignment.

Invitations to complete the survey were distributed via the authors' professional network, targeting academics and practitioners. The survey was open for one month. During this period, responses were received from 38 academics and 33 professionals. Table 10 presents descriptive statistics for the entire sample and for the two cohorts separately.

Because the dependent variables were not normally distributed (established based on visual inspection and the Shapiro-Wilks test at $p < .01$) we used non-parametric tests: we used Spearman correlations to establish relations, and medians and Mann-Whitney U tests to compare groups. As expected, these group comparisons revealed that academics tended to rate themselves as more familiar with and confident and competent in working with Petri nets, and that they had analysed and created more models in the past twelve months than professionals. Yet, they received less training. The latter result is also expected, as training is typically addressing a professional audience, e.g. in the form of continuing professional development courses, while academics rely on self-education (as confirmed by the results in Table 10), especially because the questions referred to the last twelve months. Based on this, we tested Hypothesis 1 both for the whole sample and for the two cohorts separately. The results do not support this hypothesis for the whole sample, as there is no general preference for our method: the median was zero ("neutral") in the full sample (see Table 11). However, professionals did show a preference for behavioral alignment (especially along ease of use) while academics preferred trace alignment, so Hypothesis 1 is supported for the professionals cohort only.

Further exploration of the data (in Table 12) revealed that respondents with more experience, familiarity, confidence and competence in working with Petri nets tended to have a stronger preference towards trace alignment. These results were in support of Hypothesis 2.

In summary, our user evaluation revealed that academics prefer trace alignment while professionals prefer our method. Moreover, people that possess less expertise in Petri nets have a stronger preference for our method.

9 Conclusion

We proposed a method called behavioral alignment, to check the conformance between an event log capturing the actual execution of a business process, and a model capturing its expected or normative execution. The method relies on a unified representation of event logs and process models. Specifically, the log is folded into an event structure and the model is unfolded into another event structure. The two structures are then aligned, and their commonalities and divergences are represented via a partially synchronized product from which a complete set of behavioral differences between the model and the log is extracted.

We compared the proposed method to existing conformance checking methods based on trace alignment using a three-pronged evaluation method. First, a qualitative evaluation based on a real-life event log and a corresponding process model showed that the presented method produces a more compact, yet much more understandable diagnosis than conformance checking methods based on trace alignment. This qualitative evaluation also showed that the proposed method exposes behavioral differences that are difficult or impossible to identify using trace alignment.

Second, a quantitative evaluation using two real-life collections with over 700 process models in total, showed that the proposed method – while being generally slower than trace alignment – has reasonable execution times (within 10 seconds). In extreme cases involving logs with over 8,000 event occurrences (considering distinct traces only) and a high number of differences between the process model and the event log, the execution time is still below 2 minutes. The quantitative evaluation also showed that the proposed method consistently produces more compact difference diagnosis than trace alignment methods.

Third, we conducted a user evaluation of our method relative to trace alignment by means of a survey filled in by a population of researchers and professionals in BPM. The survey results showed that while researchers have a preference for trace alignment, professionals find the proposed method more useful, easier to use, and will use it more likely than alignment. Moreover, a correlation was found between lack of expertise in Petri nets and preference for our method. However, a more in-depth user evaluation, e.g. with active user participation, would be required to fully understand the advantages and disadvantages of each method. This is a direction for future work.

A limitation of the proposed method is that it treats the input log as consisting of sequences of event labels, thereby ignoring timestamps and event payloads. Possible directions for future work include designing temporal and data-aware extensions of the method, along the lines of data-aware extensions of trace alignment methods [12].

In addition to these possible extensions, there are also multiple directions to improve the proposed method. First, the method relies on a concurrency oracle when transforming sets of traces into sets of partially-ordered runs. In the empirical evaluation, we relied on a relatively simple concurrency oracle, namely the $\alpha+$ oracle. This oracle has the limitation that it sometimes cannot isolate concurrency in the presence of short loops, skipped and/or duplicated tasks. Accordingly, another direction for future work is to evaluate the performance of the proposed method with a range of more sophisticated concurrency oracles. A more accurate concurrency oracle can lead to a more accurate transformation from traces to runs, which in turn would lead to an event structure that better reflects the log.

Another direction for future work is to design a technique to summarize the diagnosis of the differences, for example by grouping related difference statements and abstracting them via higher-level statements that strike a tradeoff between accuracy and interpretability. In a similar vein, there is room for improving the interpretability of the results by complementing the generated natural language statements with visual feedback. The latter can be achieved for example by visually displaying the state in the input process model where the difference occurs as well as the involved tasks. Also, since the PSP tells us the exact configurations in the PES of the model and in the PES of the log where each difference occurs, we can extract two incomplete runs (one in the process model and one in the log), leading to the state where the difference in question is observed. These runs can be trivially serialized into traces and used to show to the user the possible sequences of events leading to a state where the difference occurs.

Acknowledgements. We thank Jorge Muñoz-Gama for his support with the conformance checking ProM plugin, Boudewijn van Dongen for his insights on trace alignment, and Abel Armas-Cervantes for his comments on early versions of this paper. This research is funded by the Australian Research Council Discovery Project DP150103356 and the Estonian Research Council Project IUT20-55.

References

[1] W. M. P. van der Aalst, *Process Mining - Discovery, Conformance and Enhancement of Business Processes*. Springer, 2011.

[2] J. C. A. M. Buijs, M. L. Rosa, H. A. Reijers, B. F. van Dongen, and W. M. P. van der Aalst, "Improving business process models using observed behavior," in *Data-Driven Process Discovery and Analysis - Second IFIP WG 2.6, 2.12 International Symposium, SIMPDA 2012, Campione d'Italia, Italy, June 18-20, 2012, Revised Selected Papers*, ser. Lecture Notes in Business Information Processing, vol. 162. Springer, 2013.

[3] D. Fahland and W. M. P. van der Aalst, "Model repair - aligning process models to reality," *Inf. Syst.*, vol. 47, 2015.

[4] N. Kleiner, "Delta analysis with workflow logs: aligning business process prescriptions and their reality," *Requir. Eng.*, vol. 10, no. 3, 2005.

[5] M. Nielsen, G. D. Plotkin, and G. Winskel, "Petri nets, event structures and domains, part I," *Theor. Comput. Sci.*, vol. 13, 1981.

[6] W. van der Aalst, A. ter Hofstede, B. Kiepuszewski, and A. Barros, "Workflow Patterns," *Distributed and Parallel Databases*, vol. 14, no. 1, 2003.

[7] A. Rozinat and W. van der Aalst, "Conformance checking of processes based on monitoring real behavior," *Inf. Syst.*, vol. 33, no. 1, 2008.

[8] A. A. de Medeiros, "Genetic process mining," Ph.D. dissertation, Eindhoven University of Technology, 2006.

[9] S. K. L. M. vanden Broucke, J. Munoz-Gama, J. Carmona, B. Baesens, and J. Vanthienen, "Event-based real-time decomposed conformance analysis," in *Proc. of OTM Conferences*. Springer, 2014.

[10] J. Munoz-Gama, J. Carmona, and W. M. P. van der Aalst, “Single-entry single-exit decomposed conformance checking,” *Inf. Syst.*, vol. 46, 2014.
[11] A. Adriansyah, B. van Dongen, and W. van der Aalst, “Conformance checking using cost-based fitness analysis,” in *Proc. of EDOC*. IEEE, 2011.
[12] F. Mannhardt, M. de Leoni, H. A. Reijers, and W. M. van der Aalst, “Balanced multi-perspective checking of process conformance,” *Computing*, 2015.
[13] J. D. Weerdt, M. D. Backer, J. Vanthienen, and B. Baesens, “A robust f-measure for evaluating discovered process models,” in *Proceedings of the CIDM 2011*, 2011.
[14] S. K. L. M. vanden Broucke, J. D. Weerdt, J. Vanthienen, and B. Baesens, “Determining process model precision and generalization with weighted artificial negative events,” *IEEE Trans. Knowl. Data Eng.*, vol. 26, no. 8, 2014.
[15] J. Munoz-Gama and J. Carmona, “A fresh look at precision in process conformance,” in *Proc. of BPM*. Springer, 2010.
[16] A. Adriansyah, J. Munoz-Gama, J. Carmona, B. F. van Dongen, and W. M. P. van der Aalst, “Measuring precision of modeled behavior,” *Inf. Syst. E-Business Management*, vol. 13, no. 1, 2015.
[17] A. Armas-Cervantes, P. Baldan, M. Dumas, and L. García-Bañuelos, “Diagnosing behavioral differences between business process models: An approach based on event structures,” *Information Systems*, vol. 56, 2016.
[18] N. van Beest, M. Dumas, L. García-Bañuelos, and M. L. Rosa, “Log delta analysis: Interpretable differencing of business process event logs,” in *Proc. of BPM 2015*. Springer, 2015.
[19] L. García-Bañuelos, N. van Beest, M. Dumas, and M. L. Rosa, “Business process conformance checking based on event structures,” in *Proc. of NWPT'2015*. Reykjavik University, 2015, 3 pages.
[20] R. M. Dijkman, M. Dumas, and C. Ouyang, “Semantics and analysis of business process models in BPMN,” *Information & Software Technology*, vol. 50, no. 12, 2008.
[21] K. L. McMillan, “A technique of state space search based on unfolding,” *Formal Methods in System Design*, vol. 6, no. 1, 1995.
[22] J. Esparza, S. Römer, and W. Vogler, “An improvement of mcmillan's unfolding algorithm,” *Formal Methods in System Design*, vol. 20, no. 3, 2002.
[23] J. Esparza, “Model checking using net unfoldings,” *Sci. Comput. Program.*, vol. 23, no. 2-3, 1994.
[24] J. E. Cook and A. L. Wolf, “Event-based detection of concurrency,” in *FSE*. ACM, 1998.
[25] W. M. P. van der Aalst, T. Weijters, and L. Maruster, “Workflow mining: discovering process models from event logs,” *IEEE TKDE*, vol. 16, no. 9, 2004.
[26] A. K. Alves de Medeiros, W. M. P. van der Aalst, and A. J. M. M. Weijters, “Workflow mining: Current status and future directions,” in *Proc. of On The Move to Meaningful Internet Systems (OTM) 2003*, 2003.
[27] P. E. Hart, N. J. Nilsson, and B. Raphael, “A formal basis for the heuristic determination of minimum cost paths,” *IEEE Trans. Systems Science and Cybernetics*, vol. 4, no. 2, 1968.
[28] V. Pratt, “Modeling concurrency with partial orders,” *International Journal of Parallel Programming*, vol. 15, no. 1, 1986.
[29] A. Valmari, “A stubborn attack on state explosion,” *Formal Methods in System Design*, vol. 1, no. 4, 1992.
[30] J. C. Tiernan, “An Efficient Search Algorithm to Find the Elementary Circuits of a Graph,” *Commun. ACM*, vol. 13, no. 12, 1970.
[31] J. L. Szwarcfiter and P. E. Lauer, “A search strategy for the elementary cycles of a directed graph,” *BIT Numerical Mathematics*, vol. 16, no. 2, 1976.
[32] M. L. Rosa, H. A. Reijers, W. M. P. van der Aalst, R. M. Dijkman, J. Mendling, M. Dumas, and L. García-Bañuelos, “APROMORE: an advanced process model repository,” *Expert Syst. Appl.*, vol. 38, no. 6, 2011.
[33] M. de Leoni and F. Mannhardt, “Road traffic fine management process,” 2015. [Online]. Available: http://dx.doi.org/10.4121/uuid:270fd440-1057-4fb9-89a9-b699b47990f5
[34] D. Fahland, C. Favre, J. Koehler, N. Lohmann, H. Völzer, and K. Wolf, “Analysis on demand: Instantaneous soundness checking of industrial business process models,” *Data Knowl. Eng.*, vol. 70, no. 5, 2011.
[35] T. Curran and G. Keller, *SAP R/3 Business Blueprint: Understanding the Business Process Reference Model*. Upper Saddle River, 1997.
[36] S. Vanden Broucke, J. De Weerdt, J. Vanthienen, and B. Baesens, “An improved process event log artificial negative event generator,” Faculty of Economics and Business, KU Leuven (Belgium), Tech. Rep. KBI_1216, 2012.
[37] R. Conforti, M. Dumas, L. García-Bañuelos, and M. La Rosa, “BPMN Miner: Automated discovery of BPMN process models with hierarchical structure,” *Information Systems*, vol. 56, 2016.
[38] A. Field, *Discovering statistics using SPSS*. Sage publications, 2009.
[39] F. D. Davis, “Perceived usefulness, perceived ease of use, and user acceptance of information technology,” *MIS quarterly*, 1989.

Luciano García-Bañuelos is Associate Professor of Software Engineering at the University of Tartu, Estonia. He obtained his PhD in 2003 from Grenoble Institute of Technology for his work on long-running transactions. His current research interests are in the fields of service-oriented computing and business process management, with a focus on formal methods for business process modeling and analysis.

Nick van Beest is a Research Scientist at Data61, CSIRO in Brisbane, Australia. He obtained his PhD in Information Systems in 2013 at the University of Groningen, The Netherlands. His research experience covers artificial intelligence, process mining, business process compliance and knowledge-intensive business processes. He currently works on deviance mining and conformance checking for the purpose of performance improvement and automated runtime anticipation of disruptions.

Marlon Dumas is Professor of Software Engineering at University of Tartu, Estonia and Adjunct Professor of Information Systems at Queensland University of Technology, Australia. His interests span across software engineering, information systems and Business Process Management. His research focuses on combining data mining and formal methods to analyze and monitor business processes. He is co-author of the textbook Fundamentals of Business Process Management (Springer, 2013).

Marcello La Rosa is Professor of Information Systems at the Queensland University of Technology, Brisbane, Australia. His research interests include business process consolidation, mining and automation. He leads the Apromore initiative (www.apromore.org), a strategic collaboration between various universities for the development of an advanced process analytics platform. He is co-author of the textbook Fundamentals of Business Process Management (Springer, 2013).

Willem Mertens is a Postdoctoral Research Fellow at the Queensland University of Technology, Australia, and a Research Fellow of Vlerick Business School, Belgium. His primary topic of interest is positive deviance: behavior that deviates from organisational norms or processes, and is more successful because of it. His other research interests span the domains of business process management, organizational behavior and innovation.

Appendix

The Appendix for this paper is available online.