

Verification of data-aware workflows via reachability: formalisation and experiments

Riccardo De Masellis

Stockholm University, 10691 Stockholm, Sweden

RICCARDO.DEMASELLIS@PHILOSOPHY.SU.SE

Chiara Di Francescomarino

Chiara Ghidini

*FBK-IRST, via Sommarive, 18
38123 Trento, Italy*

DFMCHIARA@FBK.EU

GHIDINI@FBK.EU

Sergio Tessaris

*Free University of Bozen-Bolzano, Piazza Domenicani, 3
39100 Bozen-Bolzano, Italy*

TESSARIS@INF.UNIBZ.IT

Abstract

The growing adoption of IT-systems for the modelling and execution of (business) processes or services has thrust the scientific investigation towards techniques and tools which support complex forms of process analysis. These techniques, nowadays grouped under the umbrella of the Process Mining research area, typically rely on observation of past (tracked and logged) process executions. A first important limitation in this field is the fact that the majority of the techniques paired with concrete tool support only consider activities, but lack the ability to take into account the *data objects* manipulated by these activities. Second, Process Mining techniques mainly rely on *complete* observations of terminated process executions. In many real cases, however, only incomplete log information is available.

This paper tackles these two shortcomings by proposing an approach to exploit reachability to reason on imperative data-aware process models and possibly incomplete process executions. The contribution of this paper is twofold: first, it formulates the trace completion as a reachability problem over data-aware models and second, it provides a rigorous mapping between our data-aware models and three important paradigms for reasoning about dynamic systems, namely Action Languages, Classical Planning, and Model-Checking. This allows us to exploit and extensively evaluate the available tools for the above paradigms to solve the trace repair problem.

The rigorous encoding of our data-aware models, based on a common interpretation of the semantics of Action Languages, Classical Planning, and Model-Checking in terms of transition systems, paired with a first comprehensive assessment of the performances of their tools in computing reachability for data-aware workflow net languages, provide a solid contribution to advancing the state-of-the-art on the concrete exploitation of formal verification techniques on business processes.

1. Introduction

The use of IT systems for supporting business activities has brought to a large diffusion of *process mining* (van der Aalst & et al., 2012) techniques and tools that offer business analysts the possibility to observe the current process execution, identify deviations from the model, perform individual and aggregated analysis on current and past executions. According to the process mining manifesto (van der Aalst & et al., 2012), all these techniques and

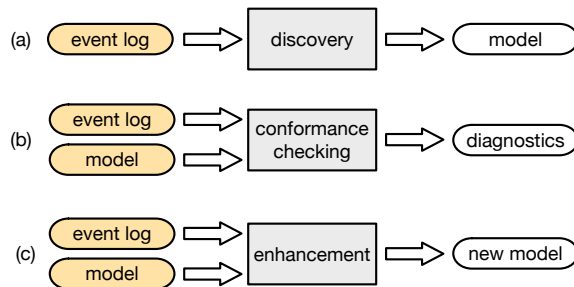


Figure 1: The three types of process mining.

tools can be grouped in three basic types: process discovery, conformance checking and process enhancement (see Figure 1), and require in input an *event log* and/or a (*process*) *model*. A log, usually described in the IEEE standard XES format (IEEE Task Force on Process Mining, 2015), is a set of execution traces (or cases) each of which is an ordered sequence of events carrying a payload as a set of attribute-value pairs. Process models instead provide a description of the scenario at hand and can be constructed using one of the available Business Process Modeling Languages, such as BPMN (Object Management Group (OMG), 2011), YAWL (van der Aalst & ter Hofstede, 2005) and DECLARE (van der Aalst, Pesic, & Schonenberg, 2009).

When it comes to conformance checking, that is the problem of assessing whether an execution trace conforms to the model, and to business process verification more in general, a well known limitation of state of the art works is the lack of concrete support provided for the handling of data-aware business processes, that is business processes for which both the control flow and the data flow are taken into account (De Masellis, Francescomarino, Ghidini, Montali, & Tessaris, 2017a). This limitation concerns both the data flow component described in the process model and the payload data contained in an event log. A second limitation of existing techniques is that they mainly rely on complete observations of terminated process executions. While the presence of non-monitorable activities (or errors in the logging procedure) makes the ability of handling incomplete event logs one of the main challenges of the business process community, as mentioned in the process mining manifesto (van der Aalst & et al., 2012), still most process mining techniques require the presence of event logs which contain complete end-to-end execution traces. Repairing incomplete execution traces by reconstructing the missing entries has therefore become an important task to enable process mining in full, as noted in recent works such as (Rogge-Solti, Ronny, van der Aalst, & Weske, 2013), (Di Francescomarino, Ghidini, Tessaris, & Sandoval, 2015). While these works deserve a praise for having motivated the importance of trace completion and having provided some basic techniques for reconstructing missing entries using the knowledge captured in process models, they again focus on the control flow only, thus completely ignoring the data flow component, and the payload data contained in an event log.

This paper tackles the issue of providing theoretically sound and practically supported data-aware business process verification techniques. It does it by focusing on the problem of automatically completing partial data-aware process executions in the presence of an imperative data-aware process model. More in detail we: (i) formulate the trace repair as a

reachability problem over data-aware models, whose solutions are completions of the partial trace that are compliant with the model and (ii) provide theoretically rigorous mappings between our data-aware models and three important paradigms for reasoning about dynamic systems: Action Languages, Classical Planning, and Model-Checking. Such mappings are based on a common interpretation of the semantics of the above paradigms in terms of transition systems. Notably, (i) and (ii) allows us to exploit the robust tools available for the paradigms above to solve the trace repair problem. Lastly (iii) we extensively evaluate the performances of such tools on the task of trace repair.

Important characteristics of our work are:

1. it is based on the well established Workflow Nets (van der Aalst, Hee, Hofstede, Sidorova, Verbeek, Voorhoeve, & Wynn, 2010), a particular class of Petri nets that have become one of the standard ways to model and analyse workflows;
2. it enriches them with rigorous data-aware constructs that reflect the expressiveness of well-established business process suites and, at the same time, are formally grounded;
3. the rigorous mapping between the data-aware workflow net language and the different reasoning paradigms is based on the solid notion of transition system. Also, by considering empty traces and complete traces, the techniques we present here can be seamlessly exploited to compute model satisfiability and trace compliance, respectively.
4. the reasoning formalisms considered in this work exhibit different characteristics, whose impact in terms of reasoning with data-aware workflow net languages is not easily predictable without a solid evaluation.

These characteristics contribute to advancing the state-of-the-art on formal verification of business processes in different ways: first, the adoption of a data-aware language formally rigorous and close to real modelling platforms is in line with recent efforts to bridge the gap between theoretical investigations and practical support for data-aware process verification (see (De Masellis et al., 2017a)). Second, the (reusable) recast of the trace repair problem in terms of reachability on a transition system is a step forward to the provision of general solutions which exploit the power of logic languages without resorting to ad hoc algorithms. Third, the encoding of the reachability problem in different verification oriented formalisms, and the comparative evaluation of their characteristics and performances, is an important step forward towards the exploitation of mature verification techniques for reasoning with data-aware processes.

The paper is structured as follows. Section 2 introduces the problem of trace completion and illustrates the importance of considering its data-aware component; Section 3 provides some background on Workflow Nets (van der Aalst et al., 2010) and describes the language of DAW-nets originally introduced in (De Masellis, Francescomarino, Ghidini, & Tessaris, 2017b); Section 4 provides a recast of the trace completion problem in terms of reachability while Section 5 describes the approach and the main steps for exploiting the notion of transition system as a bridge between DAW-net and the chosen Action Languages, Classical Planning, and Model-Checking languages. Details of the specific encodings of the DAW-nets reachability problem in the three specific formalisms are provided in Sections 6, 7, and 8;. Finally, Section 9 presents an extensive evaluation by using three well known solvers (i.e., CLINGO, FAST-DOWNWARD and NUXMV) available for the different formalisms both

on synthetic and real life logs. While the evaluation does not indicate a clear “winner”, it empirically shows that adding the data dimension dramatically affects the solvers’ ability to return a result and performance. Moreover, characteristics of the trace, such as its dimension, also have an important impact on which solver performs best. The code used and all datasets are publicly available (see Section 9 for details). The paper ends with related works and concluding remarks.

A preliminary version of this work has been published in (De Masellis et al., 2017b). Compared to the early version, where the focus was on an ad hoc encoding of the problem using a specific *Action Language*, this paper introduces a more general encoding technique based on finite transition systems and an extensive evaluation of different solvers. The new technique highlights the core dynamic properties of Workflow Nets, and enables a seamless encoding of the decision problems under investigation in a variety of formal frameworks.

2. Completing partial traces

One of the goals of process mining is to capture the as-is processes as accurately as possible: this is done by examining event logs that can be then exploited to perform process discovery, conformance checking, and process enhancement, as graphically depicted in Figure 1. In many cases, however, event logs are subject to data quality problems, resulting in *incorrect* or *missing* events in the log. In this paper we focus on the latter issue addressing the problem of **repairing execution traces that contain missing entries** (hereafter shortened in trace completion or trace repair).

The need for trace completion is motivated in depth in (Rogge-Solti et al., 2013), where missing entities are described as a frequent cause of low data quality in event logs¹. The starting point of approaches to trace completion are (partial) *execution traces* and the knowledge captured in *process models*.

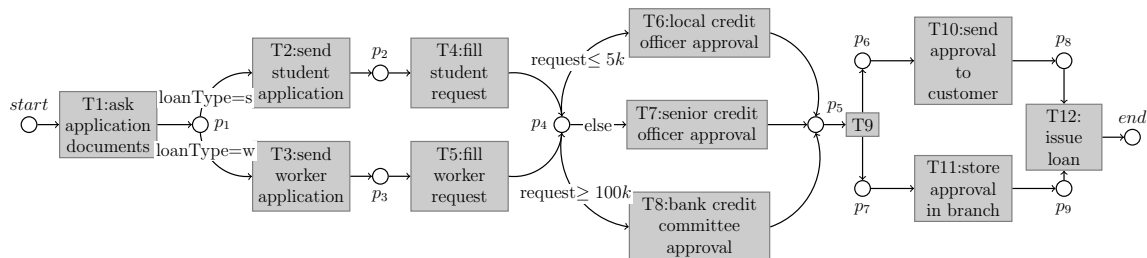


Figure 2: The LoanRequest Process Model.

1. The partiality of monitored information may come, in real settings, from different sources. First, the definitions of business processes often integrate activities that are not supported by IT systems due either to their nature (e.g. they consist of human interactions) or to the high level of abstraction of the description which may be detached from the implementation (e.g., human-oriented description versus information system oriented realisation). Indeed, this dual nature of business processes is so frequent, that widely adopted business process modelling formalisms, such as BPMN or YAWL, allow for a clear distinction between e.g. IT-supported and human-based activities. Second, logging can be subject to faults and errors. Third, IT systems may expose only partial information about their workings and the data they produce, hence preventing a complete execution tracing.

To illustrate the idea of trace completion let us consider a sample process model such as the one depicted in Figure 2 describing a *Loan Request*². The specific notation of Petri Nets used to describe this model will be introduced in the next section, but it is not essential to understand the details of trace repair presented here. The model is composed of 12 tasks $T1-T12$: The application starts with task $T1$ where the customer asks for the application forms. Then, depending on whether the customer is asking for a student loan ($loanType = s$) or a worker loan ($loanType = w$) the process continues with a mutually exclusive choice between the sequence of activities $T2-T4$ and $T3-T5$. Once the appropriate (student or worker) request is filled, another mutually exclusive choice is presented for the execution: if the amount requested is lower than $5k$, then task $T6$ is executed and only a local approval is needed; if the amount requested is greater than $100k$, then the request needs to be approved by the bank committee (task $T8$); in all the other cases the process continues with task $T7$. Once the request is approved by means of the appropriate approval task, two parallel activities are executed: the approval is sent to the customer (task $T10$), and the approval documents are stored in the branch (task $T11$). Finally the loan is issued (task $T12$).

We assume that each execution of this process, hereafter called the Loan Request (LR) process, is logged by an information system. A sample trace³ that logs the execution of a LR instance is:

$$\{T1, T3, T5, T7, T9, T10, T11, T12\}. \quad (1)$$

Multiple execution traces of the same process forms an event log. In trace (1), all the executed activities have been explicitly logged. Consider, now the (partial) execution trace

$$\{T3, T7\}. \quad (2)$$

which only records the execution of tasks $T3$ and $T7$. By aligning the trace to the model using the techniques presented in (Rogge-Solti et al., 2013) and (Di Francescomarino et al., 2015), we can exploit the events stored in the trace and the ordering of activities (hereafter control flow) specified in the model to reconstruct two possible completions: one is the trace in (1); the other is: $\{T1, T3, T5, T7, T9, T11, T10, T12\}$ which only differs from (1) for the ordering of the parallel activities.

Consider now a different scenario in which the partial trace reduces to $\{T7\}$. In this case, by using the control flow in Figure 2 we are not able to reconstruct whether the loan is a student loan or a worker loan. This increases the number of possible completions and therefore lowers the usefulness of trace repair. Assume nonetheless that the event log conforms to the XES standard and stores some observed data attached to $T7$ (enclosed in square brackets):

$$\{T7[request = 60k, loan = 50k]\} \quad (3)$$

If the process model is able to specify how transitions can read and write variables, and furthermore some constraints on how they do it, the scenario changes completely. Indeed,

-
2. For the sake of simplicity we only focus here on the, so-called, happy path, that is the successful granting of the loan.
 3. In this work we are interested in (reconstructing) the ordered sequence of events that constitutes a trace, and thus we ignore timestamps as attributes. In the illustrative examples we present the events in a trace ordered according to their execution time.

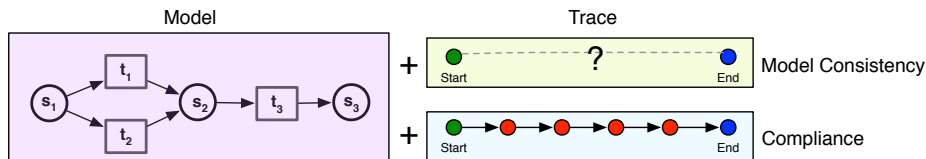


Figure 3: Corner-case incomplete execution traces.

assume that transition T_4 is empowered with the ability to write the variable *request* with a value smaller or equal than 30k (being this the maximum amount of a student loan). Using this fact, and the fact that the request examined by T_7 is greater than 30k, we can understand that the execution trace has chosen the path of the worker loan. Moreover, if the model specifies that variable *loanType* is written during the execution of T_1 , when the applicant chooses the type of loan she is interested to, we are able to infer that T_1 sets variable *loanType* to *w*.

This example, besides illustrating the idea of trace completion, also motivates why data (both in the process model and in the execution trace) are important to accomplish this task, and therefore why extending reasoning techniques beyond the mere control flow is a significant contribution to address data quality problems in event logs. Also, as already observed in (Chesani, Mello, De Masellis, Di Francescomarino, Ghidini, Montali, & Tessaris, 2018), by considering corner cases of partial traces, such as empty traces and complete traces, the ability to produce a complete(d) trace compliant with the process model can be also exploited to compute model consistency and trace compliance, respectively (see Figure 3).

3. The Framework: DAW-net

In this section we first provide the necessary background on Petri Nets and WF-nets (Section 3.1) and then suitably extend WF-nets to represent data and their evolution as transitions are performed (Section 3.2)

3.1 The Workflow Nets modeling language

Petri Nets (Petri, 1962) (PN) is a modeling language for the description of distributed systems that has widely been applied to the description and analysis of business processes (van der Aalst, 1998). The classical PN is a directed bipartite graph with two node types, called *places* and *transitions*, connected via directed arcs. Connections between two nodes of the same type are not allowed.

Definition 1 (Petri Net). *A Petri Net is a triple $\langle P, T, F \rangle$ where P is a set of places; T is a set of transitions; $F \subseteq (P \times T) \cup (T \times P)$ is the flow relation describing the arcs between places and transitions (and between transitions and places).*

The *preset* of a transition t is the set of its input places: $\bullet t = \{p \in P \mid (p, t) \in F\}$. The *postset* of t is the set of its output places: $t \bullet = \{p \in P \mid (t, p) \in F\}$. Definitions of pre- and postsets of places are analogous. Places in a PN may contain a discrete number of marks

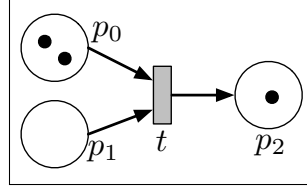


Figure 4: A Petri Net.

called tokens. Any distribution of tokens over the places, formally represented by a total mapping $M : P \mapsto \mathbb{N}$, represents a configuration of the net called a *marking*.

PNs come with a graphical notation where places are represented by means of circles, transitions by means of rectangles and tokens by means of full dots within places. Figure 4 depicts a PN with a marking $M(p_0) = 2$, $M(p_1) = 0$, $M(p_2) = 1$. The preset and postset of t are $\{p_0, p_1\}$ and $\{p_2\}$, respectively.

Process tasks are modeled in PNs as transitions while arcs and places constraint their ordering. For instance, the process in Figure 2 exemplifies how PNs can be used to model parallel and mutually exclusive choices, typical of business processes: sequences $T2;T4$ and $T3;T5$ are placed on mutually exclusive paths; the same holds for transitions $T6$ $T7$, and $T8$. Transitions $T10$ and $T11$ are instead placed on parallel paths. Finally, $T9$ is needed to prevent connections between nodes of the same type.

The expressivity of PNs exceeds, in the general case, what is needed to model business processes, which typically have a well-defined starting point and a well-defined ending point. This imposes syntactic restrictions on PNs, that result in the following definition of a workflow net (WF-net) (van der Aalst, 1998).

Definition 2 (WF-net). *A PN $\langle P, T, F \rangle$ is a WF-net if it has a single source place $start$, a single sink place end , and every place and every transition is on a path from $start$ to end , i.e., for all $n \in P \cup T$, $(start, n) \in F^*$ and $(n, end) \in F^*$, where F^* is the reflexive transitive closure of F .*

A marking in a WF-net represents the *workflow state* of a process execution. We distinguish two special markings: the *initial marking* M_s which has one token in the *start* place ($M_s(start) = 1$) and all others places are empty ($\forall p \in P \setminus \{start\}.M_s(p) = 0$) and the *final marking* M_e which has one token in the *end* place ($M_e(end) = 1$) and all others places are empty ($\forall p \in P \setminus \{end\}.M_e(p) = 0$).

The semantics of a PN/WF-net, and in particular the notion of *valid firing*, defines how transitions route tokens through the net so that they correspond to a process execution.

Definition 3 (Valid Firing). *A firing of a transition $t \in T$ from M to M' is valid, in symbols $M \xrightarrow{t} M'$, iff*

1. t is enabled in M , i.e., $\{p \in P \mid M(p) > 0\} \supseteq \bullet t$; and
2. the marking M' is such that for every $p \in P$:

$$M'(p) = \begin{cases} M(p) - 1 & \text{if } p \in \bullet t \setminus t^\bullet \\ M(p) + 1 & \text{if } p \in t^\bullet \setminus \bullet t \\ M(p) & \text{otherwise} \end{cases}$$

Condition 1. states that a transition is enabled if all its input places contain at least one token; condition 2. states that when t fires it consumes one token from each of its input places and produces one token in each of its output places.

Notationally, $t_1 t_2 t_3 \dots t_k$ is a sequence of valid firings that leads from M_0 to M_k iff $M_0 \xrightarrow{t_1} M_1 \xrightarrow{t_2} M_2 \xrightarrow{t_3} \dots \xrightarrow{t_{k-1}} M_{k-1} \xrightarrow{t_k} M_k$. In this paper we use the term *case* of a WF-Net to denote a sequence of valid firings $M_0 \xrightarrow{t_1} M_1 \xrightarrow{t_2} M_2 \xrightarrow{t_3} \dots \xrightarrow{t_{k-1}} M_{k-1} \xrightarrow{t_k} M_k$ such that $M_0 = M_s$ is the initial marking and $M_k = M_e$ is the final marking. A case is thus a sequence of valid firings that connect the initial to the final marking through the PN.

From now on we concentrate on *safe* nets, which generalise the class of *structured workflows* and are the basis for best practices in process modeling (Kiepuszewski, ter Hofstede, & Bussler, 2013). The notion of safeness is defined in terms of k -boundedness (see also (van der Aalst et al., 2010)).

Definition 4 (k -boundedness and safeness). *A marking of a PN is k -bound if the number of tokens in all places is at most k . A PN is k -bound if the initial marking is k -bound and the marking of all cases is k -bound. If $k = 1$, the PN is safe.*

It is important to notice that our approach can be seamlessly generalized to other classes of PNs, as long as they are k -bound.

Reachability on Petri Nets. The behavior of a PN can be described as a transition system where states are markings and directed edges represent firings. Intuitively, there is an edge from marking M_i to marking M_{i+1} , labeled by t_i , if $M_i \xrightarrow{t_i} M_{i+1}$ is a valid firing. Given a “goal” marking M_g , the reachability problem amounts to checking whether there is a path from the initial marking M_s to M_g . Reachability on PNs (WF-nets) is of enormous importance in process verification as it allows for checking natural behavioral properties, such as satisfiability and soundness in a natural manner (van der Aalst, 1997).

3.2 The DAW-net modeling language

DAW-net (De Masellis et al., 2017b) is a data-aware extension of WF-nets which enriches them with the capability of reasoning on data. DAW-net extends WF-nets by providing: (i) a model for representing data; (ii) a way to make decisions on actual data values; and (iii) a mechanism to express modifications to data. It does it by introducing:

- a set of variables taking values from possibly different domains (addressing (i));
- queries on such variables used as transitions preconditions (addressing (ii))
- variables updates and deletion in the specification of net transitions (addressing (iii)).

DAW-net follows the approach of state-of-the-art WF-nets with data (Sidorova, Stahl, & Trčka, 2011; de Leoni & van der Aalst, 2013), from which it borrows the above concepts, extending them by allowing reasoning on actual data values as better explained in Section 10.

Throughout the section we use the WF-net in Figure 2 extended with data as a running example.

3.2.1 DATA MODEL

The data model of WF-nets take inspiration from the data model of the IEEE XES standard for describing event logs, which represents data as a set of variables. Variables take values from specific sets on which a partial order can be defined. As customary, we distinguish between the data model, namely the intensional level, from a specific instance of data, i.e., the extensional level.

Definition 5 (Data model). *A data model is a tuple $\mathcal{D} = (\mathcal{V}, \Delta, dm, ord)$ where:*

- \mathcal{V} is a possibly infinite set of variables;
- $\Delta = \{\Delta_1, \Delta_2, \dots\}$ is a possibly infinite set of domains (not necessarily disjoint);
- $dm : \mathcal{V} \rightarrow \Delta$ is a total and surjective function which associates to each variable v its domain Δ_i ;
- ord is a partial function that, given a domain Δ_i , if $ord(\Delta_i)$ is defined, then it returns a partial order (reflexive, antisymmetric and transitive) $\leq_{\Delta_i} \subseteq \Delta_i \times \Delta_i$.⁴

A data model for the loan example is $\mathcal{V} = \{loanType, request, loan\}$, $dm(loanType) = \{\mathbf{w}, \mathbf{s}\}$, $dm(request) = \mathbb{N}$, $dm(loan) = \mathbb{N}$, with $dm(loan)$ and $dm(loanType)$ being total ordered by the natural ordering \leq in \mathbb{N} .

An actual instance of a data model is simply a partial function associating values to variables.

Definition 6 (Assignment). *Let $\mathcal{D} = \langle \mathcal{V}, \Delta, dm, ord \rangle$ be a data model. An assignment for variables in \mathcal{V} is a partial function $\eta : \mathcal{V} \rightarrow \bigcup_i \Delta_i$ such that for each $v \in \mathcal{V}$, if $\eta(v)$ is defined, i.e., $v \in img(\eta)$ where img is the image of η , then we have $\eta(v) \in dm(v)$.*

Instances of the data model are accessed through a boolean query language, which notably allows for equality and comparison. As will become clearer in Section 3.2.2, queries are used as *guards*, i.e., preconditions for the execution of transitions.

Definition 7 (Query language - syntax). *Given a data model, the language $\mathcal{L}(\mathcal{D})$ is the set of formulas Φ inductively defined according to the following grammar:*

$$\Phi \quad := \quad true \mid def(v) \mid v = t_2 \mid t_1 \leq t_2 \mid \neg\Phi_1 \mid \Phi_1 \wedge \Phi_2$$

where $v \in \mathcal{V}$ and $t_1, t_2 \in \mathcal{V} \cup \bigcup_i \Delta_i$.

Examples of queries of the loan scenarios are $request \leq 5k$ or $loanType = \mathbf{w}$. Given a formula Φ or a term t , and an assignment η , we write $\Phi[\eta]$ (respectively $t[\eta]$) for the formula Φ' (or the term t') where each occurrence of variable v for which η is defined is replaced by $\eta(v)$ (i.e. unassigned variables are not substituted).

Definition 8 (Query language - semantics). *Given a data model \mathcal{D} , an assignment η and a query $\Phi \in \mathcal{L}(\mathcal{D})$ we say that \mathcal{D}, η satisfies Φ , written $\mathcal{D}, \eta \models \Phi$ inductively on the structure of Φ as follows:*

- $\mathcal{D}, \eta \models true$;
- $\mathcal{D}, \eta \models def(v)$ iff $v[\eta] \neq v$;

4. To simplify the model we assume that if $\{(o, o'), (o', o)\} \subseteq ord(\Delta_i)$ then $o = o'$.

- $\mathcal{D}, \eta \models t_1 = t_2$ iff $t_1[\eta], t_2[\eta] \notin \mathcal{V}$ and $t_1[\eta] \equiv t_2[\eta]$;
- $\mathcal{D}, \eta \models t_1 \leq t_2$ iff $t_1[\eta], t_2[\eta] \in \Delta_i$ for some i and $\text{ord}(\Delta_i)$ is defined and $t_1[\eta] \leq_{\Delta_i} t_2[\eta]$;
- $\mathcal{D}, \eta \models \neg\Phi$ iff it is not the case that $\mathcal{D}, \eta \models \Phi$;
- $\mathcal{D}, \eta \models \Phi_1 \wedge \Phi_2$ iff $\mathcal{D}, \eta \models \Phi_1$ and $\mathcal{D}, \eta \models \Phi_2$.

Intuitively, `def` can be used to check if a variable has an associated value or not (recall that assignment η is a partial function); equality has the intended meaning and $t_1 \leq t_2$ evaluates to true iff t_1 and t_2 are values belonging to the same domain Δ_i , such a domain is ordered by a partial order \leq_{Δ_i} and t_1 is actually less or equal than t_2 according to \leq_{Δ_i} .

Lemma 1. *Let $\Phi \in \mathcal{L}(\mathcal{D})$, and $\text{adm}(\Phi)$ be its active domain – i.e. the set of constants appearing in Φ . Given an assignment η we consider the data model \mathcal{D}' as the restriction of \mathcal{D} w.r.t. $\text{adm}(\Phi) \cup \text{img}(\eta)$, then*

$$\mathcal{D}, \eta \models \Phi \text{ iff } \mathcal{D}', \eta \models \Phi$$

The last lemma can be easily proved by structural induction on the formula, and it makes the queries independent on the actual “abstract” domains enabling the use of finite subsets. Intuitively, the property holds because the query language doesn’t include quantification over the variables.

3.2.2 DATA-AWARE NET

Data-Aware nets (DAW-net) are obtained by combining the data model just introduced with a WF-net, and by formally defining how transitions are guarded by queries and how they update/delete data.

Definition 9 (DAW-net). *A DAW-net is a tuple $\langle \mathcal{D}, \mathcal{N}, \text{wr}, \text{gd} \rangle$ where:*

- $\mathcal{N} = \langle P, T, F \rangle$ is a WF-net;
- $\mathcal{D} = \langle \mathcal{V}, \Delta, \text{dm}, \text{ord} \rangle$ is a data model;
- $\text{wr} : T \mapsto \bigcup_{\mathcal{V}' \subseteq \mathcal{V}} (\mathcal{V}' \mapsto 2^{\text{dm}(\mathcal{V})})$, where $\text{dm}(\mathcal{V}) = \bigcup_{v \in \mathcal{V}} \text{dm}(v)$, is a function that associates each transition to a (partial) function mapping variables to a finite subset of their domain; that is satisfying the property that for each t , $\text{wr}(t)(v) \subseteq \text{dm}(v)$ for each v in the domain of $\text{wr}(t)$.
- $\text{gd} : T \mapsto \mathcal{L}(\mathcal{D})$ is a function that associates a guard to each transition.

Function `gd` associates a guard, namely a query, to each transition. The intuitive semantics is that a transition t can fire if its guard $\text{gd}(t)$ evaluates to true (given the current assignment of values to data). Examples are $\text{gd}(T6) = \text{request} \leq 5\text{k}$ and $\text{gd}(T8) = \neg(\text{request} \leq 99999)$.

Function `wr` is used to express how a transition t modifies data: after the firing of t , each variable v in the domain of $\text{wr}(t)$ can take any value among a specific finite subset of $\text{dm}(v)$. We have three different cases:

- $\emptyset \subset \text{wr}(t)(v) \subseteq \text{dm}(v)$: t nondeterministically assigns a value from $\text{wr}(t)(v)$ to v ;
- $\text{wr}(t)(v) = \emptyset$: t deletes the value of v (hence making v undefined);
- $v \notin \text{dom}(\text{wr}(t))$: value of v is not modified by t .

Notice that by allowing $\text{wr}(t)(v) \subseteq \text{dm}(v)$ in the first bullet above we enable the specification of restrictions for specific tasks. E.g., $\text{wr}(T4) : \{\text{request}\} \mapsto \{0 \dots 30\text{k}\}$ says that $T4$ writes the *request* variable and intuitively that students can request a maximum loan of 30k, while $\text{wr}(T5) : \{\text{request}\} \mapsto \{0 \dots 500\text{k}\}$ says that workers can request up to 500k.

The intuitive semantics of *gd* and *wr* is formalized in the notion of DAW-net valid viring, which is based on the definition of DAW-net state, and state transition. A DAW-net state includes both the state of the WF-net, namely its marking, and the state of data, namely the assignment.

Definition 10 (DAW-net state). *A state of a DAW-net $\langle \mathcal{D}, \mathcal{N}, \text{wr}, \text{gd} \rangle$ is a pair (M, η) where M is a marking for $\langle P, T, F \rangle$ and η is an assignment for \mathcal{D} .*

Analogously to traditional WF-nets, we distinguish two special states: the *initial state* (M_s, η_s) which is such that M_s is the initial marking of \mathcal{N} and η_s is empty, i.e., $\text{dom}(\eta_s) = \emptyset$, and the *final states* (M_e, η_e) are such that M_e is the final marking of \mathcal{N} .

Definition 11 (DAW-net Valid Firing). *Given a DAW-net $\langle \mathcal{D}, \mathcal{N}, \text{wr}, \text{gd} \rangle$, a firing of a transition $t \in T$ is a valid firing from (M, η) to (M', η') , written as $(M, \eta) \xrightarrow{t} (M', \eta')$, iff*

1. $M \xrightarrow{t} M'$ is a WF-Net valid firing
2. $\mathcal{D}, \eta \models \text{gd}(t)$,
3. assignment η' is such that, if $\text{WR} = \{v \mid \text{wr}(t)(v) \neq \emptyset\}$, $\text{DEL} = \{v \mid \text{wr}(t)(v) = \emptyset\}$:
 - its domain $\text{dom}(\eta') = \text{dom}(\eta) \cup \text{WR} \setminus \text{DEL}$;
 - for each $v \in \text{dom}(\eta')$:

$$\eta'(v) = \begin{cases} d \in \text{wr}(t)(v) & \text{if } v \in \text{WR} \\ \eta(v) & \text{otherwise.} \end{cases}$$

Conditions 2. and 3. extend the notion of valid firing of WF-nets imposing additional pre- and postconditions on data, and in particular preconditions on η and postconditions on η' . Specifically, 2. says that for a transition t to be fired its guard $\text{gd}(t)$ must be satisfied by the current assignment η . Condition 3. constrains the new state of data: the domain of η' is defined as the union of the domain of η with variables that are written (WR), minus the set of variables that must be deleted (DEL). Variables in $\text{dom}(\eta')$ can indeed be grouped in three sets depending on the effects of t : (i) $\text{OLD} = \text{dom}(\eta) \setminus \text{WR}$: variables whose value is unchanged after t ; (ii) $\text{NEW} = \text{WR} \setminus \text{dom}(\eta)$: variables that were undefined but have a value after t ; and (iii) $\text{OVERWR} = \text{WR} \cap \text{dom}(\eta)$: variables that did have a value and are updated with a new one after t . The final part of condition 3. says that each variable in $\text{NEW} \cup \text{OVERWR}$ takes a value in $\text{wr}(t)(v)$, while variables in OLD maintain the old value $\eta(v)$.

Similarly to regular PNs, a *case* of a DAW-net is a sequence of DAW-net valid firings $(M_0, \eta_0) \xrightarrow{t_0} (M_1, \eta_1) \xrightarrow{t_1} \dots \xrightarrow{t_{k-1}} (M_{k-1}, \eta_{k-1}) \xrightarrow{t_k} (M_k, \eta_k)$ such that (M_0, η_0) is the initial state and (M_k, η_k) is a final state.

Note that, since we assumed a finite set of transitions and $\text{wr}(t)$ is finite for any transition t , for any model W we can consider its active domain $\text{adm}(W)$ as the set of all constants in $\text{wr}(t)$ and $\text{gd}(t)$ for all the transitions t in W . Moreover, given the Lemma 1, we can consider the restriction of the model to the finite set of states defined by the active domain.

Definition 12 (Finite DAW-net models). *Let $W = \langle \mathcal{D}, \mathcal{N}, wr, gd \rangle$ be a DAW-net with $\mathcal{D} = \langle \mathcal{V}, \Delta, dm, ord \rangle$ and $\mathcal{N} = \langle P, T, F \rangle$, $adm(W)$ be its active domain. We consider the data model \mathcal{D}_W as the restriction of \mathcal{D} w.r.t. $adm(W)$.*

The finite version of W , denoted as \overline{W} , is defined as $\langle \mathcal{D}_W, \mathcal{N}, wr, gd \rangle$.

By looking at the definition of valid firing and using Lemma 1 it's easy to realise that the set of states of \overline{W} is closed w.r.t. the relation induced by the valid firing. As we'll see later on this enables us to focus on finite DAW-net models.

Reachability on DAW-nets. Similarly to regular PNs, the behaviour of a DAW-nets can be described in terms of a transition system and reachability. The next two sections are indeed devoted to reformulate the trace completion problem as a reachability problem (Section 4) and then to introduce the notion of transition system for a DAW-nets (Section 5). We conclude this section by remarking that although this work focuses on solving the trace completion problem, the technique we present for checking reachability is more general and can indeed be used to check any property on DAW-nets that can be reduced to reachability. E.g., the *soundness* property of WF-nets as defined in (van der Aalst, 1997) can be naturally lifted to DAW-nets and checked with the procedure shown here, by simply considering an empty trace and by formulating the soundness requirements of *proper termination* and *no dead transitions* (van der Aalst, 1997) as reachability properties.

4. Trace repair as reachability

In this section we provide the intuition behind our technique for solving the trace repair problem via reachability. Full details and proofs are contained in Appendix A.

A *trace* is a sequence of observed *events*, each one associated to the transition it refers to. Events can also be associated to data payloads in the form of attribute/variable-value pairs. Among these variables we can distinguish those updated by its execution (see Definition 9) and those observed. Executions (1), (2), and (3) at page 5 provide examples of traces without and with data payloads for our running example.⁵

Definition 13 (Trace). *Let $W = \langle \mathcal{D}, \mathcal{N}, wr, gd \rangle$ be a DAW-net. An event of W is a tuple $\langle t, w, w^d \rangle$ where $t \in T$ is a transition, $w \in \bigcup_{\mathcal{V}' \subseteq \mathcal{V}} (\mathcal{V}' \mapsto dm(\mathcal{V}))$ is a partial function that represents the variables written or observed by the execution of t , and $w^d \subseteq \mathcal{V}$ the set of variables deleted (undefined) or observed to be undefined by the execution of t . Obviously, w^d and the domain of w do no share any variable.*

A trace of W is a finite sequence of events $\tau = (e_1, \dots, e_n)$. In the following we indicate the i -th event of τ as τ^i . Given a set of transitions T , the set of traces is inductively defined as follows:

- ϵ is a trace;
- if τ is a trace and e an event, then $\tau \cdot e$ is a trace.

5. We remark the conceptual and practical difference between a case and a trace: a case is a semantic object intuitively representing an execution/behavior of the business process model while a trace is syntactic object that comes from the real world of which we would test the soundness by providing a possible repair.

Notice that the definition above constraints the transitions neither to be complete nor correctly ordered w.r.t. W . What we would like to check is indeed if τ can be completed so as to represent a valid case of W . Intuitively, a trace τ is *compliant* with a case C of a DAW-net W if C contains all the occurrences of the transitions observed in τ (with the corresponding variable updates) in the right order, as the following definition formalises.

Definition 14 (Trace Compliance). *An event $\langle t', w, w^d \rangle$ is compliant with a (valid) firing $(M, \eta) \xrightarrow{t'} (M', \eta')$ iff $t = t'$, $\text{dom}(w) \subseteq \text{dom}(\eta') \subseteq \mathcal{V} \setminus w^d$, and for all $v \in \text{dom}(w)$ $w(v) = \eta'(v)$.*

A trace $\tau = (e_1, \dots, e_\ell)$, $\ell > 0$, is compliant with a sequence of valid firings

$$(M_0, \eta_0) \xrightarrow{t_1} (M_1, \eta_1) \dots (M_{k-1}, \eta_{k-1}) \xrightarrow{t_k} (M_k, \eta_k)$$

iff there is an injective mapping γ between $[1 \dots \ell]$ and $[1 \dots k]$ such that:⁶

$$\forall i, j \text{ s.t. } 1 \leq i < j \leq \ell \quad \gamma(i) < \gamma(j) \quad (4)$$

$$\forall i \text{ s.t. } 1 \leq i \leq \ell \quad e_i \text{ is compliant with } (M_{\gamma(i-1)}, \eta_{\gamma(i-1)}) \xrightarrow{t_{\gamma(i)}} (M_{\gamma(i)}, \eta_{\gamma(i)}) \quad (5)$$

When $\ell = 0$ we impose that τ is compliant with any sequence of valid firings.

We say that a trace τ is compliant with a DAW-net W if there exists a case C of W such that τ is compliant with C .

Although the previous definition is general we notice that, since the last state of a case C must be a final state, τ being compliant with W means that it can be completed so as to represent a *successful* execution of the process W .

To simplify the presentation of the technical results, we assume without loss of generality that DAW-net models start with the special transition $start_t$ and terminate with the special transition end_t . Both have a single input and output places: $start_t$ input is the start place, and end_t output place is the sink. These transitions have true guards (always satisfiable) and don't modify variables; moreover, start and sink places shouldn't have any incoming and outgoing edge respectively. Every process can be reduced to such a structure as informally illustrated in the top part of of Figure 5 by the arrows labeled with (1). Note that this change would not modify the behaviour of the net: any sequence of firing valid for the original net can be extended by the firing of the additional transitions and vice versa.

We now illustrate the main idea behind our approach by means of the bottom part of Figure 5: we consider the observed events as additional transitions (in red) and we suitably “inject” them in the original DAW-net. By doing so, we obtain a new model where we aim at forcing tokens to activate the red transitions when events are observed in the trace. When, instead, there is no red counterpart, i.e., there is missing information in the trace, the tokens are free to move in the black part of the model.

More precisely, for each event e corresponding to the execution of a transition t with some data payload, we introduce a new transition t_e in the model such that:

- t_e is placed in exclusive or with the original transition t ;

6. If the trace is empty then $\ell = 0$ and γ is empty.

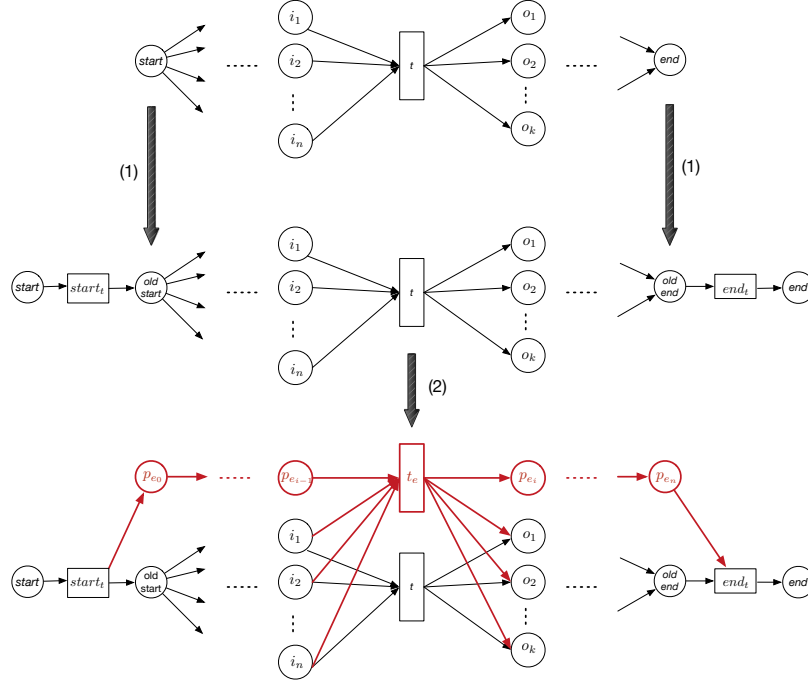


Figure 5: Outline of the trace “injection”.

- t_e includes an additional input place connected to the preceding event and an additional output place which connects it to the next event;
- $\text{gd}(t_e)$ is a conjunction of the original $\text{gd}(t)$ with the conditions that state that the variables in the event data payload that are not modified by the transition are equal to the observed ones;
- $\text{wr}(t_e)$ updates the values corresponding the variables updated by the event payload, i.e. if the event sets the value of v to d , then $\text{wr}(t_e)(v) = \{d\}$; if the event deletes the variable v , then $\text{wr}(t_e)(v) = \emptyset$.

Note that an event $\langle t, w, w^d \rangle$ cannot be compliant with any firing unless $w(v) \in \text{wr}(t)(v)$ for each variable in $\text{dom}(w) \cap \text{dom}(\text{wr}(t))$, and $w^d \cap \text{dom}(\text{wr}(t)) \subseteq \{v \mid (v, \emptyset) \in \text{wr}(t)\}$. Since these properties can be verified just by comparing the events to the model, then in the following we assume that they are both satisfied for all the events in a trace.

Given a DAW-net W , and a trace τ , the formal definition of the “injection” of τ into W , denoted with W^τ , is given below.

Definition 15 (Trace workflow). *Let $W = \langle \mathcal{D}, \mathcal{N} = \langle P, T, F \rangle, \text{wr}, \text{gd} \rangle$ be a DAW-net and $\tau = (e_1, \dots, e_n)$ – where $e_i = \langle t_i, w_i, w_i^d \rangle$ – a trace of W . The trace workflow $W^\tau = \langle \mathcal{D}, \mathcal{N}^\tau = \langle P^\tau, T^\tau, F^\tau \rangle, \text{wr}^\tau, \text{gd}^\tau \rangle$ is defined as following:*

- $P^\tau = P \cup \{p_{e_0}\} \cup \{p_e \mid e \in \tau\}$, with p_{e_0}, p_e new places;
- $T^\tau = T \cup \{t_e \mid e \in \tau\}$, with t_e new transitions;

- $F^\tau = F \cup$
 $\{(t_{e_i}, p) \mid i = 1 \dots n, (t_i, p) \in F\} \cup \{(p, t_{e_i}) \mid i = 1 \dots n, (p, t_i) \in F\} \cup$
 $\{(t_{e_i}, p_{e_i}) \mid i = 1 \dots n\} \cup \{(p_{e_{i-1}}, t_{e_i}) \mid i = 1 \dots n\} \cup \{(start_t, p_{e_0}), (p_{e_n}, end_t)\}$
- $wr^\tau(t) = \begin{cases} \{v \mapsto \{w_i(v)\} \mid v \in dom(w_i) \cap dom(wr(t_i))\} \cup \\ \{v \mapsto wr(t_i)(v) \mid v \in dom(wr(t_i)) \setminus dom(w_i)\} & \text{for } t = t_{e_i} \\ wr(t) & \text{for } t \in T \end{cases}$
- $gd^\tau(t) = \begin{cases} gd(t_i) \wedge \bigwedge_{v \in dom(w_i) \setminus dom(wr(t_i))} v = w_i(v) \wedge \\ \bigwedge_{v \in w_i^d \setminus dom(wr(t_i))} \neg def(v) & \text{for } t = t_{e_i} \\ gd(t) & \text{otherwise} \end{cases}$

It is immediate to see that W^τ is a strict extension of W (only new nodes are introduced) and, since all newly introduced nodes are in a path connecting the start and sink places, it is a DAW-net, whenever the original one is a DAW-net.

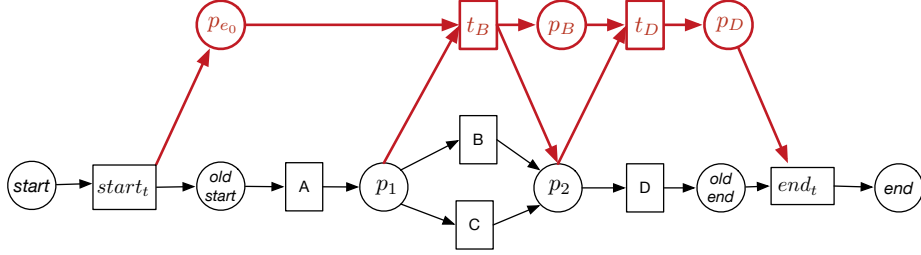


Figure 6: A sample of trace workflow obtained by adding trace $\langle B, D \rangle$.

By looking at this definition, at its graphical illustration at the bottom of Figure 5, and also at the simple example in Figure 6 it is easy to see why the red transitions must be preferred over the black ones in all the sequences of valid firings that go from $start$ to end (i.e., in all cases). Intuitively, transition end_t needs a token in the red place p_{e_n} to fire, in addition to the token in the original end of the initial net. By construction all the red places are only connected to red transitions. Thus, to properly terminate, when choosing between t_e and t , t_e must be preferred so that it produces in output all the tokens for the original black output places plus the one for the new red place. Note that whenever t was able to fire in the original net, t_e can fire as well. Indeed, by construction, it is easy to see that the token that $start_t$ generates in p_{e_0} is only consumed and produced by red transitions into red places.

We now prove the correctness and completeness of the approach by showing that W^τ characterises all and only the cases C of W to which τ is compliant with. We focus here only on the main intuitive steps of the proof, whose technical details are reported in full in Appendix A.

Roughly speaking what we need to show is that:

- (correctness) if C is an arbitrary case of W^τ , then there is a “corresponding” case C' of W such that τ is compliant with C' ; and

- (completeness) if τ is compliant with a case C of W , then there is a case C' in W^τ “corresponding” to C .

A fundamental step for our proof is therefore a way to relate cases from W^τ to the original DAW-net W . To do this we introduce a projection function Π_τ that maps elements from sequences of valid firings (cases) of the enriched DAW-net W^τ to sequences of valid firings (cases) composed of elements from the original DAW-net W . The formal definition of Π_τ is given in Definition 28 of Appendix A. In essence Π_τ maps newly introduced transitions t_e to the original transition t corresponding to the event e , and projects away the newly introduced places in the markings.

Once defined the projection function Π_τ the correctness and completeness statement can be formally formulated as follows:

Theorem 1. *Let W be a DAW-net and $\tau = (e_1, \dots, e_n)$ a trace; then W^τ characterises all and only the cases of W compatible with τ . That is*

- \Rightarrow *if C is a case of W^τ then τ is compliant with the case of W $\Pi_\tau(C)$; and*
- \Leftarrow *if τ is compliant with the case of W C then there is a case C' of W^τ s.t. $\Pi_\tau(C') = C$.*

The proof of correctness follows from two fundamental properties of W^τ and the projection function Π_τ respectively. The first property, ensured by construction of W^τ , states that τ is compliant with all cases of W^τ . The second property, instead, states that any case for W^τ can be replayed on W through the projection Π_τ , by mapping the new transitions t_e into the original ones t , as shown by Lemma 9 in Appendix A. Given these two properties, it is possible to prove that, whenever the case C of W^τ , then τ is compatible with it and with its projection $\Pi_\tau(C)$ on W (See Lemma 12 in Appendix A). The proof of completeness instead is based in the fact that we can build a case for W^τ starting from the case of W with which τ is compliant, by substituting the occurrences of firings corresponding to events in τ with the newly introduced transitions (See Lemma 13 in Appendix A).

Theorem 1 provides the main result of this section and is the basis for the reduction of the trace completion for W and τ to the reachability problem for W^τ : indeed, to check if τ is compliant with W all we have to do is finding a case of W^τ .

Corollary 1 (Trace completion as reachability). *Let W be a DAW-net and τ a trace; then τ is compliant with W iff there is a final state (M_e, η') of W^τ than can be reached from the initial state (M_s, η_s) .*

Proof. This is immediate from Theorem 1 because a final state of W^τ can be reached from its initial state iff there exists at least a case of W^τ . □

We conclude the section by stating that the transformation generating W^τ is preserving the safeness properties of the original workflow (proof in Appendix A).

Theorem 2. *Let W be a DAW-net and τ a trace of W . If W is k -safe then W^τ is k -safe as well.*

5. Reachability: from theory to practice

In this Section we illustrate our approach to complete partial trace executions by exploiting state-of-the-art techniques and tools. The paradigms we selected are: action languages, classical planning, and model-checking. The specific tools we used are CLINGO, FAST-DOWNWARD, and NUXMV, respectively; paired with the representation languages BC, PDDL, and SMV, that will be briefly recasted later in the paper.

In the following, given a DAW-net W , we show how to encode W in the specific language used by each paradigm/tool and formulate the reachability problem as in Corollary 1. We then prove that such encodings are sound and complete, i.e., if a solution is found, then it is indeed a solution for the original problem, and if it is not, then there are no solutions. Notably, we do so by providing the semantics of the above tools in terms of transition systems, which provides an intermediate representation that could support the extension of this work to new paradigms/tools whose semantics can be expressed likewise. This also allows us to use the same structure for the sound and completeness proofs for the encodings, whose main conceptual steps are now described. Notice that when DAW-net is a trace workflow W^τ , then if a solution is found, it is a completion for τ , otherwise τ is not compliant with W .

In order to exploit the notion of transition systems as a bridge between DAW-net and the three chosen paradigms, we first need to connect transition systems and a DAW-net. To do so, we observe that the behaviour of a DAW-net is possibly nondeterministic, as in each state, in general, more than one transition can fire: thus, while cases define a possible DAW-net evolution, a transition system, called *reachability graph*, captures *all* possible evolutions.

Definition 16 (DAW-net reachability graph). *Let $W = \langle \mathcal{D}, \mathcal{N}, wr, gd \rangle$ be a DAW-net with $\mathcal{D} = \langle \mathcal{V}, \Delta, dm, ord \rangle$ and $\mathcal{N} = \langle P, T, F \rangle$, then its reachability graph is a transition system $RG_W = (T, \bar{S}, \bar{s}_0, \bar{\delta})$ where:*

- \bar{S} is the set of states, each representing a DAW-net (M, η) state;
- $\bar{s}_0 = (M_s, \eta_s)$ is the initial state;
- \bar{S} and $\bar{\delta} \subseteq \bar{S} \times T \times \bar{S}$ are defined by induction as follows:
 - $\bar{s}_0 \in \bar{S}$;
 - if $(M, \eta) \in \bar{S}$ and $(M, \eta) \xrightarrow{t} (M', \eta')$ is a valid firing then $s' = (M', \eta') \in \bar{S}$ and $(s, t, s') \in \bar{\delta}$ (and we write $s \xrightarrow{t} s' \in \bar{\delta}$).

Let us consider the finite version of W as defined in Definition 12; clearly the initial state is among its states and, since the set of its states is closed w.r.t. the relation induced by the valid firings we can state the following:

Lemma 2. *Let $W = \langle \mathcal{D}, \mathcal{N}, wr, gd \rangle$ be a DAW-net model and \bar{W} its finite version, then their reachability graphs are identical.*

In virtue of the above lemma, in the rest of the paper when we mention the states of a DAW-net W we intend the states of \bar{W} , that is, we restrict to the finite set of states defined by the active domain.

As a second observation the semantics of our chosen paradigms can be given by means of transition systems. Given a DAW-net W , let us denote with $BC(W)$, $PDDL(W)$ and $SMV(W)$ the encodings of W in the respective languages and with $TS_{BC(W)}$, $TS_{PDDL(W)}$ and $TS_{SMV(W)}$ the transition systems generated by the above encodings. A *path* of a transition system RG_W is a sequence of states obtained by starting from the initial state of RG_W and traversing its transitions, thus it represents a possible evolution or behaviour (indeed, each path of RG_W is a trace of W by definition). If we can show that for each path in TS_* , with $*$ $\in \{BC(W), PDDL(W), SMV(W)\}$, there is a *equivalent* (see later) path in RG_W , then the encoding of the reachability problem for W in the $*$ language is sound, as the tool only generates DAW-net behaviours. Likewise, if for each path in RG_W there is a path in TS_* , then the encoding is complete as all DAW-net behaviors are captured by TS_* . This property is called *trace equivalence* (Baier & Katoen, 2008), and is of utmost importance for us as trace-equivalent transition systems are indistinguishable by linear properties, hence including reachability properties. In other words, if RG_W is trace equivalent with TS_* , a reachability property in TS_* is true if and only if it is true in RG_W , which allows us to use the above tools for solving our trace repair problem.

In order to formalize trace equivalence we shall first define what does it mean for two paths, generated by different transition system, to be equivalent.

Let $\pi : s_0 \xrightarrow{t_1} s_1 \xrightarrow{t_2} \dots \xrightarrow{t_n} s_n$ and $\pi' : s'_0 \xrightarrow{t'_1} s'_1 \xrightarrow{t'_2} \dots \xrightarrow{t'_n} s'_n$ be paths generated by RG and TS_* respectively. We say that π and π' are *equivalent* if $s_0 = s'_0$ and for each $i \in \{1, \dots, n\}$ we have that $s_i = s'_i$ and $t_i = t'_i$. Technically however, states of TS_* encode information on the state of the net (M, η) in different ways, and possibly also use some additional information for technical reasons. Let us then call $enc_*(M, \eta)$ the encoding of (M, η) , and $enc_*(t)$ the encoding of t in the language $*$ $\in \{BC, PDDL, SMV\}$. Then π is equivalent to π' (w.r.t. enc_*) if $enc_*(s_0) = s'_0$ and for each $i \in \{1, \dots, n\}$ we have that $enc_*(s_i) = s'_i$ and $enc_*(t_i) = t'_i$.

Definition 17 (Trace equivalence). *Let RG_W be the reachability graph of a DAW-net W and TS_* a transition system, then RG_W and TS_* are trace equivalent iff there is an injective function enc_* from the states and transitions of RG_W to those of TS_* s.t.*

1. for each path $s_0 \xrightarrow{t_1} s_1 \xrightarrow{t_2} \dots \xrightarrow{t_n} s_n$ of RG_W ,

$$enc_*(s_0) \xrightarrow{enc_*(t_1)} enc_*(s_1) \xrightarrow{enc_*(t_2)} \dots \xrightarrow{enc_*(t_n)} enc_*(s_n)$$

is a path of TS_* ; and

2. for each path $s'_0 \xrightarrow{t'_1} s'_1 \xrightarrow{t'_2} \dots \xrightarrow{t'_n} s'_n$ of TS_* there is a path $s_0 \xrightarrow{t_1} s_1 \xrightarrow{t_2} \dots \xrightarrow{t_n} s_n$ of RG_W s.t. $enc_*(s_0) = s'_0$, and $enc_*(s_i) = s'_i$ $enc_*(t_i) = t'_i$ for $1 \leq i \leq n$.

Recall that the reachability problem, as defined in Corollary 1, expresses whether a specific state (M_e, η') can be reached from the initial one (M_s, η_s) . Given Definition 16, it immediately follows that if RG_W and TS_* are trace equivalent, then, if a state is reachable in RG_W , it is also reachable in TS_* by performing the very same transitions and vice-versa, thus they satisfy the same reachability properties.

In order to exploit the specific algorithms to verify reachability in the “target” system for solving the equivalent problem in DAW-net we need also to show that the corresponding

language for specifying the problem is expressive enough to capture all and only the states corresponding to e.g. the final states of DAW-net.⁷

In the next sections, for each language $* \in \{\text{BC}, \text{PDDL}, \text{SMV}\}$ we prove that RG_W and TS_* are trace equivalent, and therefore they can be used to solve the above mentioned reachability problem.

6. The encoding in Action Languages

In this section will show how the transition system underlying DAW-net can be represented by means of the so called *action languages* (Lifschitz, 1999), in order to use the ASP based solver CLINGO to verify reachability properties. In our work will be focusing on the BC language as introduced in (Lee, Lifschitz, & Yang, 2013), but the same idea can be applied to different frameworks (e.g. DLV^K (Eiter, Faber, Leone, Pfeifer, & Polleres, 2003)).

6.1 The BC action language

An action description B in the language BC includes a finite set of symbols: *fluent*, and *action constants*. Each fluent constant is associated to a finite set, of cardinality greater or equal than 2, called the domain. Boolean fluents have the domain equal to $\{\text{TRUE}, \text{FALSE}\}$. An atom is an expression of the form $f = v$, where f is a fluent constant, and v is an element of its domain.

A BC description B contains *static* and *dynamic laws*; the former are expressions of the type

$$A_0 \text{ if } A_1, \dots, A_n \text{ ifcons } A_{n+1}, \dots, A_m \quad (6)$$

while the second are of the form

$$A_0 \text{ after } A'_1, \dots, A'_n \text{ ifcons } A_{n+1}, \dots, A_m \quad (7)$$

with $0 \leq n \leq m$, where A_0, A_1, \dots, A_m are atoms, and A'_1, \dots, A'_n can be atoms or action constants.

Intuitively, static laws assert that each state satisfying A_1, \dots, A_n and where A_{n+1}, \dots, A_m can be assumed satisfies A_0 as well. In dynamic laws, instead, A'_1, \dots, A'_n must be satisfied in the preceding state. An action constant is satisfied if the action is being “executed” in a specific state.

Initial states can be specified with statements of the form

$$\text{initially } A \quad (8)$$

where A is an atom, that is, an expression of the form $f = v$.

Semantics of a BC description B is given in terms of a translation into a sequence $P_\ell(B)$ of ASP programs where atoms are of the form $i : A$ with $0 \leq i \leq \ell$, and A being a BC fluent or action constant. Each static law (6) is translated as the set of ASP rules with strong negation

$$i : A_0 \leftarrow i : A_1, \dots, i : A_n, \sim \neg(i : A_{n+1}), \dots, \sim \neg(i : A_m)$$

7. In general we might be interested to verify reachability of a specific state which is not necessarily terminal, or a specific subset of the terminal states – e.g. the ones with a specific value for one of the variables – however, this is outside the scope of this paper.

where $0 \leq i \leq \ell$. Each dynamic law (7) is translated as

$$i + 1 : A_0 \leftarrow i : A_1, \dots, i : A_n, \sim \neg(i + 1 : A_{n+1}), \dots, \sim \neg(i + 1 : A_m)$$

where $0 \leq i < \ell$. If A_0 is **false** the corresponding rules are constraints.

Initial state constraints of the form **initially** $f=v$ force the value of fluent f in the first state and are therefore translated as

$$0 : f = v$$

In addition to the rules corresponding to static / dynamic laws and initial constraints, the ASP program contains rules to set the initial state by nondeterministically assigning the values of fluents

$$0 : f = v \vee \neg(0 : f = v)$$

for each fluent f and elements v of its domain; force total knowledge on action execution

$$i : a \vee \neg(i : a)$$

for each action constant a and $i < \ell$; and finally existence and uniqueness of values for the fluents

$$\begin{aligned} & \leftarrow \sim(i : f = v_1), \dots, \sim(i : f = v_k) \\ \neg i : f = v & \leftarrow i : f = w \end{aligned}$$

for any fluent f , v, v_1, \dots, v_k, w elements of its domain, $v \neq w$, and $i \leq \ell$.

Given a stable model M_ℓ of $P_\ell(\mathbf{B})$, we indicate as $\nu_{\text{BC}}^i(M_\ell)$ the variable assignments in the corresponding state $0 \leq i \leq \ell$:

$$\nu_{\text{BC}}^i(M_\ell) = \{f \mapsto o \mid f \in \mathcal{F}, i:f=o \in M_\ell\}$$

Note that the above rules on fluents guarantee that $\nu_{\text{BC}}^i(M_\ell)$ is well defined: total and functional.

Although not originally introduced in (Lee et al., 2013), we introduce the implicit transition system induced by a BC description in order to relate it with the DAW-net reachability graph as explained in Section 5.

Definition 18 (BC transition system). *Let \mathbf{B} be a BC description over the set \mathcal{F} of fluent, and \mathcal{A} of action constants. Let $S_{\mathbf{B}}$ be the set of all total assignments for fluents in \mathcal{F} satisfying the domain restrictions. Then its transition system $TS_{\text{BC}(\mathbf{B})} = (\mathcal{A}, S, S_0, \delta)$ is defined as:*

- $S_0 \subseteq S_{\mathbf{B}}$ is the set of initial states

$$S_0 = \{\nu_{\text{BC}}^0(M_0) \mid M_0 \text{ is a stable model of } P_0(\mathbf{B})\}$$

- $S \subseteq S_{\mathbf{B}}$ and $\delta \subseteq S_{\mathbf{B}} \times 2^{\mathcal{A}} \times S_{\mathbf{B}}$ are the minimum sets s.t.:
 - $S_0 \subseteq S$;
 - if $M_{\ell+1}$ is a stable model of $P_{\ell+1}(\mathbf{B})$ for some $\ell \geq 0$, and $\nu_{\text{BC}}^\ell(M_{\ell+1}) \in S$; then $\nu_{\text{BC}}^{\ell+1}(M_{\ell+1}) \in S$ and $(\nu_{\text{BC}}^\ell(M_{\ell+1}), \{a \in \mathcal{A} \mid \ell:a \in M_{\ell+1}\}, \nu_{\text{BC}}^{\ell+1}(M_{\ell+1})) \in \delta$

While states of a BC transition system can be easily compared with DAW-net reachability graph states, BC enables concurrent and “empty” transitions ($\delta \subseteq S_B \times 2^A \times S_B$); therefore the encoding should take that into account (see Section 6.2 for details).

The tool we use in Section 9 to compute reachability on problems expressed with the BC language is the Coala compiler and CLINGO ASP solver (Gebser, Grote, & Schaub, 2010). Within the tool, the final states to define the reachability problem can be specified using a set of (possibly negated) atoms using **finally** statements analogous to the **initially** previously described in Equation (8).

6.2 Encoding DAW-net in BC

Given a DAW-net W , its encoding in the BC, denoted as $BC(W)$, introduces: (i) a fluent for each variable in W , with the domain corresponding to the variable domain plus a special constant `null` (representing the fact that a variable is not assigned); (ii) a boolean fluent for each place; and (iii) an action constant for each transition. To simplify the notation we will use the same constant names. Moreover, without loss of generality we assume that guards are in conjunctive normal form.⁸

Let $\langle \mathcal{D}, \mathcal{N}, wr, gd \rangle$ be a DAW-net as defined in Definition 9, where $\mathcal{N} = \langle P, T, F \rangle$ and $\mathcal{D} = \langle \mathcal{V}, \Delta, dm, ord \rangle$. Let $\mathcal{V}' \subseteq \mathcal{V}$ be the (finite) set of variables appearing in the model, and $adm(v)$ be the *active domain* of $v \in \mathcal{V}'$, that is, $adm(v) = \bigcup_{t \in T} wr(t)(v)$. The encoding $BC(W)$ starts by representing in BC key properties of transactions in a DAW-net:

- All fluents are “inertial”, that is their value propagates through states unless their value is changed by the transitions:

$$v=o \text{ after } v=o \text{ ifcons } v=o \quad \text{for } v \in \mathcal{V}' \text{ and } o \in adm(v) \cup \{\text{null}\} \quad (9)$$

$$p=o \text{ after } p=o \text{ ifcons } p=o \quad \text{for } p \in P \text{ and } o \in \{\text{true}, \text{false}\} \quad (10)$$

- Transitions modify the value of places and variables; for each transition $t \in T$:

$$p=\text{false after } t \quad \text{for all } p \in \bullet t \setminus t^\bullet \quad (11)$$

$$p=\text{true after } t \quad \text{for all } p \in t^\bullet \setminus \bullet t \quad (12)$$

$$v=d \text{ after } t \text{ ifcons } v=d \quad \text{for all } (v,d) \in wr(t) \quad (13)$$

$$v=\text{null after } t \quad \text{for all } v \text{ s.t. } wr(t)(v) = \emptyset \quad (14)$$

$$\text{false after } t \text{ ifcons } v=d \quad \text{for all } v \in \mathcal{V}' \text{ s.t. } wr(t)(v) \neq \emptyset, d \in \{\text{null}\} \cup adm(v) \setminus wr(t)(v) \quad (15)$$

Note that equation (15) ensures that after the execution of a transition the value of a variable modified by it must be among the legal values.

- Transitions cannot be executed in parallel:

$$\text{false after } t, s \quad \text{for } (t, s) \in T \times T \text{ and } t \neq s \quad (16)$$

- Transitions cannot be executed unless input places are active:

$$\text{false after } t, p=\text{false} \quad \text{for } t \in T \text{ and } p \in \bullet t \quad (17)$$

8. We do not expect huge guards, therefore we ignore the potential exponential blowup due to the conversion to CNF of arbitrary queries.

- Initially, all the places but the start one are false and variables are unassigned

$$\mathbf{initially\ start=true} \quad (18)$$

$$\mathbf{initially\ p=false} \quad \text{for } p \in P \text{ and } p \neq \text{start} \quad (19)$$

$$\mathbf{initially\ v=null} \quad \text{for } v \in \mathcal{V}' \quad (20)$$

- In BC a state progression does not require an actual action. To disallow this scenario we need to track states that have been reached without the execution of an action. This is achieved by including an additional boolean fluent `trans` that is true if a state has been reached after a transition:

$$\mathbf{trans=true\ after\ t} \quad \text{for } t \in T \quad (21)$$

$$\mathbf{initially\ trans=true} \quad (22)$$

We now describe the encoding in BC of DAW-net queries. Due to the rule-based essence of BC we consider a Disjunctive Normal Form characterisation of the DAW-net queries defined in Definition 7): for any DAW-net query Φ there is an equivalent normal form $\bar{\Phi}^9$ s.t. $\bar{\Phi} = \bigvee_{i=1}^k t_1^i \wedge \dots \wedge t_{\ell_i}^i$ where each term t_j^i is either of the form $v = o$ or $\neg\text{def}(v)$. To prove the existence of such formula $\bar{\Phi}$ we consider the disjunctive normal form of Φ , and to each term not in the form $v = o$ or $\neg\text{def}(v)$ we apply the following equivalences that are valid when considering the finite domain data model \mathcal{D}_W (which we can consider without loss of generality because of Lemma 2):

$$\begin{aligned} v_1 = v_2 &\equiv \bigvee_{o \in \text{dm}(v_1) \cap \text{dm}(v_2)} v_1 = o \wedge v_2 = o \\ \neg(v_1 = v_2) &\equiv \bigvee_{(o_1, o_2) \in \text{dm}(v_1) \times \text{dm}(v_2), o_1 \neq o_2} v_1 = o_1 \wedge v_2 = o_2 \\ \neg(v = o) &\equiv \bigvee_{o \in \text{dm}(v) \setminus \{o\}} v = o \\ v \leq o &\equiv \bigvee_{(o', o) \in \text{ord}(v)} v = o \\ o \leq v &\equiv \bigvee_{(o, o') \in \text{ord}(v)} v = o \\ v_1 \leq v_2 &\equiv \bigvee_{(o_1, o_2) \in \text{ord}(v_1) \cap \text{ord}(v_2)} v_1 = o_1 \wedge v_2 = o_2 \\ \text{def}(v) &\equiv \bigvee_{o \in \text{dm}(v)} v = o \end{aligned}$$

Given a DAW-net query Φ , and its Disjunctive Normal Form characterisation $\bar{\Phi} = \bigvee_{i=1}^k t_1^i \wedge \dots \wedge t_{\ell_i}^i$, we define the translation Φ^{BC} of Φ as $\bigvee_{i=1}^k t_1^{i\text{BC}} \wedge \dots \wedge t_{\ell_i}^{i\text{BC}}$; where $v = o^{\text{BC}} \mapsto v = o$ and $\neg\text{def}(v)^{\text{BC}} \mapsto v = \text{null}$.¹⁰

9. $\bar{\Phi}$ is not unique.
 10. This rather inefficient explicit representation is convenient for the proofs, and it can be optimised in the actual ASP implementation via grounding techniques and explicit representation of the orderings by means of ad hoc predicates.

Having encoded the queries we can now encode the fact that transitions cannot be executed unless the guards are satisfied. This condition is encoded as a set of dynamic constraints preventing the execution unless at least one of the clauses of the guard is not satisfied. For each transition t s.t. $\mathbf{gd}(t) \not\equiv \text{true}$ we consider $\neg\mathbf{gd}(t)^{\text{BC}} = \bigvee_{i=1}^k t_1^{i \text{ BC}} \wedge \dots \wedge t_{\ell_i}^{i \text{ BC}}$ and include the following dynamic constraints:

$$\begin{aligned}
 & \mathbf{false \ after \ } t, t_1^{1 \text{ BC}}, \dots, t_{\ell_1}^{1 \text{ BC}} \\
 & \dots \\
 & \mathbf{false \ after \ } t, t_1^{k \text{ BC}}, \dots, t_{\ell_k}^{k \text{ BC}}
 \end{aligned} \tag{23}$$

The set of final states for the reachability problem can be specified as the ones in which all place fluents are FALSE but the one corresponding to the sink:

$$\begin{aligned}
 & \mathbf{finally \ sink = true} \\
 & \mathbf{finally \ } p = \mathbf{false} \quad \text{for } p \in P \text{ and } p \neq \text{sink}
 \end{aligned} \tag{24}$$

6.3 Correctness and completeness of the BC encoding

Let be $\text{BC}(W)$ be the BC encoding of a DAW-net W . In this section we are going to show that $TS_{\text{BC}(W)}$ and the reachability graph RG_W of W are trace equivalent as per Definition 17. This section reports all the main steps involved, while further technical details and proofs are presented in Appendix B.

To simplify the proofs we first relate sequences of valid firings in W of arbitrary length ℓ to stable models of $P_\ell(\text{BC}(W))$. In fact, the definition of $TS_{\text{BC}(W)}$ (Definition 18) is based on the notion of stable models.

To map sequences of valid firings into stable models we observe that stable models $P_\ell(\text{BC}(W))$ include both the fluent assignments and the information regarding the action constants that “cause” a specific state transition; therefore we split the mapping in two parts to separate these two distinct aspects:

Definition 19. *Let $\rho = (M_0, \eta_0) \xrightarrow{t_0} (M_1, \eta_1) \xrightarrow{t_2} \dots \xrightarrow{t_{\ell-1}} (M_\ell, \eta_\ell)$ be a sequence of valid firings in W (with $\ell \geq 0$).¹¹ The mapping $\Phi_i(\rho)$ (for $0 \leq i \leq \ell$) is defined as following:*

$$\begin{aligned}
 \Phi_i^\nu(\rho) &= \{i : p = \text{TRUE}, \neg(i : p = \text{FALSE}) \mid p \in P, M_i(p) > 0\} \cup \\
 & \quad \{i : p = \text{FALSE}, \neg(i : p = \text{TRUE}) \mid p \in P, M_i(p) = 0\} \cup \\
 & \quad \{i : v = o, \neg(i : v = \text{null}) \mid v \in \mathcal{V}', \eta_i(v) = o\} \cup \\
 & \quad \{i : v = \text{null} \mid v \in \mathcal{V}', \eta_i(v) \text{ is undefined}\} \cup \\
 & \quad \{\neg(i : v = o) \mid v \in \mathcal{V}', o \in \text{adm}(v), \eta_i(v) \neq o \text{ or } \eta_i(v) \text{ is undefined}\} \cup \\
 & \quad \{i : \text{trans} = \text{TRUE}, \neg(i : \text{trans} = \text{FALSE})\} \\
 \Phi_i^\tau(\rho) &= \begin{cases} \emptyset & \text{for } i \geq \ell \\ \{i : t_i\} \cup \{\neg(i : t) \mid t \in T, t \neq t_i\} & \text{for } i < \ell \end{cases} \\
 \Phi_i(\rho) &= \Phi_i^\nu(\rho) \cup \Phi_i^\tau(\rho)
 \end{aligned}$$

11. We consider (M_0, η_0) a sequence of size 0.

Note that Definition 19, together with Definition 18, enable us to say that the firing $(M_i, \eta_i) \xrightarrow{t_i} (M_{i+1}, \eta_{i+1})$ in ρ corresponds to the $TS_{BC(W)}$ transition $\nu_{BC}^i(\Phi_i^\nu(\rho)) \xrightarrow{\{t \mid t \in \Phi_i^+(\rho)\}} \nu_{BC}^{i+1}(\Phi_{i+1}^\nu(\rho))$. Moreover, the definition of $\Phi_i^\nu(\cdot)$ depends only on the state (M_i, η_i) and not on the selected path ρ , which provides just the “witness” for being a state of the reachability graph.

Before proving completeness and correctness we show that the states of the two transition systems preserve the satisfiability of $\mathcal{L}(\mathcal{D}_W)$ queries:

Lemma 3. *Let W be a DAW-net model, and $\rho = (M_0, \eta_0) \xrightarrow{t_0} (M_1, \eta_1) \xrightarrow{t_1} \dots \xrightarrow{t_{\ell-1}} (M_\ell, \eta_\ell)$ be a sequence of valid firings of W (with $\ell \geq 0$). For every state (M_i, η_i) ($0 \leq i \leq \ell$) and query $\Phi \in \mathcal{L}(\mathcal{D}_W)$*

$$\mathcal{D}_i, \eta_i \models \Phi \text{ iff } \Phi_i(\rho) \models i : \Phi^{BC}$$

where $i : \Phi^{BC}$ is the formula where $i : \cdot$ is placed in front of each term t^{BC} , and $\Phi_i(\rho) \models i : t^{BC}$ iff $i : t^{BC} \in \Phi_i(\rho)$

The proof is by induction on $\mathcal{L}(\mathcal{D}_W)$.

We are now ready to state the completeness of BC encoding.

Lemma 4 (Completeness of BC encoding). *Let W be a DAW-net and $BC(W)$ its encoding, then for any $\ell \geq 0$ if ρ is a sequence of valid firings of W of length ℓ , then $\bigcup_{i=0}^{\ell} \Phi_i(\rho)$ is a stable model of $P_\ell(BC(W))$.*

To prove correctness we first prove few properties about the reachability graph $TS_{BC(W)}$.

Lemma 5. *Let $TS_{BC(W)} = (\mathcal{A}, S, S_0, \delta)$, then*

1. $S_0 = \{s_0\}$;
2. if $(s, A, s') \in \delta$ then $|A| = 1$.

By using this lemma, given a DAW-net model W , we can compare the transition system corresponding to its reachability graph $RG_W = (T, \bar{S}, \bar{s}_0, \bar{\delta})$ to $TS_{BC(W)}$ (Definition 18). This because it guarantees that there is a single initial state, and the transitions (the δ of Definition 18) contain exactly one action each.¹² The following definition introduces the mapping from stable models to sequences of valid firings.

Definition 20. *For any stable model M_ℓ of $P_\ell(BC(W))$ with $\ell \geq 0$ we also define the “inverse” relation $\Phi_i^-(M_\ell) = (M_i, \eta_i)$ for $0 \leq i \leq \ell$:*

$$\begin{aligned} M_i &= \{(p, 1) \mid p \in P, \nu_{BC}^i(M_\ell)(p) = \text{TRUE}\} \cup \\ &\quad \{(p, 0) \mid p \in P, \nu_{BC}^i(M_\ell)(p) = \text{FALSE}\} \\ \eta_i &= \{(v, o) \mid v \in V', \nu_{BC}^i(M_\ell)(v) = o, o \neq \text{null}\} \end{aligned}$$

and the transition $\tau_i(M_\ell)$ for $0 \leq i < \ell$:

$$\tau_i(M_\ell) = a \quad \text{s.t. } (i : a) \in M_\ell$$

where a is an action constant in $P_\ell(BC(W))$.

¹². BC allows transitions without or with multiple parallel actions.

Lemma 6 (Correctness of BC encoding). *Let W be a DAW-net and $\text{BC}(W)$ its encoding, then for any $\ell \geq 0$ if M_ℓ is a stable model of $P_\ell(\text{BC}(W))$, then*

$$\rho = \Phi_0^-(M_\ell) \xrightarrow{\tau_0(M_\ell)} \Phi_1^-(M_\ell) \xrightarrow{\tau_1(M_\ell)} \dots \xrightarrow{\tau_{\ell-1}(M_\ell)} \Phi_\ell^-(M_\ell)$$

is a sequence of valid firings of W of length ℓ .

We are now ready to state the main result of this Section, namely trace equivalence between $\mathcal{G}_{\text{BC}}(\text{BC}(W))$ and $RG_W = (T, \bar{S}, \bar{s}_0, \bar{\delta})$. We do that by instantiating the abstract notion of enc_* introduced in Section 5 to BC, and by using the completeness and correctness lemmata 4 and 6.

Theorem 3. *Let W be a DAW-net model, $RG_W = (T, \bar{S}, \bar{s}_0, \bar{\delta})$ its reachability graph, then $TS_{\text{BC}(W)}$ and RG_W are trace equivalent.*

Proof. To demonstrate trace equivalence we define the mapping $\text{enc}_{\text{BC}}(\cdot)$ as following: for each (M, η)

$$\begin{aligned} \text{enc}_{\text{BC}}((M, \eta)) = & \{(p, \text{TRUE}) \mid (p, 1) \in M\} \cup \\ & \{(p, \text{FALSE}) \mid (p, 0) \in M\} \cup \\ & \{(v, \text{null}) \mid v \text{ not in the domain of } \eta\} \cup \\ & \eta \end{aligned}$$

and $\text{enc}_{\text{BC}}(t) = \{t\}$ for each transition in T .

Note that, by construction, for each sequence of valid firings $(M_0, \eta_0) \xrightarrow{t_0} (M_1, \eta_1) \xrightarrow{t_1} \dots \xrightarrow{t_{\ell-1}} (M_\ell, \eta_\ell)$, $\text{enc}_{\text{BC}}((M_i, \eta_i)) = \nu_{\text{BC}}^i(\Phi_i^-(\rho))$, and for any stable model M_ℓ of $P_\ell(\text{BC}(W))$ $\text{enc}_{\text{BC}}(\Phi_\ell^-(M_\ell)) = \nu_{\text{BC}}^\ell(M_\ell)$. Therefore the proof is a consequence of the lemmata 4 and 6.

Moreover, it is immediate to verify that the set of final states as characterised in Equation (24) corresponds to all and only the final states of the DAW-net model mapped via $\text{enc}_{\text{BC}}(\cdot)$. \square

The trace equivalence demonstrated by Theorem 3 enables the use of ASP based planning to verify the existence of a plan satisfying the termination conditions expressed in Equation 24. In the case that the model is a trace workflow W^τ , then the plan would exhibit a valid completion of the trace τ ; while the failure of finding such a plan would imply that the trace do not admit any completion.

7. The encoding in Classical Planning

Automated planning has been widely employed as a reasoning framework to verify reachability properties of finite transition systems. It is therefore natural to include it among the verification tools that we considered for our evaluation.

In order to be as general as possible without committing ourselves to a specific planner we based our encoding on the de facto standard namely PDDL (Ghallab, Knoblock, Wilkins, Barrett, Christianson, Friedman, Kwok, Golden, Penberthy, Smith, Sun, & Weld, 1998), however to simplify the formal discussion we will adopt the *state-variable* representation (Ghallab, Nau, & Traverso, 2016) instead of the more common *classical*. There is no loss of generality, since is known that state-variable planning domains can be translated to classical planning domains with at most a linear increase in size.

7.1 State-Variable Representation of Planning Problems

A detailed description of the planning formalism is outside the scope of this paper and the reader is referred to the relevant literature – e.g. (Ghallab et al., 2016) – however in this section we briefly outline the main concepts for our setting.

Given a set of objects B , a *state variable* is a symbol v and an associated domain $\text{Range}(v) \subseteq B$.¹³ The (finite) set of all the variables is denoted by X , a *variable assignment* over X maps each variable v to one of the objects in $\text{Range}(v)$. In addition, we will consider a set of *rigid relations* R over B that represent non-temporal relations among elements of the domain.

An *atom* is a term of the form TRUE , $r(z_1, \dots, z_n)$, or $v = z_1$; where $r \in R$, $v \in X$, and z_i are either constants or parameters.¹⁴ A *literal* is a positive or negated atom.

An *action template* is a tuple $\alpha = (\text{head}(\alpha), \text{pre}(\alpha), \text{eff}(\alpha))$: $\text{head}(\alpha)$ is an expression $\text{act}(z_1, \dots, z_k)$ where act is an (unique) action name and z_i are parameters with an associated finite $\text{Range}(z_i) \subseteq B$, $\text{pre}(\alpha)$ is a set of literal (the preconditions), and $\text{eff}(\alpha)$ is a set of assignments $v \leftarrow z$ where z is a constant or parameter. All parameters in α must be included in $\text{head}(\alpha)$. A *state-variable action* (or just action) is a ground template action where the parameters are substituted by constants from their corresponding ranges.

Definition 21. A classical planning domain is a triple $\Sigma = (S, A, \gamma)$ where

- S the set of variable assignments over X ;
- A is the set of all actions;
- $\gamma : S' \times A' \rightarrow S$ with $S' \times A' = \{(s, a) \in S \times A \mid \text{pre}(a) \text{ is satisfied in } S\}$ and

$$\begin{aligned} \gamma(s, a) = & \{(v, w) \mid v \leftarrow w \in \text{eff}(a)\} \\ & \{(v, w) \in s \mid \forall w' \in B \ v \leftarrow w' \notin \text{eff}(a)\}. \end{aligned}$$

Definition 22. A plan is a finite sequence of actions $\pi = \langle a_1, \dots, a_n \rangle$, or the empty plan $\langle \rangle$. A plan is applicable to a state s_0 if there are states s_1, \dots, s_n s.t. $\gamma(s_{i-1}, a_i) = s_i$ (empty plans are always applicable).

If π is applicable to s , we extend the notation for γ to include plans $\gamma(s, \pi)$: where $\gamma(s, \langle \rangle) = s$ and $\gamma(s, \pi.a) = \gamma(\gamma(s, \pi), a)$.¹⁵

Definition 23. A state-variable planning problem is a triple $P = (\Sigma, s_0, g)$ where Σ is a planning domain, s_0 the initial state, and g a set of ground literals called the goal.

A solution for P is a plan π s.t. $\gamma(s_0, \pi)$ satisfies g .

To simplify the encoding introduced in Section 7.2 we will consider an extended language for the precondition of action templates. This extended language enables the use of boolean combinations as well as atoms in which state variables can be used in place of constants, that is, $v = v'$, or $r(v_1, \dots, v_n)$. Planning domains using the extended language can be encoded into classical planning domains by considering the disjunctive normal form of the

13. In this context we restrict to state variables as 0-ary functions.

14. We use the term *parameter* to denote mathematical variables, to avoid confusion with state variables.

15. $\langle a_1, \dots, a_n \rangle.a$ is defined as $\langle a_1, \dots, a_n, a \rangle$.

original preconditions and introducing a new template for each conjunct.¹⁶ The extended syntax concerning state variables can be dealt by introducing a (new) parameter for each variable in the precondition appearing within a rigid relation or in the left hand side of an equality. This new parameter z_v for the variable v will be substituted to the original occurrence of v and the new term $v = z_v$ is included in the precondition.

The above definitions can be re-casted for the extended language, and plans for the extended and corresponding classical planning domains can be related via a surjection projecting the “duplicated” template (introduced by the DNF conversion) into the original ones.

7.2 Encoding DAW-net in PDDL

The main idea behind the encoding of a DAW-net W in PDDL, denoted with $\text{PDDL}(W)$, is similar to the one presented in Section 6.2, where state variables are used to represent both places and DAW-net variables, and actions represent transitions. The main differences lays in the fact that: (i) nondeterminism in the assignments must be encoded by means of grounding of action schemata; and (ii) the language for preconditions does not include disjunction, and thus more than a schema may correspond to a single transition.

Let $W = \langle \mathcal{D}, \mathcal{N}, \text{wr}, \text{gd} \rangle$ be a DAW-net as defined in Definition 9, where $\mathcal{N} = \langle P, T, F \rangle$ and $\mathcal{D} = \langle \mathcal{V}, \Delta, \text{dm}, \text{ord} \rangle$. $\mathcal{V}' \subseteq \mathcal{V}$ be the (finite) set of variables appearing in the model, and $\text{adm}(v)$ be the *active domain* of $v \in \mathcal{V}'$, that is, $\text{adm}(v) = \bigcup_{t \in T} \text{wr}(t)(v)$. The encoding $\text{PDDL}(W)$ is defined by introducing: a state variable for each $v \in \mathcal{V}'$, with domain corresponding to the domain of v plus a special constant `null`; and a boolean state variable for each place $p \in P$. To simplify the notation we will use the same constant names. In addition to the state variables we introduce a rigid binary relation `ord` to encode the partial ordering of Definition 5:

$$\text{ord} = \bigcup_{\Delta_i \in \Delta} \{(o, o') \in \Delta_i^2 \mid o \leq_{\Delta_i} o'\} \quad (25)$$

and unary relations $\text{wr}_{t,v}$ to encode the range $\text{wr}(t)(v)$ of Definition 9:

$$\begin{aligned} \text{wr}_{t,v} &= \{o \mid o \in \text{wr}(t)(v)\} && \text{for all } t \in T \text{ and } v \in \mathcal{V} \text{ s.t. } \text{wr}(t)(v) \neq \emptyset \\ \text{wr}_{t,v} &= \{\text{null}\} && \text{for all } t \in T \text{ and } v \in \mathcal{V} \text{ s.t. } \text{wr}(t)(v) = \emptyset \end{aligned} \quad (26)$$

Note that Lemma 2 guarantees the finiteness of `ord`, while $\text{wr}(t)(v)$ are finite by definition.

¹⁶. This transformation might introduce an exponential blow-up of the domain due to the DNF transformation. However, this can be prevented in the overall conversion from DAW-net by imposing the DNF restriction in the original model as well.

Given a DAW-net query Φ (Definition 7) we denote with Φ^{PDDL} its translation in the PDDL preconditions, defined as:

$$\begin{aligned}
 true^{\text{PDDL}} &= \text{TRUE} \\
 \text{def}(v)^{\text{PDDL}} &= \neg(v = \text{null}) \\
 v = t_2^{\text{PDDL}} &= \neg(v = \text{null}) \wedge (v = t_2) \\
 t_1 \leq t_2^{\text{PDDL}} &= \text{ord}(t_1, t_2) \\
 \neg\Phi_1^{\text{PDDL}} &= \neg\Phi_1^{\text{PDDL}} \\
 \Phi_1 \wedge \Phi_2^{\text{PDDL}} &= \Phi_1^{\text{PDDL}} \wedge \Phi_2^{\text{PDDL}}
 \end{aligned} \tag{27}$$

Analogously to the previous section we introduce an action template for each transition in W where preconditions, in addition to the guard, include the marking of the network (token in the places). The selection of the actual value to be assigned to the updated variables is selected by means of action parameters (one for each variable). Formally, let t be a transition in W , and $\{v_1, \dots, v_k\}$ the variables in the domain of $\text{wr}(t)$; then $\text{PDDL}(W)$ includes the action template $\alpha_t = (\text{head}(\alpha_t), \text{pre}(\alpha_t), \text{eff}(\alpha_t))$ defined as following:

$$\begin{aligned}
 \text{head}(\alpha_t) &= t(z_{v_1}, \dots, z_{v_k}) \\
 \text{pre}(\alpha_t) &= \text{gd}(t)^{\text{PDDL}} \wedge \bigwedge_{v \in \{v_1, \dots, v_k\}} \text{wr}_{t,v}(z_v) \wedge \bigwedge_{p \in \bullet t} (p = \text{TRUE}) \\
 \text{eff}(\alpha_t) &= \bigwedge_{v \in \{v_1, \dots, v_k\}} (v = z_v) \wedge \bigwedge_{p \in \bullet t \setminus t^\bullet} (p = \text{FALSE}) \wedge \bigwedge_{p \in t^\bullet} (p = \text{TRUE})
 \end{aligned} \tag{28}$$

Analogously to the BC case, the goal of the planning problem is described by the set of ground literals corresponding to the state in which all place fluents are FALSE but the one corresponding to the sink:

$$\begin{aligned}
 \text{sink} &= \text{TRUE} \\
 p &= \text{FALSE} \quad \text{for } p \in P \setminus \{\text{sink}\}
 \end{aligned} \tag{29}$$

7.3 Correctness and completeness of the PDDL encoding

To show that the transition systems underlying a DAW-net W and its encoding $\text{PDDL}(W)$ are equivalent we define a bijective mapping between the states of the two systems and we show that for each valid firing in W there is a ground action “connecting” the image of the states in γ and vice-versa.

Let (M, η) be an arbitrary state of W , we define the mapping Ψ to states of $\text{PDDL}(W)$ as:

$$\begin{aligned}
 \Psi(M, \eta) &= \{v \mapsto v[\eta] \mid v \in \mathcal{V}, v[\eta] \neq v\} \cup \\
 &\quad \{v \mapsto \text{null} \mid v \in \mathcal{V}, v[\eta] = v\} \cup \\
 &\quad \{p \mapsto \text{TRUE} \mid p \in P, M(p) > 0\} \cup \\
 &\quad \{p \mapsto \text{FALSE} \mid p \in P, M(p) = 0\}
 \end{aligned} \tag{30}$$

Since we restrict to 1-safe nets it is easy to show that Ψ is a bijection, moreover:

Lemma 7. *Let W be a DAW-net model, for every state (M, η) of W and query $\Phi \in \mathcal{L}(\mathcal{D}_W)$*

$$\mathcal{D}, \eta \models \Phi \text{ iff } \Psi(M, \eta) \models \Phi^{\text{PDDL}}$$

Proof. The statement can be easily proved by induction on $\mathcal{L}(\mathcal{D}_W)$ formulae and by using Equation (27) for base cases. \square

Now we can show that valid firings of DAW-net corresponds to transitions in the planning domain

Lemma 8. *Let W a DAW-net model and $\text{PDDL}(W)$ its encoding into a planning domain:*

1. *if $(M, \eta) \xrightarrow{t} (M', \eta')$ is a valid firing of W , then there is a ground action a_t of α_t s.t. $(\Psi(M, \eta), a_t, \Psi(M', \eta')) \in \gamma$;*
2. *if there is a ground action a_t of α_t s.t. $(s, a_t, s') \in \gamma$ then $\Psi^-(s) \xrightarrow{t} \Psi^-(s')$ is a valid firing of W .*

The proof is in Appendix C.

To compare the PDDL encoding of W with the DAW-net reachability graph $RG_W = (T, \bar{S}, \bar{s}_0, \bar{\delta})$ introduced in Definition 16) we consider the transition system $TS_{\text{PDDL}(W)}$ derived from the planning domain $\text{PDDL}(W)$.

Definition 24. *Let W a DAW-net model. The transition system $TS_{\text{PDDL}(W)} = (T, S_{\text{PDDL}}, s_0, \delta_{\text{PDDL}})$ is defined as follow*

- $s_0 = \Psi(\bar{s}_0)$
- S_{PDDL} is the minimum set that includes s_0 and if $s \in S_{\text{PDDL}}$ and $(s, a_t, s') \in \gamma$ then $s' \in S_{\text{PDDL}}$
- $\delta_{\text{PDDL}} = \{(s, t, s') \in S_{\text{PDDL}} \times T \times S_{\text{PDDL}} \mid (s, a_t, s') \in \gamma, a_t \text{ is a ground action of } \alpha_t\}$

Theorem 4. *Let W be a DAW-net model, then RG_W and $TS_{\text{PDDL}(W)}$ are trace equivalent.*

Proof. We define $\text{enc}_{\text{PDDL}}(\cdot)$ as the restriction of Ψ (as defined in Definition. 30) to the set of states of RG_W , and the identity w.r.t. the transitions. We show that $\text{enc}_{\text{PDDL}}(\cdot)$ satisfies the properties of Definition 17 by induction on the length of the paths.

The base case of a path of length 0 is satisfied because $s_0 = \text{enc}_{\text{PDDL}}(\bar{s}_0) = \Psi(\bar{s}_0)$ is a state in $TS_{\text{PDDL}(W)}$ by Definition 24.

For the inductive step we just need to consider the last element of the paths:

1. If $s_0 \xrightarrow{t_1} s_1 \xrightarrow{t_2} \dots s_{n-1} \xrightarrow{t} s_n$ is a path in RG_W , then $s_{n-1} \xrightarrow{t} s_n$ is a valid firing, therefore by Lemma 8 there is a ground action a_t of α_t s.t. $(\Psi(s_{n-1}), a_t, \Psi(s_n)) \in \gamma$. By Definition 24 $(\Psi(s_{n-1}), t, \Psi(s_n))$ is a transition of $TS_{\text{PDDL}(W)}$, and so is $(\text{enc}_{\text{PDDL}}(s_{n-1}), \text{enc}_{\text{PDDL}}(t), \text{enc}_{\text{PDDL}}(s_n))$.
2. If $\{(s_0, t_1, s_1), \dots, (s_{n-1}, t, s_n)\} \subseteq \delta_{\text{PDDL}}$ then there is a ground action a_t of α_t s.t. $(s_{n-1}, a_t, s_n) \in \gamma$; therefore $\Psi^-(s_{n-1}) \xrightarrow{t} \Psi^-(s_n)$ is a valid firing of W by Lemma 8. So $\Psi^-(s_{n-1})$ and $\Psi^-(s_n)$ are states of RG_W and $\text{enc}_{\text{PDDL}}(\Psi^-(s_i)) = s_i$ by construction.

Moreover, it is immediate to verify that the set of states as characterised by the goal in 29 corresponds to all and only the final states of the DAW-net model mapped via $\text{enc}_{\text{PDDL}}(\cdot)$. \square

The trace equivalence demonstrated by Theorem 7.3 enables the use of classical planning tools to verify the existence of a plan satisfying the termination conditions expressed in Equation 29. In the case that the model is a trace workflow W^τ , then the plan would exhibit a valid completion of the trace τ ; while the failure of finding such a plan would imply that the trace do not admit any completion.

8. The encoding in nuXmv

Model checking tools take as input a specification of system executions and a temporal property to be checked on such executions. The output is *true* when the system satisfies the property, otherwise *false* is returned as well as a *counterexample* showing one (among the possibly many) evolution that falsifies the formula. In NUXMV, the system is specified in an input language inspired by SMV (McMillan, 1993) and the properties in a temporal language, such as linear-time temporal logic (LTL, see (Baier & Katoen, 2008) Chapter 5, for an introduction).

Let W be a DAW-net and RG_W its reachability graph. We show how to encode RG_W in the NUXMV input language and the reachability problem as a LTL formula such that, when processed by NUXMV, an execution reaching the final state is generated in the form of a counterexample, if any. When the DAW-net is a trace workflow W^τ , then the counterexample is a possible completion of τ .

8.1 The nuXmv input language

Many of the available tools for traditional model checking, including NUXMV, deal with infinite-executions systems, and hence they require systems transition relation to be left-total. On the contrary, our DAW-SYS models are based on workflow-nets which by definition represent process which usually terminates. Also, in NUXMV transitions have no label. For this reason, we adapt the reachability graph so as to have infinite paths and no transition labels, and we call such a modified reachability graph RG_W^\sim . This kind of transformations are customary using tools requiring infinite-executions to describe finite-execution ones. Despite RG_W^\sim being semantically different from RG_W , it is also possible to tweak the formal properties to be verified in such a way that the adaptation is sound and complete. From here on we abstract from such technical subtleties and refer to (De Giacomo, De Masellis, & Montali, 2014) for the details.

Intuitively, RG_W^\sim is obtained from RG_W by: (i) moving transition labels in the source state; (ii) adding a fresh self-looping (with transition *ended*) state s_ϵ and (iii) adding a (*last*) transition from every sink state of RG_W to s_ϵ , as the following definition formalizes.

Definition 25. Let $RG_W = (T, \bar{S}, \bar{s}_0, \bar{\delta})$ be a reachability graph of a DAW-net. We build a transition system $RG_W^\sim = (T \cup \{\textit{last}, \textit{ended}\}, S, S_0, \delta)$ as follows:

- the set of states are triples (M, η, t) with $t \in T \cup \{\textit{last}, \textit{ended}\}$ where we distinguish the special state $s_\epsilon = (M_\epsilon, \eta_\epsilon, \textit{ended})$ where M_ϵ and η_ϵ are arbitrary;
- set S_0 is built as follows. If exist $(M_0, \eta_0) \xrightarrow{t} (M', \eta') \in \bar{\delta}$ then $S_0 = \{(M_0, \eta_0, t) \mid \exists t, M', \eta'. (M_0, \eta_0) \xrightarrow{t} (M', \eta') \in \bar{\delta}\}$, otherwise $S_0 = \{(M_0, \eta_0, \textit{last})\}$.

- S and $\delta \subseteq S \times S$ are the least sets s.t.:
- $S_0 \cup \{s_\epsilon\} \subseteq S$;
- $s_\epsilon \rightarrow s_\epsilon \in \delta$;
- if $(M, \eta) \xrightarrow{t} (M', \eta') \in \bar{\delta}$ then:
 - if $(M', \eta') \xrightarrow{t'} (M'', \eta'') \in \bar{\delta}$ then $s = (M, \eta, t) \rightarrow s' = (M', \eta', t') \in \delta$ and $s, s' \in S$;
 - otherwise $s = (M, \eta, t) \rightarrow s' = (M', \eta', \text{last}) \in \delta$ and $(s', s_\epsilon) \in \delta$ and $s, s' \in S$;
- if $\bar{\delta} = \emptyset$ then $(M_0, \eta_0, \text{last}) \rightarrow s_\epsilon \in \delta$ (notice that if $\bar{\delta} = \emptyset$ then $S_0 = \{(M_0, \eta_0, \text{last})\}$).

We now describe from a very high-level perspective the syntax and semantics of NUXMV, i.e., how a transition system $TS_{\text{SMV}(m)}$ is built given syntactic description of a so-called *module* m . In general, NUXMV allows for several modules descriptions, but in our case one is enough to represent $RG_{\tilde{W}}$.

A module is specified by means of several sections and, in particular:

- A VAR section where a set of variables and their domains are defined. Each complete assignment of values to variables is a state of the module.
- A INIT section where a formula representing the initial state(s) of the system is defined.
- A TRANS section where the transition relation between states is specified in a declarative way by means of a list of ordered rules $\langle r_1, \dots, r_n \rangle$ in a switch-case block¹⁷. Each rule r_i has the form $(pre_i, post_i)$, where pre_i is (a formula representing) a precondition on the current state and $post_i$ is (a formula representing) a postcondition on the next state (and possibly the current state). In NUXMV, postconditions are expressed by means of primed variables, namely $\hat{v}' = \text{value}$ means that in the next state variable \hat{v} is assigned to value . Intuitively, state q' is a successor of q iff i is the smallest index such that q models pre_i and (q, q') models $post_i$. With slight notational abuse we write $q \models \varphi$ when the current state (recall that a state is an assignment of values to variables) satisfies a formula containing only unprimed variables and we write $(q, q') \models \varphi'$ when the pair current state, next state satisfies a formula containing both primed and unprimed variables.

Definition 26. *Let m be a NUXMV module and $stVar$ be the set of variables in its VAR section. Module m defines a transition system $TS_{\text{SMV}(m)} = (stVar, Q, Q_0, \rho)$ where:*

- Q is the set of states, each representing a total assignment of domain values to variables;
- the initial states Q_0 are those satisfying the formula in the INIT section;
- States Q and transition function $\rho \subseteq Q \times Q$ are defined by mutual induction as follows:
 - $Q_0 \subseteq Q$;
 - if $q \in Q$ and i is the smallest index such that $q \models pre_i$ (such an i is required to exist by NUXMV), then $(q, q') \in \rho$ and q' is such that $(q, q') \models post_i$.

17. There are several equivalent ways to describe a the transition relation in NUXMV, we just choose one.

8.2 Encoding of DAW-net in nuXmv

We present our encoding of W in NUXMV each section of the module $\text{SMV}(W)$.

VAR section. We have:

- one boolean variable $\hat{p} \in \hat{P}$ for each place $p \in P$;
- one \hat{tr} variable which ranges over $T \cup \{\text{last}, \text{ended}\}$;
- one variable $\hat{v} \in \hat{V}$ for each $v \in \mathcal{V}$ ranging over $\text{dm}(v) \cup \{\text{undef}\}$ where undef it is a special value not contained in any of the $\text{dm}(v)$.

Let $\text{enc}_{\text{SMV}}(M, \eta)$ be the encoding for DAW-net markings and assignments to formulas in NUXMV of the kind:

$$\bigwedge_{p \in P} \hat{p} = \{\text{TRUE}, \text{FALSE}\} \bigwedge_{v \in \mathcal{V}} \hat{v}_i = d$$

where $\hat{p} = \text{TRUE}$ iff $M(p) = 1$ and $\hat{p} = \text{FALSE}$ otherwise and $\hat{v}_i = d$ iff $\eta(v)$ is defined and $\eta(v) = d$ or $d = \text{undef}$ if $\eta(v)$ is undefined. Besides, let $\text{enc}_{\text{SMV}}(t)$ be the encoding of DAW-net transition t to the NUXMV formula $\hat{tr} = t$, for $t \in T \cup \{\text{last}, \text{ended}\}$.

Remark 1. *The pair of functions $\text{enc}_{\text{SMV}}(M, \eta)$ and $\text{enc}_{\text{SMV}}(t)$ are a bijection between states S of $RG_{\tilde{W}}$ and states Q of $TS_{\text{SMV}(W)}$.*

INIT section. The init section is the following formula:

$$\begin{aligned} & \text{enc}_{\text{SMV}}(M_0, \eta_0) \wedge \\ & \hat{tr} \neq \text{ended} \wedge \\ & \bigwedge_{t \in T} \hat{tr} = t \rightarrow \text{pre}(t) \wedge \\ & \hat{tr} = \text{last} \rightarrow (\bigwedge_{t \in T} \neg \text{pre}(t)) \end{aligned}$$

where $\text{pre}(t)$ is a formula representing the preconditions for executing t according to the semantics of the DAW-net: it thus consists of guard $\text{gd}(t)$ as well as preconditions on the input set $\bullet t$.

TRANS section. We have:

- one rule for each $t \in T$ which has:

$$\begin{aligned} & \text{pre} := \text{enc}_{\text{SMV}}(t) \\ \\ & \text{post} := \quad \forall p \in \bullet t \setminus t^\bullet. \text{next}(\hat{p}) = \text{FALSE} \wedge \\ & \quad \forall p \in t^\bullet \setminus \bullet t. \text{next}(\hat{p}) = \text{TRUE} \wedge \\ & \quad \text{next}(\hat{p}) = \hat{p} \text{ otherwise } \wedge \\ & \quad \forall v \in \text{dm}(\text{wr}(t)) \wedge \text{wr}(t)(v) \neq \emptyset. \text{next}(\hat{v}) \in \text{wr}(t)(v) \wedge \\ & \quad \forall v \in \text{dm}(\text{wr}(t)) \wedge \text{wr}(t)(v) = \emptyset. \text{next}(\hat{v}) = \text{undef} \wedge \\ & \quad \text{next}(\hat{v}) = \hat{v} \text{ otherwise } \wedge \\ & \quad \bigwedge_{t \in T} \text{next}(\hat{tr}) = t \rightarrow \text{next}(\text{pre}(t)) \wedge \\ & \quad \text{next}(\hat{tr}) = \text{last} \rightarrow (\bigwedge_{t \in T} \neg \text{next}(\text{pre}(t))) \wedge \\ & \quad \text{next}(\hat{tr}) = \text{ended} \rightarrow \hat{tr} = \text{last} \end{aligned}$$

where $\text{next}(\text{pre}(t))$ is likewise formula $\text{pre}(t)$ but with all its variables primed (thus referring to the next state). Intuitively, the postcondition contains: (i) a condition on the (next) values of pre- and post-sets of t ; (ii) a condition on the (next) values

of the variables according to function wr of the DAW-net and (iii) conditions on the (next) value of variable \hat{tr} : next value of \hat{tr} is t , only if the (next) state is consistent with the preconditions for executing t .

- a rule for $\hat{tr} = \text{last}$:

$$pre := \text{enc}_{\text{SMV}}(\text{last})$$

$$post := \text{next}(\hat{tr}) = \text{ended} \wedge \\ \text{next}(\text{enc}_{\text{SMV}}(M_\epsilon, \eta_\epsilon))$$

- a rule for looping in the ended state with $\hat{tr} = \text{ended}$:

$$pre := \text{enc}_{\text{SMV}}(\text{ended})$$

$$post := \text{next}(\hat{tr}) = \text{ended} \wedge \\ \text{next}(\text{enc}_{\text{SMV}}(M_\epsilon, \eta_\epsilon))$$

We notice that since $\bigcup pre_i$ form a partition for all possible values of \hat{tr} we have that for each assignment to stVar , i.e., state q , there exists a *unique* i such that $q \models pre_i$. In other words, the *ordering* of rules does not matter.

8.3 Correctness and completeness of the nuXmv encoding

As introduced in Section 5, we prove soundness and completeness of the encoding by proving trace equivalence between $RG_{\tilde{W}}$ and the transition system $TS_{\text{SMV}(W)}$ generated by module $\text{SMV}(W)$. A standard technique to prove trace equivalence amounts, in fact, to provide a bisimulation relation between states of $RG_{\tilde{W}}$ and $TS_{\text{SMV}(W)}$ as the latter implies the former (Baier & Katoen, 2008). A *bisimulation relation* has a co-inductive definition which says that states of the two transition systems are bisimilar if: (i) they satisfy the same *local* properties; (ii) every transition from one state can be mimicked by the other state and their arrival states are again in bisimulation (and vice-versa). If every initial states of $RG_{\tilde{W}}$ is in bisimulation with an initial state of $TS_{\text{SMV}(W)}$, then they are trace equivalent, as any trace start from an initial state.

Definition 27. $RG_{\tilde{W}} = (T, S, S_0, \delta)$ be the reachability graph of a DAW-net and $TS_{\text{SMV}(W)} = (\text{stVar}, Q, q_0, \rho)$ be the nuXMV transition system of its encoding. We define bisimulation relation $B \subseteq S \times Q$ as follows: $(s = (M, \eta, t), q) \in B$ implies:

1. $q = \text{enc}_{\text{SMV}}(M, \eta) \wedge \text{enc}_{\text{SMV}}(t)$;
2. for all $s \rightarrow s' \in \delta$ there exists $q' \in Q$ such that $q \rightarrow q' \in \rho$ and $(s', q') \in B$ and
3. for all $q \rightarrow q' \in \rho$ there exists $s' \in S$ such that $s \rightarrow s' \in \delta$ and $(s', q') \in B$.

We prove that there exists a bisimulation between $RG_{\tilde{W}}$ and $TS_{\text{SMV}(W)}$ by showing one.

Theorem 5. Let $R \subseteq S \times Q$ be such that $((s = (M, \eta, t), q) \in R)$ iff $q = \text{enc}_{\text{SMV}}(M, \eta) \wedge \text{enc}_{\text{SMV}}(t)$. Then:

1. R is a bisimulation relation;

2. for each $s_0 \in S_0$ there exists $q_0 \in Q_0$ such that $(s_0, q_0) \in R$ and vice versa.

Proof. We first prove 1. Let $s' = (M', \eta', t')$ be such that $s \rightarrow s' \in \delta$, then we have to prove that there exists a q' such that $q \rightarrow q' \in \rho$ and $(s', q') \in R$. We have three cases:

- $t \in T$: then, by inspecting the corresponding rule in the TRANS section, it is immediate to see that q and s are such that $M'(p) \leftrightarrow \hat{p}$ (such values are deterministic). Let us now consider the value of a generic variable $v \in \mathcal{V}$ assigned nondeterministically by η' : we have again three cases: either $v \in \text{dm}(\text{wr}(t))$ and $\text{wr}(t)(v) \neq \emptyset$ hence its new value will be in $\text{wr}(t)(v)$; or $\text{dm}(\text{wr}(t))$ and $\text{wr}(t)(v) = \emptyset$ hence its new value will be undef; or $v \notin \text{dm}(\text{wr}(t))$ and its new value will be the same as the old one. Such cases are coded in conjunction in the rule: let us assume the first one applies with value $v' = c$, by Definition 26 then there exists q' such that $(q, q') \models \text{post}_i$ and $q'(v) = c$; for the other two cases the choice is deterministic, either $q'(v) = \text{undef}$ or $q'(v) = q(v)$. It follows that there exists a q' that agrees on $\text{enc}_{\text{SMV}}(M', \eta')$. It remains to show that there exists a q' such that also $q'(\hat{t}r) = t'$: if $t' \in T$, then it means that M' and η' satisfy the preconditions of the firing of t' , meaning that $t \rightarrow \text{next}(\text{pre}(t))$ is satisfied, hence there exists such a q' (there exists at least one, the choice of t' is nondeterministic); if $t' = \text{last}$ then, from Definition 25 we have that no t'', M'', η'' such that $(M', \eta') \xrightarrow{t''} (M'', \eta'')$ exist, i.e., $\bigwedge_{t \in T} \neg \text{next}(\text{pre}(t))$ holds, hence $q'(\hat{t}r) = \text{last}$. Notice that we assumed $t \in T$, hence it is never the case that $t' = \text{ended}$. Since q' agrees with s' on the markings, the values of variables and the transition, it follows that $(s', q') \in R$.
- $t = \text{last}$, then from Definition 25 we have that $s' = s_\epsilon$ and $t' = \text{ended}$. Since $q(\hat{t}r) = \text{last}$, then rule $t = \text{last}$ applies in the encoding and from its definition, $(s', q') \in R$ immediately follows.
- $t = \text{ended}$, then from Definition 25 we have that $s' = s_\epsilon$ and $t' = \text{last}$. Since $q(\hat{t}r) = \text{ended}$, then rule $t = \text{ended}$ applies in the encoding and from its definition, $(s', q') \in R$ immediately follows.

Now, let q' such that $q \rightarrow q' \in \rho$. We have to prove that there exists a $s' = (M', \eta', t')$ such that $s \rightarrow s' \in \delta$ and $q' = \text{enc}_{\text{SMV}}(M', t') \wedge \text{enc}_{\text{smvts}}(t)$. We have three cases: $t \in T$, $t = \text{last}$ and $t = \text{ended}$. Analogous considerations of the previous case hold, as the rules in the encoding are not only sound, but also complete.

We now show 2. Given Remark 1, it is enough to show that $(s, q) \in R$ implies $s \in S_0 \leftrightarrow q \in Q_0$. We first prove the left to right direction, i.e., that $(s = (M_0, \eta_0, t), q) \in R$ implies $q \in Q_0$, namely that formula $\text{enc}_{\text{SMV}}(M_0, \eta_0) \wedge \text{enc}_{\text{SMV}}(t)$ satisfies the first conjunct of the INIT section. By definition of INIT, the first conjunct is satisfied, thus it remains to prove $\text{enc}_{\text{SMV}}(t)$. By Definition 16 and Definition 25 either (1) $t \in T$ or (2) $t = \text{last}$. If (1) then there is no valid firing in (M, η) , which entails that $\bigwedge_{t \in T} \neg \text{pre}(t)$ and if (2), then t can fire in (M_0, η_0) hence its preconditions are satisfied, which means $\text{enc}_{\text{SMV}}(M_0, \eta_0)$ satisfies $\text{pre}(t)$. In both cases the formula in the INIT section is satisfied and $q \in Q_0$.

The right to left direction is analogous by noticing that $\text{pre}()$ trivially encodes the preconditions for firing t . \square

We conclude the section by observing that, thanks to Remark 1, we are able to capture with the simple formula $\Phi := \text{end} = \text{TRUE} \bigwedge_{p \in P \setminus \{\text{end}\}} \hat{p} = \text{FALSE}$ all and only states in $TS_{\text{SMV}(W)}$ that correspond to final states in the DAW-net. Therefore, when NUXMV checks

the LTL formula expressing property: “a Φ -state is never visited”, it returns *true* if no final state can be reached and *false* and a path to a final state otherwise. When the DAW-net is a trace workflow $TS_{SMV(W)}$, then “true” means that τ is not compliant with W , otherwise the counterexample is a possible completion of τ .

9. Evaluation

In this section we aim at evaluating the three investigated solvers (i.e., CLINGO, FAST-DOWNWARD and NUXMV), and the respective paradigms (i.e., answer set programming, automated planning and model checking), on the task of completing incomplete traces of data-aware workflows. In detail, we are interested to answer the following two research questions:

- RQ1.** What are the performance of the three solvers when accomplishing the task of completing incomplete traces by leveraging also data payloads?
- RQ2.** What are the main factors impacting the solvers’ capability to deal with the task of completing incomplete traces?

The first research question aims at evaluating and comparing the performance of the three solvers, both in terms of capability of returning a result within a reasonable amount of time and, when a result is returned, in terms of the time required for solving the task. The second research question, instead, aims at investigating which factors could influence the solvers’ capability to deal with such a task (e.g., the level of incompleteness of the execution traces, the characteristics of the traces, the size of the process model, the domain of the data payload associated to the events, or the data payload itself).

In order to answer the above research questions by evaluating the three investigated solvers in different scenarios, we carried out two different evaluations: one based on synthetic logs and a second one based on real-life logs. Indeed, on the one hand, the synthetic log evaluation (Subsection 9.1) allows us to investigate the solvers’ capabilities on logs with specific characteristics. On the other hand, the real-life log evaluation (Subsection 9.2) allows us to investigate the capability of the solvers to deal with real-life problems. In the next subsections, we detail the two evaluations (Subsections 9.1 and 9.2) and we investigate some threats to the validity of the obtained results (Subsection 9.3).

All the experiments have been carried out on a Kubernetes cluster running over a pool of virtual machines hosted on hardware based on Intel Xeon X5650 2.67GHz processors. For the container running the experiments a single CPU has been allocated since the reasoning systems we used do not exploit multiprocessing. For each run we set up a (real) time limit of one hour, and a memory limit of 8GiB.¹⁸ The code used for the experiments is available for download from <https://doi.org/10.5281/zenodo.3459656>. The directory `experiments` includes the models and experiments descriptions, as well as a `README.md` file with the details on how to run them.

18. Given the single task nature of the container, the CPU time is only marginally lower than the real time.

9.1 Synthetic Log Evaluation

In order to investigate the performance of the three solvers in different settings, we built synthetic models with different characteristics. For each of them, we identified different traces with different features and, in turn, for each trace, we considered different levels of trace completeness ranging from the complete trace (100% complete) to the empty trace (0% complete). In the next subsections we describe the datasets, the procedure and the metrics used, as well as the obtained results.

9.1.1 DATASET, PROCEDURE AND METRICS

We built five different synthetic models $M1 - M5$ with different characteristics and of different size. In detail, $M1$ is the model reported in Figure 7. The model deals with 5 different variables: *number*, *first*, *second*, *third*, *fourth* and *fifth*. Activity A is able to write four variables (*first*, *second*, *third* and *fourth*), activity I writes variable *number* and, finally, activity G writes variable *fifth*.

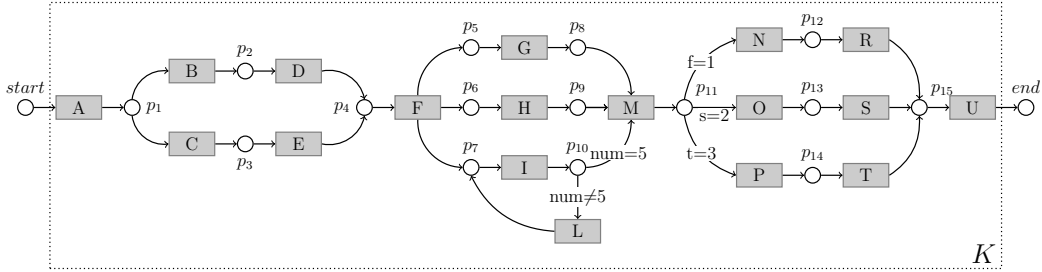


Figure 7: Model $M1$ and its core K

Figure 7 also shows the guards related to model $M1$. In detail, the exit condition from the loop ($L - I$) is variable *number* equals to 5 at I , while whenever the value of *number* written by I is different from 5, the loop is exited and M is executed. Moreover, variables *first*, *second* and *third* determine the branch to be taken at the end of the model. In detail, if at M *first* is equal to 1, the branch with N is taken, if *second* is equal to 2, the branch O is taken, while if *third* is equal to 3, the branch starting with P is taken¹⁹. Whenever more than one of the three conditions are verified, one of the three branches can be non-deterministically taken.

We built the synthetic models $M2 \dots M5$ starting from model $M1$ and creating replicas of $M1$. For notation simplicity, we denote with K the core part of $M1$, i.e., the model obtained by excluding the start and the end event (see Figure 8(a)) and we enumerate the replicas of K in $M2 \dots M5$ with progressive number. In detail, $M2$ concatenates two replicas of K , labelled as $K0$ and $K1$ (see Figure 8(b)), $M3$ uses the parallel construct to compose together two replicas of K , namely $K0$ and $K1$ (see Figure 8(c)). $M4$ concatenates three replicas of K , namely $K0$, $K1$ and $K2$ (see Figure 8(d)). Finally, in $M5$, three replicas of K are composed in an exclusive choice. In detail, one of the two alternative branches concatenates two replicas of K , i.e., $K0$ and $K1$, while the other exclusive branch is composed of a single replica of K , named $K2$ (see Figure 8(e)). Different replicas of K

¹⁹. For the sake of readability *first*, *second* and *third* are shortened in f , s and t in Figure 7.

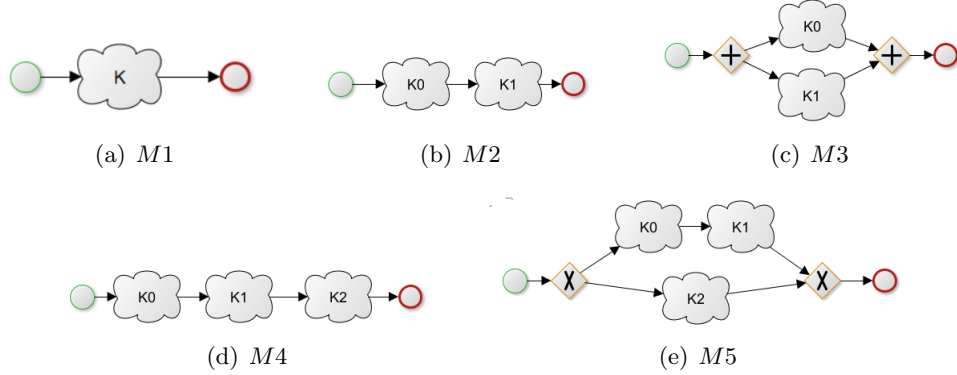


Figure 8: Models $M1, \dots, M5$ used in the evaluation

do not share variables, which are renamed in order to avoid interactions that could limit the compatible traces.

For each of the five models $M1 \dots M5$, we generated 8 different execution traces of different length and with different characteristics. In detail, we generated 4 compliant traces and 4 non-compliant traces. Out of the 4 compliant traces, in two traces (trace types $T3$ and $T4$) the loop(s) of K replica(s) is(are) never executed, while in the remaining two, the loop is executed twice. Moreover, in two (trace types $T1$ and $T3$) out of the four cases of the compliant traces, the path followed by the trace is deterministically driven by the values associated by the activity A to the variables *first*, *second* and *third*, while in the remaining two traces more than one branch could be activated based on the values of the variables at activity M . Also in the case of non-compliant traces, two traces (trace types $T5$ and $T7$) do not execute any loop iteration, while the other 2 traces iterated over the loop two times. Finally, out of the four non-compliant traces, two (trace type $T5$ and $T6$) were not compliant only because of control flow violation (mutually exclusive paths were executed in the same trace), while the remaining two traces were non-compliant due to data constraint violations (a path that was not activated by the corresponding guard is executed). Table 9.1.1 reports, for each $tracet_{ij}$ of model M_i and trace type T_j , some information about its characteristics.

Finally, for each of the 40 traces, we considered five different levels of completeness²⁰, ranging from 100% (i.e., the complete trace) to 0% (i.e., the empty trace). Table 9.1.1, column *Completeness* shows four level of completeness (100%, 75%, 50% and 25%) for each of the considered traces and the corresponding characteristics. In the table we did not report the information related to the empty traces (0% of completeness) which is the same for all the 8 different types of traces. For each of the three solvers, hence, we carried on, in total, 165 runs (5 models, 8 traces for each model and 4 levels of incompleteness for each trace and model, as well as 5 empty traces per model).

We evaluated the first research question (**RQ1**) by computing, for each solver (i) the number/percentage of runs the solver is able to return; (ii) when a solution is returned, the

20. The execution traces have been made incomplete by removing events so as to preserve the non-compliance of traces (compliance is always preserved), as well as the capability for a solver to reconstruct the missing events.

Model	Trace type	Trace	Cycle It.	Det.	Conf.	Cause	Compl.	Length	Trace type	Trace	Cycle It.	Det.	Conf.	Cause	Compl.	Length
M1	T ₁	t ₁₁	N	Y	C	100%	11	T ₅	t ₁₅	Y	N	NC	CF	100%	13	
						75%	8							25%	10	
						50%	5							50%	7	
						25%	3							75%	4	
M1	T ₂	t ₁₂	N	N	C	100%	11	T ₆	t ₁₆	2	Y	NC	CF	100%	17	
						75%	8							75%	13	
						50%	5							50%	10	
						25%	3							25%	6	
M1	T ₃	t ₁₃	2	Y	C	100%	15	T ₇	t ₁₇	N	Y	NC	CF+DF	100%	11	
						75%	11							75%	8	
						50%	8							50%	5	
						25%	4							25%	3	
M1	T ₄	t ₁₄	2	N	C	100%	15	T ₈	t ₁₈	2	Y	NC	CF+DF	100%	15	
						75%	11							75%	11	
						50%	8							50%	8	
						25%	4							25%	4	
M2	T ₁	t ₂₁	N	Y	C	100%	22	T ₅	t ₂₅	N	Y	NC	CF	100%	26	
						75%	16							75%	20	
						50%	10							50%	14	
						25%	6							25%	8	
M2	T ₂	t ₂₂	N	N	C	100%	22	T ₆	t ₂₆	2	Y	NC	CF	100%	34	
						75%	16							75%	26	
						50%	10							50%	20	
						25%	6							25%	12	
M2	T ₃	t ₂₃	2	Y	C	100%	30	T ₇	t ₂₇	N	Y	NC	CF+DF	100%	22	
						75%	22							75%	16	
						50%	16							50%	10	
						25%	8							25%	6	
M2	T ₄	t ₂₄	2	N	C	100%	30	T ₈	t ₂₈	2	Y	NC	CF+DF	100%	30	
						75%	22							75%	22	
						50%	16							50%	16	
						25%	8							25%	8	
M3	T ₁	t ₃₁	Y	Y	C	100%	22	T ₅	t ₃₅	N	Y	NC	CF	100%	26	
						75%	16							75%	20	
						50%	10							50%	14	
						25%	6							25%	8	
M3	T ₂	t ₃₂	N	Y	C	100%	22	T ₆	t ₃₆	2	Y	NC	CF	100%	34	
						75%	16							75%	26	
						50%	10							50%	20	
						25%	6							25%	12	
M3	T ₃	t ₃₃	2	Y	C	100%	30	T ₇	t ₃₇	N	Y	NC	CF+DF	100%	22	
						75%	22							75%	16	
						50%	16							50%	10	
						25%	8							25%	6	
M3	T ₄	t ₃₄	2	N	C	100%	30	T ₈	t ₃₈	2	Y	NC	CF+DF	100%	30	
						75%	22							75%	22	
						50%	16							50%	16	
						25%	8							25%	8	
M4	T ₁	t ₄₁	Y	Y	C	100%	33	T ₅	t ₄₅	N	Y	NC	CF	100%	39	
						75%	24							75%	30	
						50%	15							50%	21	
						25%	9							25%	12	
M4	T ₂	t ₄₂	N	Y	C	100%	33	T ₆	t ₄₆	2	Y	NC	CF	100%	51	
						75%	24							75%	39	
						50%	15							50%	30	
						25%	9							25%	18	
M4	T ₃	t ₄₃	2	Y	C	100%	45	T ₇	t ₄₇	N	Y	NC	CF+DF	100%	33	
						75%	33							75%	24	
						50%	24							50%	15	
						25%	12							25%	9	
M4	T ₄	t ₄₄	2	N	C	0%	45	T ₈	t ₄₈	2	Y	NC	CF+DF	0%	45	
						25%	33							25%	33	
						50%	24							50%	24	
						75%	12							75%	12	
M5	T ₁	t ₅₁	Y	Y	C	0%	22	T ₅	t ₅₅	N	Y	NC	CF	0%	13	
						25%	8							25%	10	
						50%	10							50%	7	
						75%	3							75%	4	
M5	T ₂	t ₅₂	N	Y	C	0%	11	T ₆	t ₅₆	2	Y	NC	CF	0%	34	
						25%	16							25%	26	
						50%	10							50%	10	
						75%	6							75%	12	
M5	T ₃	t ₅₃	2	Y	C	0%	30	T ₇	t ₅₇	N	Y	NC	CF+DF	0%	11	
						25%	11							25%	8	
						50%	16							50%	10	
						75%	8							75%	6	
M5	T ₄	t ₅₄	2	N	C	0%	30	T ₈	t ₅₈	2	Y	NC	CF+DF	0%	15	
						25%	22							25%	22	
						50%	16							50%	8	
						75%	8							75%	4	

Table 1: Synthetic trace characteristics

time required by the solver for returning the solution. For **RQ2**, instead, we evaluated the metrics reported above, when changing factors, such as the level of incompleteness of the traces, the characteristics of the traces, as well as the model used for generating the traces.

9.1.2 RESULTS

Only 2 runs out of the 495 runs carried out for the synthetic datasets did not return a result because they exceeded the maximum time threshold. Differently from NUXMV, which exceeded the maximum time limit in 2 cases ($\sim 1\%$ of the runs), both CLINGO and FAST-DOWNWARD were always able to return results.

Figure 9, besides the small percentage of timed-out runs of NUXMV, shows the average CPU time (and the corresponding variance) required by each of the three solvers to return results. We can clearly observe that CLINGO and FAST-DOWNWARD are the solvers with the best performance. Differently from NUXMV, which fails in returning results within the time-out threshold for two runs, the two solvers are able to return results for each of the runs. Moreover, CLINGO and FAST-DOWNWARD have comparable performance in terms of time required for completing incomplete traces (see the overlapping lines of CLINGO and FAST-DOWNWARD in Figure 9). The time required by the two solvers (few seconds on average) is always lower than the time required by NUXMV, which takes, on average, about 8 minutes.

This analysis allows us to assess that, overall, in case of synthetic datasets (of medium size), CLINGO and FAST-DOWNWARD are the most reliable in finding solutions in a reasonable amount of time, and the fastest ones among the three considered solvers. NUXMV has instead lower performance both in terms of capability to return results in a reasonable amount of time, as well as in terms of the average amount of time required to return computations (**RQ1**).

In detail, Figure 9(a) shows the performance of the three solvers for different levels of incompleteness. By looking at the figure, we can observe that the level of incompleteness has an impact on NUXMV: the more the trace is empty, the more the solver is able to return results (i.e., the less number of timed-out runs occur), the faster it is, and the lower the variance of the runs is.

The plot in Figure 9(b), which reports solvers' performance for traces with different characteristics (the empty trace and the 8 types of traces considered), shows that the only runs for which NUXMV is unable to return results relate to traces of type $T5$ and $T6$, i.e., traces non-compliant due to issues related to the control flow. However, the characteristics of these two types of traces, do not seem to impact the average time required by NUXMV to return results. Differently from what happens for timed-out runs, the types of traces that require more time to NUXMV seem to be $T3$ and $T8$.

Finally, Figure 9(c) investigates the impact of different models on the solvers. The plot suggests that all the solvers are able to manage well small and medium models (e.g., $M1$, $M2$ and $M3$). NUXMV has instead difficulties with larger models. In particular, both the two timed-out runs relate to model $M4$, which is the largest model, i.e., the model whose traces are always the longest traces (see Table 9.1.1). Moreover, the runs related to this model that are able to complete, require more time than other runs for returning results.

Differently from NUXMV, none of the three investigated factors affects the performance of CLINGO and FAST-DOWNWARD, which are always able to return a solution in few seconds.

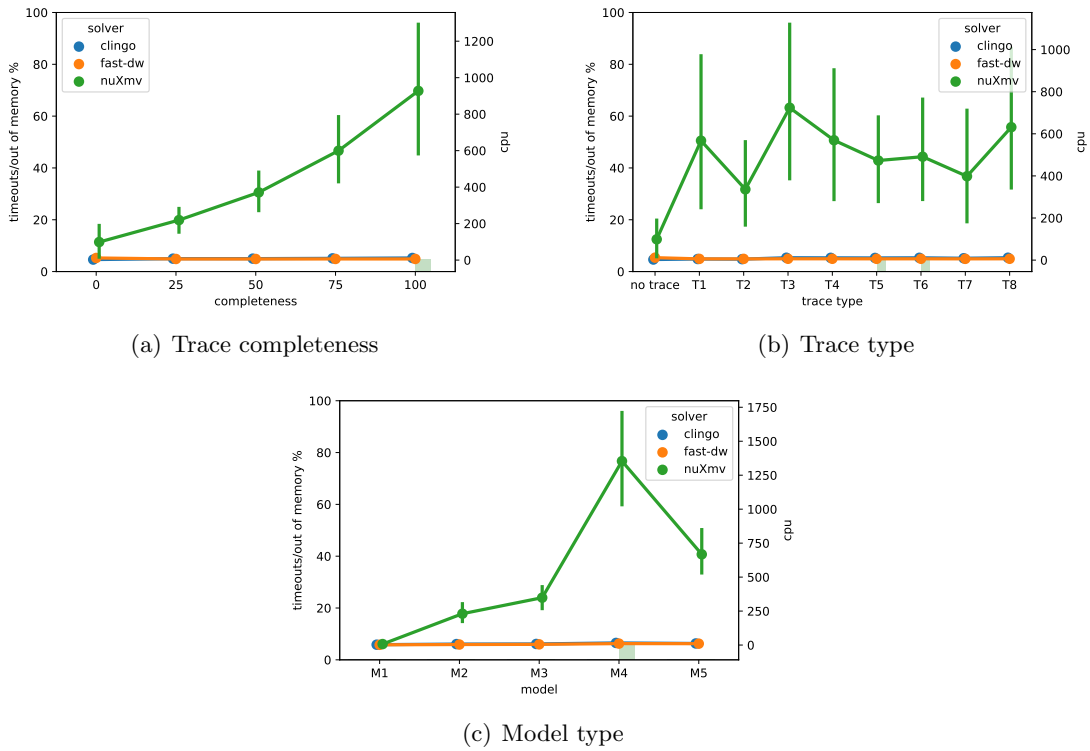


Figure 9: Solver performance for different levels of completeness, different traces and different models. In the plots, lines correspond to CPU time (right scale) and bars to interrupted computations.

Overall, we can assess that, in case of synthetic datasets (of medium size), trace completeness has a strong impact on NUXMV, which performs better on empty traces. NUXMV performance are also influenced by the model size and, in particular, by the length of the traces that it generates: the larger the model and the longer the traces, the worse the performance of the solver. On the contrary, none of the investigated factors has an impact on the CLINGO and FAST-DOWNWARD performance (**RQ2**).

9.2 Real-life Log Evaluation

In order to evaluate the capability of the three solvers to deal with real-life settings, we exercised them on a real world event log. In the next subsections we first describe the considered dataset, the procedure and the metrics used for its evaluation; we then show the obtained results.

9.2.1 DATASET, PROCEDURE AND METRICS

In order to exercise the three solvers in a real-life setting, we resorted to the BPI Challenge 2011 (3TU Data Center, 2011) (BPIC2011) event log. The log pertains to a health-care process and, in particular, contains the executions of a process related to the treatment of patients diagnosed with cancer in a large Dutch academic hospital. The whole event log contains 1143 execution cases and 150291 events distributed across 623 event classes (activities). Each case refers to the treatment of a different patient. The event log contains domain specific attributes that are both case attributes and event attributes. For example, *Age*, *Diagnosis*, and *Treatment_code* are case attributes and *Activity_code*, *Number_of_executions*, *Producer_code*, *Section* and *Group* are event attributes.

The inputs required for accomplishing the task of completing incomplete traces are a process model and incomplete traces. We hence extracted a data-aware Petri Net model and we generated a set of incomplete traces by applying the following procedure:

1. We discovered the data-aware Petri net from the BPIC2011 event log by applying the ProM DATA-FLOW DISCOVERY plugin (de Leoni & van der Aalst, 2013).²¹ The discovered data-aware Petri Net has 355 transitions, 61 places and 710 edges and 4 variables (*Activity_code*, *Producer_code*, *Section* and *Group*), which are read and written by the Petri net transitions, are constrained in 25 guards throughout the data-aware Petri net. An example of a discovered guard is reported in 31.

$$(Producer_code="CHE2")\&\&(Activity_code=="370422") \quad (31)$$

2. For the testing set generation, we randomly selected 9 complete traces - details on the traces and their length are reported in columns *Trace ID* and *Trace length* in Table 2 - from the dataset.
3. For each of the 9 randomly selected traces, 3 incomplete traces - obtained by randomly removing 25%, 50% and 75% of the events are added to the testing set.

21. As input for the DATA-FLOW DISCOVERY plugin we used the Petri Net discovered with the INDUCTIVE MINER algorithm by setting the *Noise threshold* parameter of the plugin to 0.20.

Trace ID	Trace length
<i>t</i> 207	14
<i>t</i> 296	515
<i>t</i> 365	30
<i>t</i> 533	565
<i>t</i> 610	127
<i>t</i> 679	5
<i>t</i> 729	3
<i>t</i> 847	56
<i>t</i> 1132	34

Table 2: BPIC2011 testing set trace length

Scenario	total	CLINGO	FAST-DOWNWARD	NUXMV
full data	37	37 (out-of-memory)	37 (out-of-memory)	22 (timed-out)
restricted domain	37	35 (out-of-memory)	37 (out-of-memory)	12 (timed-out)
no data	37	20 (out-of-memory)	0	12 (timed-out)

Table 3: Timed-out and out-of-memory runs for the three solvers on the BPIC2011

4. The empty trace is added to the testing set.

We evaluated the 3 solvers on each trace of the testing set for a total of 37 runs per solver (9 different traces for 4 different levels of completeness and the empty trace).

Also in this case, we evaluated **RQ1** by computing, for each solver (i) the number/percentage of runs the solver is able to return; (ii) the time required for returning the solution. For the second research question (**RQ2**), instead, we focused on the factors that could impact the results. Besides looking at the degree of trace completeness and at different traces, we also investigated the impact of data and of its domain on the performance of the solvers on large real-life datasets. To this aim, we evaluated the three solvers by computing the metrics reported above in three different settings:

- **full data**: in this setting we considered the data payloads associated to events and we did not apply any restriction to their domain;
- **restricted domain**: in this setting we reduced the domain of the data to the only constant values occurring in the model;²²
- **no data**: in this setting we do not use data at all.

For each setting, we performed the 37 runs described above, we collected the metrics and we compared the results.

9.2.2 RESULTS

Table 3 reports, for each of the three scenarios, the number of timed-out and out-of-memory runs for the three solvers on the BPIC2011. When applying the three solvers to the task of completing incomplete traces of a large real-life event log in the **full data** scenario, the only solver able to return results (in 15 runs out of the 37 ones carried out) is NUXMV; CLINGO and FAST-DOWNWARD fail in each of the executed runs due to out-of-memory issues. Indeed, when applied to large datasets, the CLINGO and FAST-DOWNWARD grounding phase, which likely allows the solver to complete the runs in short time for small datasets, is unable to complete due to a memory space explosion.

²². Please notice that we were able to reduce the domain of the variables to the sole constant values, as we only have guards with equality constraints.

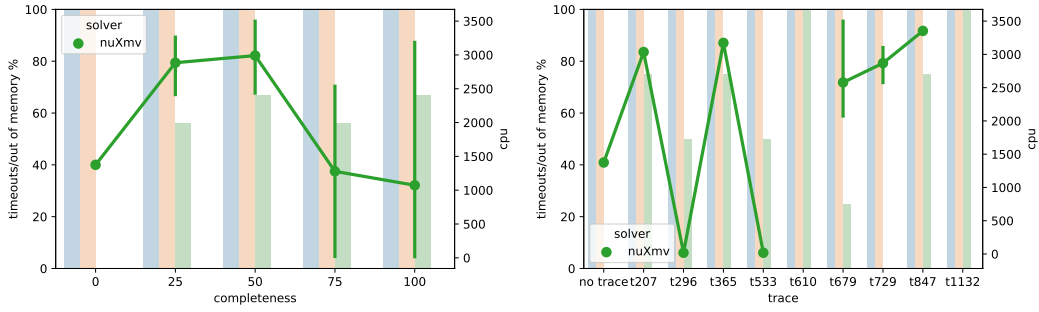


Figure 10: Solver performance for different levels of completeness and different traces of the BPIC2011 event log in the **full data** scenario

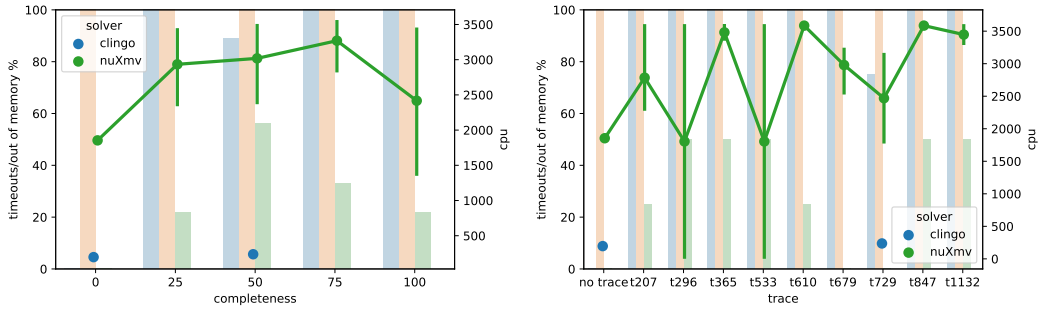


Figure 11: Solver performance for different levels of completeness and different traces of the BPIC2011 event log in the **restricted domain** scenario

Figure 10 plots the performance (percentage of timed-out/out-of-memory runs and average time) of the three solvers for different levels of trace completeness and for the different considered traces (the empty trace and the 9 traces of the BPIC2011). Overall, no general trend or pattern related to the trace completeness and to the trace type can be identified for the 15 runs for which NUXMV was able to return a result and the required time ranges from ~ 20 seconds to one hour. Differently from the synthetic datasets, an opposite trend can be observed with large real-life logs: performance improves with more complete traces. For instance, the timed-out runs and the average time required for 75%-complete traces are less than the timed-out runs and the average time required for the 50%-complete traces. Moreover, we can notice that for a couple of traces - t_{610} and t_{1132} - NUXMV is also not able to provide any result, as the other two solvers.

Figure 11 shows the obtained results for the **restricted domain** scenario. Despite the domain restriction, CLINGO and FAST-DOWNWARD are still unable to return results, while an improvement can be observed in the the capability of NUXMV to return results: the number of timed-out runs decreases of about 50% (from 22 to 12 runs), as reported in Table 3. The trend observed for the timed-out runs is similar to the one of the **full data** scenario (the percentage of timed-out runs is null for the empty traces and reaches its peak

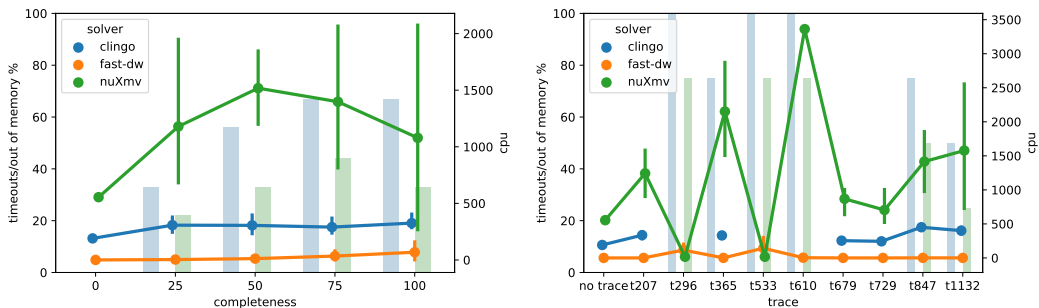


Figure 12: Solver performance for different levels of completeness and different traces of the BPIC2011 event log in the `no data` scenario

for 50%-complete traces), while some differences are registered in terms of required time (e.g., the average time required for 75%-complete traces increases almost of a factor of 3 with respect to the `full data` scenario). If we look at different traces, trends are instead quite different from the `full data` scenario. For instance, in the `restricted domain` scenario, NUXMV was unable to deal only with 20% of the runs for the trace *t610*, which was a problematic trace in the `full data` scenario. On the contrary, the percentage of timed-out runs for trace *t296* (50%) remains unchanged moving from the `full data` to the `restricted domain` scenario.

Finally, the three solvers have been evaluated in the `no data` scenario. In this setting, each of the three solvers is able to return some results (see the last row in Table 3). When data are not taken into account, the performance of the solvers are close to the results obtained with the syntetic datasets: FAST-DOWNWARD is the only one able to deal with all the 37 runs, followed by NUXMV registering 12 timed-out runs and, finally, by CLINGO with 20 out-of-memory runs. Moving from a `restricted domain` to a `no data` scenario has hence a strong impact on the capability of FAST-DOWNWARD and CLINGO to return results, while this is not the case for NUXMV, whose number of timed-out runs remains the same of the `restricted domain` scenario.

The plots in Figure 12 report, besides the percentage of timed-out and out-of-memory runs, also the time required by the solvers for the returned runs. As for the syntetic datasets, FAST-DOWNWARD is the solver with the best performance (it takes from few seconds to few minutes) followed by CLINGO (~ 300 seconds, for the runs it is able to complete) and NUXMV (~ 20 minutes with a high variance). The time required by FAST-DOWNWARD seems to slightly depend on the level of trace incompleteness: the more the trace is complete, the larger the encoded net, the more time FAST-DOWNWARD requires to complete the task.

Overall, when dealing with real-life datasets enriched with data (both with and without domain restrictions), NUXMV is the best solver, while CLINGO and FAST-DOWNWARD are unable to deal with this scenario. When, instead data are not taken into account, FAST-DOWNWARD outperforms the other two solvers both in terms of returned results and of required time. Although NUXMV fails in returning results in less runs (about half) than CLINGO, CLINGO has better performance than NUXMV in terms of required time (**RQ1**).

Moreover, although no general pattern and trend can be identified with respect to trace incompleteness and trace type, data has a strong impact on CLINGO and FAST-DOWNWARD and the data domain affects the capability of NUXMV of returning results (**RQ2**).

Summarizing the above results, we can conclude that NUXMV is the most robust approach as it is able to deal reasonably well with very complex settings as real-life event logs with data. On the other hand, FAST-DOWNWARD, which on large search spaces ends up in out-of-memory errors, on synthetic and on real-life medium-size search spaces outperforms its competitors, carrying on the trace completion task in few seconds (**RQ1**). Concerning **RQ2**, for each of the three solvers, no global trends or patterns related to the level of trace completeness and to the characteristics of the trace have been identified. The only exception is NUXMV and its capability to better deal with incomplete rather than with complete traces for synthetic event logs and in part viceversa for large real-life event logs. The size of the resulting search space and, in particular, leveraging data, has a huge impact on the difficulties of CLINGO and FAST-DOWNWARD to solve the trace completion task.

We remark that, although the focus of the evaluation is on the specific task of completing incomplete traces of data-aware workflows, since the task is recasted to a reachability problem (see Section 4), the problem we are tackling is actually a wider one. The evaluation carried out shows the performance of the three solvers on the reachability problem for large automata, as for instance, the one obtained from the analysis of the DAW-net discovered from the BPI2011 event log.

9.3 Threats to Validity

The main threats affecting the validity of the reported results are *external* validity issues, which hamper the generalization of the findings. Indeed, we evaluated the solvers on a single real-life log and tested them only on a small number of log traces. Nevertheless, we tried to mitigate the above threat, by carrying on a systematic evaluation on the synthetic datasets and randomly selecting the subset of traces in the real-life event log used in the evaluation. A second threat to the external validity of the results relates to the choice of the data and of the guards. However, also in this case, we tried to reduce the threat by considering different data attributes and values in the synthetic evaluation, and relying on data and guards discovered by the ProM DATA-FLOW DISCOVERY plugin. Finally, we used the solvers as informed users, i.e., we did not apply to them special optimizations. Dedicated encodings and tuning strategies could improve the performance of the solvers. However, the threat is mitigated by the fact that none of the three solvers has been optimized.

10. Related Work

The problem of completing incomplete traces has been investigated in a number of works of trace alignment in the field of process mining, where it still represents one of the challenges (van der Aalst & et al., 2012). Several works have addressed the problem of aligning event logs and procedural models, without (Adriansyah, van Dongen, & van der Aalst, 2011) and with (de Leoni, van der Aalst, & van Dongen, 2012; de Leoni & van der Aalst, 2013; Mannhardt, de Leoni, Reijers, & van der Aalst, 2016) data. All these works, however, explore the search space of possible moves in order to find the best one aligning the log and the model. Differently from them, in this work (i) we assume that the model is correct and

we focus on the repair of incomplete execution traces; (ii) we want to exploit state-of-the-art formal verification techniques to solve the problem as a reachability problem on control and data flow rather than solving an optimisation problem.

The specific problem of reconstructing flows of activities of a model given a partial set of information on it can be related to several fields of research in which the dynamics of a system are perceived only to a limited extent and hence it is needed to reconstruct missing information. Most of the approaches facing this problem share the common conceptual view that a model is taken as a reference to construct possible model-compliant “worlds”, out of the set of observations that convey limited data. When focusing on work in the field of business processes can overall divide the existing proposals in two groups: quantitative (i.e., relying on the availability of a probabilistic model of execution and knowledge) and qualitative (i.e., relying on equally likely “alternative worlds”) approaches. The former, typically based on Partially Observable Markov Decision Models, rely on the availability of a probabilistic model of execution and knowledge. For example, in (Rogge-Solti et al., 2013), the authors exploit stochastic PNs and Bayesian Networks to recover missing information (activities and their durations) of process execution traces. The latter stand on the idea of describing “possible outcomes” regardless of likelihood; hence, knowledge about the world will consist of equally likely “alternative worlds” given the available observations in time, as in this work. For example, in (Bertoli, Di Francescomarino, Dragoni, & Ghidini, 2013) the same issue of reconstructing missing information has been tackled by reformulating it in terms of a Satisfiability (SAT) problem rather than as a planning problem. In (Chesani et al., 2018), the authors propose a framework based on abduction, able to provide key reasoning services, including the completion of incomplete traces when dealing with workflows enriched with temporal workflows, i.e., workflows enriched with time conditions (Kumar, Sabbella, & Barton, 2015). The framework offers different key reasoning services (consistency, compliance, runtime monitoring, prediction) for the analysis of process executions and process models, that can be easily customized to different types of knowledge enforced in the model (e.g., time constraints, business rules).

When focusing on verification techniques for data-aware processes, a number of theoretical works exist both in the area of data-aware processes and of (variants of) PNs. Unfortunately, when combining processes and data, verification problems suddenly become undecidable (Calvanese, De Giacomo, & Montali, 2013). We can divide this literature in two streams. In the first stream, PNs are enriched, by making tokens able to carry various forms of data and by making transitions aware of such data, such in CPNs (van der Aalst & Stahl, 2011) or in data variants such as (Structured) Data Nets (Badouel, H elou et, & Morvan, 2015; Lazi c, Newcomb, Ouaknine, Roscoe, & Worrell, 2007), ν -PNs (Rosa-Velardo & de Frutos-Escrig, 2011) and Conceptual WF-nets with data (Sidorova et al., 2011). For full CPNs, reachability is undecidable and usually obtained by imposing finiteness of color domains. Data variants instead weaken data-related aspects. Specifically Data Nets and ν -PNs consider data as unary relations, while semistructured data tokens are limited to tree-shaped data structures. Also, for these models coverability is decidable, but reachability is not. The work in (Sidorova et al., 2011) considers data elements (e.g., *Price*) that can be used on transitions’ preconditions. However, these nets do not consider data values (e.g., in the example of Section 2 we would not be aware of the values of the variable *request* that T_4 is enabled to write) but only whether the value is “defined” or “undefined”, thus

limiting the reasoning capabilities that can be provided on top of them. For instance, in the example of Section 2, we would not be able to discriminate between the worker and the student loan for the trace in (3), as we would only be aware that *request* is **defined** after $T4$. The second stream contains proposals that take a different approach: instead of making the control-flow model increasingly data-aware, they consider standard data models (such as relational databases and XML repositories) and make them increasingly “dynamics-aware”. Notable examples are relational transducers (Abiteboul, Vianu, Fordham, & Yesha, 2000), active XML (Abiteboul, Segoufin, & Vianu, 2009), the artifact-centric paradigm (Gerede, Bhattacharya, & Su, 2007; Damaggio, Deutsch, & Vianu, 2011; Bagheri Hariri, Calvanese, De Giacomo, Deutsch, & Montali, 2013), and DCDSs (Bagheri Hariri et al., 2013). Such works differ on the limitations imposed to achieve decidability, but they all lack an intuitive control-flow perspective. A further recent work presents RAW-SYS (De Masellis et al., 2017a), a framework that joins the two streams above by directly combining a control-flow model based on PNs and standard data models (à la DCDS) as first class citizens, in which activities/tasks are expressed in a STRIP-like fashion. Verification in RAW-SYS is shown to be decidable when the number of objects accumulated in the same state is bounded (but still infinitely many values may occur in a run) and first-order quantification is restricted on the active domain of the current state. The definition of DAW-net follows the approach of RAW-SYS in combining control-flow model based on PNs and standard data models (à la DCDS) as first class citizens. However, the current work differentiates from that of RAW-SYS in that we consider a data model with variables only which, on the one hand does not require complex restrictions to achieve decidability of reasoning tasks and on the other it is still expressive enough to address the trace completion problem with events carrying data. We indeed remark that the payload of XES standard for traces is a set of attribute-value pairs, thus making data-aware process model with relational data structures needlessly cumbersome for solving reasoning tasks on concrete logs.

The VERIFAS system (Li, Deutsch, & Vianu, 2017) leverages model checking techniques for the verification of temporal properties of infinite-state transition systems arising from processes that carry and manipulate unbounded data. Although in principle the system can be used to verify reachability properties of DAW-net models, it’s difficult to compare because the representation language for the models is more expressive, and the language for specifying the properties that can be verified can be used for more than reachability. Moreover the software used for their experiments (Li, 2017) cannot be used directly for our purposes and it would require nontrivial modifications in order to be included in our experiments.

Moving to the exploitation of verification techniques in the context of business processes, we can notice that planning techniques have already been used in this context, e.g., for verifying process constraints (Regis, Ricci, Aguirre, & Maibaum, 2012), for accomplishing business process reengineering (Rodríguez-Moreno, Borrajo, Cesta, & Oddi, 2007), for conformance checking (de Leoni & Marrella, 2017) as well as for the construction and adaptation of autonomous process models (da Silva & de Lemos, 2011; Marrella, Russo, & Mecella, 2012).

In (De Giacomo, Maggi, Marrella, & Sardiña, 2016; de Leoni, Lanciano, & Marrella, 2018) automated planning techniques have been applied for aligning execution traces and process models. As in this work, in (Di Francescomarino et al., 2015), planning techniques

have been used for addressing the problem of incomplete execution traces with respect to procedural models. However, differently from the two approaches above, this work uses for the first time planning techniques to target the problem of completing incomplete execution traces with respect to a procedural model that also takes into account data and the value they can assume. In (Giordano & Theseider Dupr, 2018) an action language similar to BC (see Section 6), and its encoding in ASP, is exploited for verification of Business Processes enriched with a lightweight ontology language for describing semantic relations among data objects. As the works above mentioned its objectives are different from the ones presented in this paper where we focus on trace completion and evaluating the scalability of different approaches.

11. Conclusions

The work focuses on the problem of the automatic completion of incomplete traces by leveraging workflow nets enriched with rigorous data-aware constructs. Such a problem is recasted to a problem of reachability on a transition system. Besides tackling the specific problem of trace completion, the work proposed in this paper contributes to advancing the state-of-the-art on the concrete exploitation of formal verification techniques on business processes in two different ways: first it presents a rigorous encoding of a data-aware workflow nets based language into Action Languages, Classical Planning, and Model-Checking, based on a common interpretation in terms of transition systems; second it provides a first comprehensive assessment of the performances of different solvers, one for each language, in terms of reasoning with data-aware workflow net languages on both synthetic and real-life data. The evaluation shows that NUXMV is the most robust approach as it is able to deal reasonably well with diverse types of event logs with data. On the other hand, FAST-DOWNWARD, outperforms its competitors on medium-size search spaces.

In the future, we would like to explore and evaluate the effectiveness of the proposed encodings and solvers also for the variants of the reachability problem aimed at addressing other key reasoning services, such as workflow net property verification. Furthermore we aim at investigating the impact of dedicated encodings and tuning strategies that improve the performance of the solvers so as to further support the usage of these techniques in practice.

References

- 3TU Data Center (2011). BPI Challenge 2011 Event Log.. doi:10.4121/uuid:d9769f3d-0ab0-4fb8-803b-0d1120ffc54.
- Abiteboul, S., Segoufin, L., & Vianu, V. (2009). Modeling and verifying Active XML artifacts.. *32(3)*, 10–15.
- Abiteboul, S., Vianu, V., Fordham, B. S., & Yesha, Y. (2000). Relational transducers for electronic commerce. *J. Comput. Syst. Sci.*, *61(2)*, 236–269.
- Adriansyah, A., van Dongen, B. F., & van der Aalst, W. (2011). Conformance checking using cost-based fitness analysis. In *Proc. of EDOC*, pp. 55–64.

- Badouel, E., Hélouët, L., & Morvan, C. (2015). Petri nets with semi-structured data. In *Proc. of 36th International Conference on Application and Theory of Petri Nets and Concurrency*.
- Bagheri Hariri, B., Calvanese, D., De Giacomo, G., Deutsch, A., & Montali, M. (2013). Verification of relational data-centric dynamic systems with external services.. pp. 163–174.
- Baier, C., & Katoen, J. (2008). *Principles of model checking*. MIT Press.
- Bertoli, P., Di Francescomarino, C., Dragoni, M., & Ghidini, C. (2013). Reasoning-based techniques for dealing with incomplete business process execution traces. In *AI*IA*, Vol. 8249 of *LNCS*, pp. 469–480. Springer.
- Calvanese, D., De Giacomo, G., & Montali, M. (2013). Foundations of data-aware process analysis: A database theory perspective.. pp. 1–12.
- Chesani, F., Mello, P., De Masellis, R., Di Francescomarino, C., Ghidini, C., Montali, M., & Tessaris, S. (2018). Compliance in business processes with incomplete information and time constraints: a general framework based on abductive reasoning. *Fundam. Inform.*, 161(1-2), 75–111.
- da Silva, C. E., & de Lemos, R. (2011). A framework for automatic generation of processes for self-adaptive software systems. *Informatica (Slov.)*, 35(1), 3–13.
- Damaggio, E., Deutsch, A., & Vianu, V. (2011). Artifact systems with data dependencies and arithmetic.. pp. 66–77.
- De Giacomo, G., De Masellis, R., & Montali, M. (2014). Reasoning on LTL on finite traces: Insensitivity to infiniteness. In *Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence, July 27 -31, 2014, Québec City, Québec, Canada.*, pp. 1027–1033.
- De Giacomo, G., Maggi, F. M., Marrella, A., & Sardiña, S. (2016). Computing trace alignment against declarative process models through planning. In *ICAPS*, pp. 367–375.
- de Leoni, M., Lanciano, G., & Marrella, A. (2018). Aligning partially-ordered process-execution traces and models using automated planning. In *Proceedings of the Twenty-Eighth International Conference on Automated Planning and Scheduling, ICAPS 2018, Delft, The Netherlands, June 24-29, 2018.*, pp. 321–329.
- de Leoni, M., & Marrella, A. (2017). Aligning real process executions and prescriptive process models through automated planning. *Expert Syst. Appl.*, 82, 162–183.
- de Leoni, M., & van der Aalst, W. (2013). Data-aware Process Mining: Discovering Decisions in Processes Using Alignments. In *Proc of ACM SAC*, pp. 1454–1461.
- de Leoni, M., van der Aalst, W., & van Dongen, B. F. (2012). Data- and resource-aware conformance checking of business processes. In *BIS*, Vol. 117 of *LNBIP*, pp. 48–59.
- de Leoni, M., & van der Aalst, W. M. P. (2013). Data-aware process mining: discovering decisions in processes using alignments. In *Proceedings of the 28th Annual ACM Symposium on Applied Computing, SAC '13, Coimbra, Portugal, March 18-22, 2013*, pp. 1454–1461.

- De Masellis, R., Francescomarino, C. D., Ghidini, C., Montali, M., & Tessaris, S. (2017a). Add data into business process verification: Bridging the gap between theory and practice. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence, February 4-9, 2017, San Francisco, California, USA.*, pp. 1091–1099.
- De Masellis, R., Francescomarino, C. D., Ghidini, C., & Tessaris, S. (2017b). Enhancing workflow-nets with data for trace completion. In Teniente, E., & Weidlich, M. (Eds.), *Business Process Management Workshops - BPM 2017 International Workshops, Barcelona, Spain, September 10-11, 2017, Revised Papers*, Vol. 308 of *Lecture Notes in Business Information Processing*, pp. 89–106. Springer.
- Di Francescomarino, C., Ghidini, C., Tessaris, S., & Sandoval, I. V. (2015). Completing workflow traces using action languages. In *CAiSE*, Vol. 9097 of *LNCS*, pp. 314–330. Springer.
- Eiter, T., Faber, W., Leone, N., Pfeifer, G., & Polleres, A. (2003). A logic programming approach to knowledge-state planning, II: The DLVK system. *Art. Intell.*, 144(1–2), 157–211.
- Gebser, M., Grote, T., & Schaub, T. (2010). Coala: A Compiler from Action Languages to ASP. In Hutchison, D., Kanade, T., Kittler, J., Kleinberg, J. M., Mattern, F., Mitchell, J. C., Naor, M., Nierstrasz, O., Pandu Rangan, C., Steffen, B., Sudan, M., Terzopoulos, D., Tygar, D., Vardi, M. Y., Weikum, G., Janhunen, T., & Niemelä, I. (Eds.), *Logics in Artificial Intelligence*, Vol. 6341, pp. 360–364. Springer Berlin Heidelberg, Berlin, Heidelberg.
- Gerede, C. E., Bhattacharya, K., & Su, J. (2007). Static analysis of business artifact-centric operational models.. In *SOCA*, pp. 133–140. IEEE Computer Society.
- Ghallab, M., Knoblock, C., Wilkins, D., Barrett, A., Christianson, D., Friedman, M., Kwok, C., Golden, K., Penberthy, S., Smith, D., Sun, Y., & Weld, D. (1998). PDDL - The Planning Domain Definition Language..
- Ghallab, M., Nau, D., & Traverso, P. (2016). *Automated Planning and Acting* (1st edition). Cambridge University Press, New York, NY, USA.
- Giordano, L., & Theseider Dupr, D. (2018). Enriched Modeling and Reasoning on Business Processes with Ontologies and Answer Set Programming. In Weske, M., Montali, M., Weber, I., & vom Brocke, J. (Eds.), *Business Process Management Forum*, Lecture Notes in Business Information Processing, pp. 71–88. Springer International Publishing.
- IEEE Task Force on Process Mining (2015). XES Standard Definition. <http://www.xes-standard.org/>.
- Kiepuszewski, B., ter Hofstede, A. H. M., & Bussler, C. J. (2013). On structured workflow modelling. In *Seminal Contributions to Information Systems Engineering*.
- Kumar, A., Sabbella, S. R., & Barton, R. R. (2015). Managing controlled violation of temporal process constraints. In Motahari-Nezhad, H. R., Recker, J., & Weidlich, M. (Eds.), *Business Process Management*, pp. 280–296, Cham. Springer International Publishing.

- Lazić, R., Newcomb, T., Ouaknine, J., Roscoe, A. W., & Worrell, J. (2007). Nets with Tokens Which Carry Data. In *Proceedings of the 28th International Conference on Applications and Theory of Petri Nets and Other Models of Concurrency, (ICATPN 2007)*, LNCS, pp. 301–320. Springer.
- Lee, J., Lifschitz, V., & Yang, F. (2013). Action Language BC: Preliminary Report. In *Proceedings of the Twenty-Third International Joint Conference on Artificial Intelligence, IJCAI '13*, pp. 983–989, Beijing, China. AAAI Press.
- Li, Y. (2017). Has-verifier: Implementation of Verifiers for Hierarchical Artifact Systems..
- Li, Y., Deutsch, A., & Vianu, V. (2017). VERIFAS: A Practical Verifier for Artifact Systems. *Proc. VLDB Endow.*, 11(3), 283–296.
- Lifschitz, V. (1999). Action languages, answer sets and planning. In *The Logic Programming Paradigm: a 25-Year Perspective*, pp. 357–373. Springer.
- Lifschitz, V., & Turner, H. (1994). Splitting a logic program. In *Proceedings of the Eleventh International Conference on Logic Programming*, pp. 23–37, Cambridge, MA, USA. MIT Press.
- Mannhardt, F., de Leoni, M., Reijers, H. A., & van der Aalst, W. M. P. (2016). Balanced multi-perspective checking of process conformance. *Computing*, 98(4), 407–437.
- Marrella, A., Russo, A., & Mecella, M. (2012). Planlets: Automatically recovering dynamic processes in YAWL. In *OTM Conferences (1)*, pp. 268–286.
- McMillan, K. L. (1993). *Symbolic model checking*. Kluwer.
- Object Management Group (OMG) (2011). Business process model and notation (BPMN) version 2.0. Standard.
- Petri, C. A. (1962). *Kommunikation mit Automaten*. Ph.D. thesis, Institut für instrumentelle Mathematik, Bonn.
- Regis, G., Ricci, N., Aguirre, N., & Maibaum, T. S. E. (2012). Specifying and verifying declarative fluent temporal logic properties of workflows. In *Proc. of SBMF*, pp. 147–162.
- Rodríguez-Moreno, M. D., Borrajo, D., Cesta, A., & Oddi, A. (2007). Integrating planning and scheduling in workflow domains. *Expert Systems with Applications*, 33(2), 389–406.
- Rogge-Solti, A., Ronny, S., van der Aalst, W., & Weske, M. (2013). Improving documentation by repairing event logs. In *The Practice of Enterprise Modeling*, Vol. 165 of *LNBIP*, pp. 129–144. Springer.
- Rosa-Velardo, F., & de Frutos-Escrig, D. (2011). Decidability and complexity of petri nets with unordered data. *Theoretical Computer Science*, 412(34), 4439 – 4451.
- Sidorova, N., Stahl, C., & Trčka, N. (2011). Soundness verification for conceptual workflow nets with data: Early detection of errors with the most precision possible. *Inf. Syst.*, 36(7), 1026–1043.
- van der Aalst, W. (1998). The application of petri nets to workflow management. *J. of Circuits, Sys. and Comp.*, 08, 21–66.

- van der Aalst, W., & et al. (2012). Process mining manifesto. In *BPM Workshops*. Springer.
- van der Aalst, W., Hee, K. v., Hofstede, A. t., Sidorova, N., Verbeek, H., Voorhoeve, M., & Wynn, M. (2010). Soundness of workflow nets: classification, decidability, and analysis. *Formal Aspects of Comp.*, *23*(3), 333–363.
- van der Aalst, W., Pesic, M., & Schonenberg, H. (2009). Declarative workflows: Balancing between flexibility and support. *Computer Science - Research and Development*, *23*, 99–113.
- van der Aalst, W., & Stahl, C. (2011). *Modeling Business Processes: A Petri Net-Oriented Approach*. Cooperative information systems. MIT Press.
- van der Aalst, W., & ter Hofstede, A. (2005). Yawl: Yet another workflow language. *Inf. Syst.*, *30*(4), 245–275.
- van der Aalst, W. M. P. (1997). Verification of workflow nets. In *Proc. of ICATPN*, pp. 407–426.

Appendix A. Technical details of Section 4

To relate cases from W^τ to the original workflow W we introduce a “projection” function Π_τ that maps elements from cases of the enriched workflow to cases using only elements from the original workflow. To simplify the notation we will use the same name to indicate mappings from states, firings and cases.

Definition 28. *Let $W = \langle \mathcal{D}, \mathcal{N} = \langle P, T, F \rangle, wr, gd \rangle$ be a DAW-net, $\tau = (e_1, \dots, e_n)$ – where $e_i = \langle t_i, w_i, w_i^d \rangle$ a trace of W , and $W^\tau = \langle \mathcal{D}, \mathcal{N}^\tau = \langle P^\tau, T^\tau, F^\tau \rangle, wr^\tau, gd^\tau \rangle$ the corresponding trace workflow. The mapping Π_τ is defined as following:*

1. let M' be a marking of W^τ , then

$$\Pi_\tau(M') = M' \cap P \times \mathbb{N}$$

is a marking of W ;

2. let (M', η') be a state of W^τ , then

$$\Pi_\tau((M', \eta')) = (\Pi_\tau(M'), \eta')$$

is a state of W ;

3. let t be a transition in T^τ , then

$$\Pi_\tau(t) = \begin{cases} t_i & \text{for } t = t_{e_i} \\ t & \text{for } t \in T \end{cases}$$

is a transition of W ;

4. let $(M, \eta) \xrightarrow{t} (M', \eta')$ be a firing in W^τ , then

$$\Pi_\tau((M, \eta) \xrightarrow{t} (M', \eta')) = \Pi_\tau((M, \eta)) \xrightarrow{\Pi_\tau(t)} \Pi_\tau((M', \eta'))$$

is a firing in W ;

5. let $C = f_0, \dots, f_k$ be a sequence of valid firings in W^τ , then

$$\Pi_\tau(C) = \Pi_\tau(f_0), \dots, \Pi_\tau(f_k)$$

is a sequence of valid firings in W .

In the following we consider a DAW-net $W = \langle \mathcal{D}, \mathcal{N} = \langle P, T, F \rangle, wr, gd \rangle$ and a trace $\tau = (e_1, \dots, e_n)$ of W – where $e_i = \langle t_i, w_i, w_i^d \rangle$. Let $W^\tau = \langle \mathcal{D}, \mathcal{N}^\tau = \langle P^\tau, T^\tau, F^\tau \rangle, wr^\tau, gd^\tau \rangle$ be the corresponding trace workflow. To simplify the notation, in the following we will use t_{e_0} as a synonymous for $start_t$ and $t_{e_{n+1}}$ as end_t ; as if they were part of the trace.

Lemma 9. *If C is a sequence of valid firings in W^τ then $\Pi_\tau(C)$ is a sequence of valid firings in W .*

Proof. Let $C = (M_0, \eta_0) \xrightarrow{t_1} (M_1, \eta_1) \dots (M_{k-1}, \eta_{k-1}) \xrightarrow{t_k} (M_k, \eta_k)$, to show that $\Pi_\tau(C)$ is a case of W we need to prove that (i) $\Pi_\tau((M_0, \eta_0))$ is an initial state of W ; and (ii) the firing $\Pi_\tau((M_{i-1}, \eta_{i-1}) \xrightarrow{t_i} (M_i, \eta_i))$ is valid w.r.t. W for all $1 \leq i \leq n$.

- i) By definition $\Pi_\tau((M_0, \eta_0)) = (\Pi_\tau(M_0), \eta')$ and $\Pi_\tau(M_0) \subseteq M_0$. Since the start place is in P , then start is the only place with a token in $\Pi_\tau(M_0)$.
- ii) Let consider an arbitrary firing $f_i = (M_{i-1}, \eta_{i-1}) \xrightarrow{t_i} (M_i, \eta_i)$ in C (valid by definition), then $\Pi_\tau(f_i) = (\Pi_\tau(M_{i-1}), \eta_{i-1}) \xrightarrow{\Pi_\tau(t_i)} (\Pi_\tau(M_i), \eta_i)$.

Note that – by construction – $\mathbf{gd}(t_i)$ is a restriction of $\mathbf{gd}(\Pi_\tau(t_i))$, $\Pi_\tau(t_i)^\bullet = t_i^\bullet \cap P$, ${}^\bullet\Pi_\tau(t_i) = {}^\bullet t_i \cap P$, $\mathbf{dom}(\mathbf{wr}(t_i)) = \mathbf{dom}(\mathbf{wr}(\Pi_\tau(t_i)))$ and $\mathbf{wr}(t_i)(v) \subseteq \mathbf{wr}(\Pi_\tau(t_i))(v)$; therefore

- $\{p \in P^\tau \mid M_{i-1} > 0\} \cap P = \{p \in P \mid \Pi_\tau(M_{i-1}) > 0\} \supseteq {}^\bullet\Pi_\tau(t_i)$ because $\{p \in P^\tau \mid M_{i-1} > 0\} \supseteq {}^\bullet t_i$;
- $\mathcal{D}, \eta \models \mathbf{gd}(\Pi_\tau(t_i))$ because $\mathcal{D}, \eta \models \mathbf{gd}(t_i)$
- for all $p \in P$ $\Pi_\tau(M_j)(p) = M_j(p)$, therefore:

$$M_i(p) = \Pi_\tau(M_i)(p) = \begin{cases} M_{i-1}(p) - 1 = \Pi_\tau(M_{i-1})(p) - 1 & \text{if } p \in {}^\bullet\Pi_\tau(t_i) \setminus \Pi_\tau(t_i)^\bullet \\ M_{i-1}(p) + 1 = \Pi_\tau(M_{i-1})(p) + 1 & \text{if } p \in \Pi_\tau(t_i)^\bullet \setminus {}^\bullet\Pi_\tau(t_i) \\ M_{i-1}(p) = \Pi_\tau(M_{i-1})(p) & \text{otherwise} \end{cases}$$

because f_i is valid w.r.t. W^τ ;

- the assignment η_i satisfies the properties that its domain is

$$\mathbf{dom}(\eta_i) = \mathbf{dom}(\eta_{i-1}) \cup \{v \mid \mathbf{wr}(\Pi_\tau(t_i))(v) \neq \emptyset\} \setminus \{v \mid \mathbf{wr}(\Pi_\tau(t_i))(v) = \emptyset\}$$

and for each $v \in \mathbf{dom}(\eta_i)$:

$$\eta_i(v) = \begin{cases} d \in \mathbf{wr}(t_i)(v) \subseteq \mathbf{wr}(\Pi_\tau(t_i))(v) & \text{if } v \in \mathbf{dom}(\mathbf{wr}(t_i)) = \mathbf{dom}(\mathbf{wr}(\Pi_\tau(t_i))) \\ \eta_{i-1}(v) & \text{otherwise.} \end{cases}$$

because f_i is valid. □

Before going into details, we will consider some properties of the “trace” workflow.

Lemma 10. *Let $W = \langle \mathcal{D}, \mathcal{N} = \langle P, T, F \rangle, \mathbf{wr}, \mathbf{gd} \rangle$ be a DAW-net and $\tau = (e_1, \dots, e_n)$ – where $e_i = \langle t_i, w_i, w_i^d \rangle$ – a trace of W . If $C = (M_0, \eta_0) \xrightarrow{t_1} (M_1, \eta_1) \dots (M_{k-1}, \eta_{k-1}) \xrightarrow{t_k} (M_k, \eta_k)$ is a sequence of valid firings in W^τ then for all $0 \leq i \leq k$:*

$$\sum_{p \in P^\tau \setminus P} M_i(p) \leq M_0(\mathit{start})$$

Proof. By induction on the length of C .

- For $k = 1$ then the only executable transition is start_t , therefore $t_1 = \mathit{start}_t$ which – by assumption – has two output places and – by construction – $\mathit{start}_t^\bullet \setminus P = \{p_{e_0}\}$. Since the firing is valid, then $M_1(p_{e_0}) = M_0(p_{e_0}) + 1 = 1 \leq M_0(\mathit{start})$.

- Let's assume that the property is true for a sequence C of length n and consider $C' = C(M_n, \eta_n) \xrightarrow{t_{n+1}} (M_{n+1}, \eta_{n+1})$. By construction, each $p \in P^\tau \setminus P$ has a single incoming edge and $\{t \in T^\tau \mid e_i \in t^\bullet\} = \{t_{e_i}\}$ and $\{t \in T^\tau \mid e_i \in \bullet t\} = \{t_{e_{i+1}}\}$. Therefore the only occurrence in which a $p_{e_i} \in P^\tau \setminus P$ can *increase* its value is when $t_{n+1} = t_{e_i}$. Since the transition is valid, then $M_{n+1}(p_{e_i}) = M_n(p_{e_i}) + 1$ and $M_{n+1}(p_{e_{i-1}}) = M_n(p_{e_{i-1}}) - 1$; therefore $\sum_{p \in P^\tau \setminus P} M_i(p) = \sum_{p \in P^\tau \setminus P} M_{i-1}(p) \leq M_0(\text{start}) - 1$ by the inductive hypothesis. □

Lemma 11. *Let $W = \langle \mathcal{D}, \mathcal{N} = \langle P, T, F \rangle, wr, gd \rangle$ be a DAW-net and $\tau = (e_1, \dots, e_n)$ – where $e_i = \langle t_i, w_i, w_i^d \rangle$ – a trace of W , $C = (M_0, \eta_0) \xrightarrow{t_1} (M_1, \eta_1) \dots (M_{k-1}, \eta_{k-1}) \xrightarrow{t_k} (M_k, \eta_k)$ a sequence of valid firings in W^τ , and t_{e_i} is a transition of a firing f_m in C with $1 \leq i \leq n$, then (i) $t_{e_{i-1}}$ is in a transition of a firing in C that precedes f_m , (ii) and if $M_0(\text{start}) = 1$ then there is a single occurrence of t_{e_i} in C .*

Proof. The proof for the first part follows from the structure of the workflow net; because – by construction – each $p \in P^\tau \setminus P$ has a single incoming edge and $\{t \in T^\tau \mid e_i \in t^\bullet\} = \{t_{e_i}\}$ and $\{t \in T^\tau \mid e_i \in \bullet t\} = \{t_{e_{i+1}}\}$. Since each firing must be valid – if $f_m = (M_{m-1}, \eta_{m-1}) \xrightarrow{t_{e_i}} (M_m, \eta_m)$ is in C , then $M_{m-1}(p_{e_{i-1}}) \geq 1$ and this can only be true if there is a firing $f_r = (M_{r-1}, \eta_{r-1}) \xrightarrow{t_{e_{i-1}}} (M_r, \eta_r)$ in C s.t. $r < m$.

To prove the second part is enough to show that for each $1 \leq i \leq n$, if t_{e_i} appears more than once in C then there must be multiple occurrences of $t_{e_{i-1}}$ as well. In fact, if this is the fact, then we can use the previous part to show that there must be multiple occurrences of $t_{e_0} = \text{start}$, and this is only possible if $M_0(\text{start}) > 1$.

By contradiction let's assume that there are two firings f_m and f'_m , with $m < m'$, with the same transition t_{e_i} , but there is only a single occurrence of $t_{e_{i-1}}$ in a firing f_r . Using the previous part of this lemma we conclude that $r < m < m'$, therefore $M_{m-1}(p_{e_{i-1}}) = 1$ because a token could be transferred into $p_{e_{i-1}}$ only by $t_{e_{i-1}}$, so $M_m(p_{e_{i-1}}) = 0$. In the firings between m and m' there are no occurrences of $t_{e_{i-1}}$, so $M_{m'-1}(p_{e_{i-1}}) = M_m(p_{e_{i-1}}) = 0$ which is in contradiction with the assumption that f'_m is a valid firing. □

Now we're ready to show that the “trace” workflow characterises all and only the cases compliant wrt the given trace. We divide the proof into correctness and completeness.

Lemma 12 (Correctness). *Let $W = \langle \mathcal{D}, \mathcal{N} = \langle P, T, F \rangle, wr, gd \rangle$ be a DAW-net, $\tau = (e_1, \dots, e_n)$ – where $e_i = \langle t_i, w_i, w_i^d \rangle$ – a trace of W , and $C = (M_0, \eta_0) \xrightarrow{t_1} (M_1, \eta_1) \dots (M_{k-1}, \eta_{k-1}) \xrightarrow{t_k} (M_k, \eta_k)$ be a sequence of valid firings in W^τ s.t. $M_0(\text{start}) = 1$; given $\ell = \max(\{i \mid t_{e_i} \text{ is in a firing of } C\} \cup \{0\})$, then the trace $\tau' = (e_1, \dots, e_\ell)$ if $\ell > 0$ or $\tau' = \emptyset$ if $\ell = 0$ is compatible with the sequence $\Pi_\tau(C)$ in W .*

Proof. By induction on the length of C .

- If $C = (M_0, \eta_0) \xrightarrow{t_1} (M_1, \eta_1)$ then $t_1 = \text{start}_t$ because the firing is valid and the only place with a token in M_0 is start ; therefore $\ell = 0$ and τ' is the empty trace, and C trivially satisfy the empty trace.

- Let $C = (M_0, \eta_0) \xrightarrow{t_1} (M_1, \eta_1) \dots (M_{k-1}, \eta_{k-1}) \xrightarrow{t_k} (M_k, \eta_k)$ s.t. $\Pi_\tau(C)$ is compliant with τ' . Let's consider $C' = C \cdot (M_k, \eta_k) \xrightarrow{t_{k+1}} (M_{k+1}, \eta_{k+1})$: either $t_{k+1} \in T^\tau \setminus T$ or $t_{k+1} \in T$. In the first case $t_{k+1} = t_{e_\ell}$ for some $1 \leq \ell \leq n$, and – by using Lemma 11 – in C there are occurrences of all the t_{e_i} for $1 \leq i < \ell$ and it's the only occurrence of t_{e_ℓ} . This means that $\ell = \max(\{i \mid t_{e_i} \text{ is in a firing of } C\} \cup \{0\})$ and we can extend γ to γ' by adding the mapping from ℓ to $k+1$. The mapping is well defined because of the single occurrence of t_{e_ℓ} . By definition of t_{e_ℓ} , $(M_k, \eta_k) \xrightarrow{t_{k+1}} (M_{k+1}, \eta_{k+1})$ is compliant with e_ℓ because the additional conditions in $\text{gd}^{\tau'}(t_{k+1})$ guarantee the proper assignments for variables that are not assigned by the transition (Def. 15). Moreover the mapping Π_τ preserve the assignments, therefore $\Pi_\tau(M_k, \eta_k) \xrightarrow{t_{k+1}} (M_{k+1}, \eta_{k+1})$ is compliant with e_ℓ as well. By using the inductive hypnoses we can show that C' is compliant as well. In the second case the mapping is not modified, therefore the inductive hypothesis can be used to provide evidence of the first two conditions for trace compliance of Definition 14. □

Lemma 13 (Completeness). *Let $W = \langle \mathcal{D}, \mathcal{N} = \langle P, T, F \rangle, wr, gd \rangle$ be a DAW-net, $\tau = (e_1, \dots, e_n)$ – where $e_i = \langle t_i, w_i, w_i^d \rangle$ – a trace of W , and $C = (M_0, \eta_0) \xrightarrow{t_1} (M_1, \eta_1) \dots (M_{k-1}, \eta_{k-1}) \xrightarrow{t_k} (M_k, \eta_k)$ be a sequence of valid firings in W such that τ is compliant with it, then there is a sequence of valid firings C' in W^τ s.t. $\Pi_\tau(C') = C$.*

Proof. Since C is compliant with τ , then there is a mapping γ satisfying the conditions of Definition 14. Let $C' = (M'_0, \eta_0) \xrightarrow{t'_1} (M'_1, \eta_1) \dots (M'_{k-1}, \eta_{k-1}) \xrightarrow{t'_k} (M'_k, \eta_k)$ a sequence of firing of W^τ defined as following:

- $M'_0 = M_0 \cup \{(p_{e_i}, 0) \mid 0 \leq i \leq n\}$
- $t'_1 = t_1$ and $M'_1 = M_1 \cup \{(p_{e_j}, 0) \mid 1 \leq j \leq n\} \cup \{(p_{e_0}, 1)\}$
- for each $(M'_{i-1}, \eta_{i-1}) \xrightarrow{t'_i} (M'_i, \eta_i)$, $2 \leq i \leq n$:
 - if there is ℓ s.t. $\gamma(\ell) = i$ then $t'_i = t_{e_\ell}$ and

$$M'_i = M_i \cup \{(p_{e_j}, 0) \mid 0 \leq j \leq n, j \neq \ell\} \cup \{(p_{e_\ell}, 1)\}$$

- otherwise $t'_i = t_i$ and

$$M'_i = M_i \cup (M'_{i-1} \cap (P^\tau \setminus P) \times \mathbb{N})$$

It's not difficult to realise that by construction $\Pi_\tau(C') = C$.

To conclude the proof we need to show that C' is a sequence of valid firing in W^τ . Clearly (M'_0, η_0) is a starting state, so we need to show that all the firings are valid. The conditions involving variables – guards and update of the assignment – follows from the fact that the original firings are valid and the newly introduced transitions are restricted according to the trace data.

Conditions on input and output places that are both in W and W^τ are satisfied because of the validity of the original firing and the additional conditions for the guards of newly introduced transitions are satisfied because of τ is compliant with C . The newly introduced places satisfy the conditions because of the compliance wrt the trace, which guarantees that for each firing with transition t_{e_ℓ} there is the preceding firing with transition $t_{e_{\ell-1}}$ that put a token in the $p_{e_{\ell-1}}$ place. \square

Theorem 1. *Let W be a DAW-net and $\tau = (e_1, \dots, e_n)$ a trace; then W^τ characterises all and only the cases of W compatible with τ . That is*

- \Rightarrow if C is a case of W^τ then τ is compliant with the case of W $\Pi_\tau(C)$; and
- \Leftarrow if τ is compliant with the case of W C then there is a case C' of W^τ s.t. $\Pi_\tau(C') = C$.

Proof.

- \Rightarrow If C is a case of W^τ then it must contain t_{e_n} (the only transition adding a token in the sink), therefore the ℓ of Lemma 12 is n so $\tau' = \tau$ and τ is compliant with $\Pi_\tau(C)$.
- \Leftarrow If τ is compliant with C then by Lemma 13 there is a case C' of W^τ s.t. $\Pi_\tau(C') = C$. Moreover, since the last state of C is a final state, so must be the final state of C' because by construction there cannot be any token in the newly introduced places after the last transition t_{e_n} .

\square

Theorem 2. *Let W be a DAW-net and τ a trace of W . If W is k -safe then W^τ is k -safe as well.*

Proof. We prove the theorem by induction on the length of an arbitrary sequence of valid firings $C = (M_0, \eta_0) \xrightarrow{t_1} (M_1, \eta_1) \dots (M_{k-1}, \eta_{k-1}) \xrightarrow{t_k} (M_k, \eta_k)$ in W^τ . Note that by construction, for any marking M' of W^τ and $p \in P$, $M'(p) = \Pi_\tau(M')(p)$.

- For a case of length 1 the property trivially holds because by definition $M_0(start) \leq k$ and for each $p \in P^\tau$ (different from $start$) $M_0(start) = 0$, and since $(M_0, \eta_0) \xrightarrow{t_1} (M_1, \eta_1)$ is valid the only case in which the number of tokens in a place is increased is for $p \in t_1^\bullet \setminus \bullet t_1$. For any p different from $start$ this becomes $1 \leq k$; while since the $start$ place – by assumption – doesn't have any incoming arc therefore $M_1(start) = M_0(start) - 1 \leq k$.
- For the inductive step we assume that each marking M_0, \dots, M_{m-1} is k -safe. By contradiction we assume that M_m is not k -safe; therefore there is a place $p \in P^\tau$ s.t. $M_m > k$. There are two cases, either $p \in P^\tau \setminus P$ or $p \in P$. In the first case there is a contradiction because, by Lemma 10, $\sum_{p \in P^\tau \setminus P} M_i(p) \leq M_0(start) = k$. In the second case, since $\Pi_\tau(C)$ is a case of W and $\Pi_\tau(M_m)(p) = M_m(p)$, there is a contradiction with the hypothesis that W is k -safe.

\square

Appendix B. Technical details of 6

Lemma 3. *Let W be a DAW-net model, and $\rho = (M_0, \eta_0) \xrightarrow{t_0} (M_1, \eta_1) \xrightarrow{t_1} \dots \xrightarrow{t_{\ell-1}} (M_\ell, \eta_\ell)$ be a sequence of valid firings of W (with $\ell \geq 0$). For every state (M_i, η_i) ($0 \leq i \leq \ell$) and query $\Phi \in \mathcal{L}(\mathcal{D}_W)$*

$$\mathcal{D}_i, \eta_i \models \Phi \text{ iff } \Phi_i(\rho) \models i : \Phi^{\text{BC}}$$

where $i : \Phi^{\text{BC}}$ is the formula where $i : \cdot$ is placed in front of each term t^{BC} , and $\Phi_i(\rho) \models i : t^{\text{BC}}$ iff $i : t^{\text{BC}} \in \Phi_i(\rho)$

Proof. Without loss of generality we can assume that Φ is in the DNF form described in Section 6; i.e. $\bigvee_{i=1}^k t_1^i \wedge \dots \wedge t_{\ell_i}^i$ where each term t_j^i is either of the form $v = o$ or $\neg \text{def}(v)$. In this case to prove the lemma is sufficient to show that for the base cases $\mathcal{D}_i, \eta_i \models t$ iff $i : t^{\text{BC}} \in \Phi_i(\rho)$:

$$(v = o) \quad \mathcal{D}_i, \eta_i \models v = o \text{ iff } \eta_i(v) = o \text{ iff (by Definition 19) } \{i : v = o\} \subseteq \Phi_i'(\rho) \subseteq \Phi_i(\rho);$$

$$(\neg \text{def}(v)) \quad \mathcal{D}_i, \eta_i \models \neg \text{def}(v) \text{ iff } \eta_i(v) \text{ is undefined iff (by Definition 19) } \{i : v = \text{null}\} \subseteq \Phi_i'(\rho) \subseteq \Phi_i(\rho).$$

□

Lemma 4 (Completeness of BC encoding). *Let W be a DAW-net and $\text{BC}(W)$ its encoding, then for any $\ell \geq 0$ if ρ is a sequence of valid firings of W of length ℓ , then $\bigcup_{i=0}^{\ell} \Phi_i(\rho)$ is a stable model of $P_\ell(\text{BC}(W))$.*

Proof. We prove the lemma by induction on the length of the case ρ that $\bigcup_{i=0}^{\ell} \Phi_i(\rho)$ is a stable model of $P_\ell(\text{BC}(W))$. For $\ell = 0$ there is only the initial state and $P_0(\text{BC}(W))$ includes only rules derived from the static laws (none), the **initially** statements:

$$\begin{aligned} 0 : \text{start} &= \text{TRUE} \\ 0 : p &= \text{FALSE} && \text{for } p \in P \text{ and } p \neq \text{start} \\ 0 : v &= \text{null} && \text{for } v \in \mathcal{V}' \\ 0 : \text{trans} &= \text{TRUE} \end{aligned}$$

and the encoding of BC into ASP:

- for all the places and the additional fluent for enforcing actions $p \in P \cup \{\text{trans}\}$

$$\begin{aligned} 0 : p &= d \vee \neg(0 : p = d) && \text{for } d \in \{\text{TRUE}, \text{FALSE}\} \\ \neg 0 : p &= d' \leftarrow 0 : p = d && \text{for } d, d' \in \{\text{TRUE}, \text{FALSE}\} \text{ and } d \neq d' \end{aligned}$$

$$\begin{aligned} \leftarrow 0 : p &= d, \neg 0 : p = d && \text{for } d \in \{\text{TRUE}, \text{FALSE}\} \\ \leftarrow \sim(0 : p &= \text{TRUE}), \sim(0 : p = \text{FALSE}) \end{aligned}$$

- for all the variables $v \in V'$

$$\begin{aligned} 0 : v &= d \vee \neg(0 : v = d) && \text{for } d \in \text{adm}(v) \cup \{\text{null}\} \\ \neg 0 : v &= d' \leftarrow 0 : v = d && \text{for } d, d' \in \text{adm}(v) \cup \{\text{null}\} \text{ and } d \neq d' \end{aligned}$$

$$\begin{aligned}
 \leftarrow 0 : v = d, \neg 0 : v = d & \quad \text{for } d \in \text{adm}(v) \cup \{\text{null}\} \\
 \leftarrow \sim(0 : v = d_1), \dots, \sim(0 : v = d_k), \sim(0 : v = \text{null}) & \quad \text{s.t. } \{d_1, \dots, d_k\} = \text{adm}(v)
 \end{aligned}$$

By the definition of (M_0, η_0)

$$\begin{aligned}
 \Phi_0^{\nu}((M_0, \eta_0)) &= \{0:\text{start}=\text{TRUE}, \neg 0:\text{start}=\text{FALSE}\} \cup \\
 &\quad \{0:p=\text{FALSE}, \neg 0:p=\text{TRUE} \mid p \in P \setminus \{\text{start}\}\} \cup \\
 &\quad \{0:v=\text{null} \mid v \in \mathcal{V}'\} \cup \{\neg 0:v=o \mid v \in \mathcal{V}', o \in \text{adm}(v)\} \cup \\
 &\quad \{0:\text{trans}=\text{TRUE}, \neg 0:\text{trans}=\text{FALSE}\} \\
 \Phi_0^{\tau}((M_0, \eta_0)) &= \emptyset
 \end{aligned}$$

It's not difficult to see that all the above rules are satisfied by $\Phi_0((M_0, \eta_0))$. We need to show also that it's the minimal model of the reduct $P_0(\text{BC}(W))^{\Phi_0((M_0, \eta_0))}$ of $P_0(\text{BC}(W))$ w.r.t. $\Phi_0((M_0, \eta_0))$. To this end we note that in the reduct all the rules including NAF literals are removed because each fluent is assigned to a value, and the remaining constraints

$$\leftarrow 0:f=o, \neg 0:f=o$$

are satisfied by construction because $\{0:f=o, \neg 0:f=o\} \not\subseteq \Phi_0((M_0, \eta_0))$.²³

In $\Phi_0((M_0, \eta_0))$ there are just literals $0:f=c$ and $\neg 0:f=c$. Removing any positive (i.e. $0:f=c$) literal would falsify the the **initially** statements, and removing the negative ones would contradict the corresponding rules:

$$\neg 0:f=o' \leftarrow 0:f=o$$

where $0:f=o \in \Phi_0((M_0, \eta_0))$ and $o' \neq o$. This concludes that $\Phi_0((M_0, \eta_0))$ is a minimal model for $P_0(\text{BC}(W))^{\Phi_0((M_0, \eta_0))}$.

For the inductive step we assume that for an arbitrary case ρ of length ℓ , $\bigcup_{i=0}^{\ell} \Phi_i(\rho)$ is a stable model of $P_{\ell}(\text{BC}(W))$ and we show that $\bigcup_{i=0}^{\ell+1} \Phi_i(\rho \xrightarrow{t_{\ell}} (M_{\ell+1}, \eta_{\ell+1}))$ is a stable model of $P_{\ell+1}(\text{BC}(W))$.

Note that, by construction,

$$\bigcup_{i=0}^{\ell+1} \Phi_i(\rho \xrightarrow{t_{\ell}} (M_{\ell+1}, \eta_{\ell+1})) = \bigcup_{i=0}^{\ell} \Phi_i(\rho) \cup \Phi_{\ell}^{\tau}(\rho \xrightarrow{t_{\ell}} (M_{\ell+1}, \eta_{\ell+1})) \cup \Phi_{\ell+1}^{\nu}(\rho \xrightarrow{t_{\ell}} (M_{\ell+1}, \eta_{\ell+1}))$$

23. Remember that a constraint $\leftarrow \ell_1, \dots, \ell_n$ corresponds to the rule $f \leftarrow \sim f, \ell_1, \dots, \ell_n$.

and $P_{\ell+1}(\text{BC}(W))$ is equal to $P_{\ell}(\text{BC}(W))$ plus the rules:

$\ell:t \vee \neg\ell:t$	for all $t \in T$
$\ell + 1:v=o \leftarrow \ell:v=o, \sim\neg\ell + 1:v=o$	for $v \in \mathcal{V}'$ and $o \in \text{adm}(v) \cup \{\text{null}\}$
$\ell + 1:p=o \leftarrow \ell:p=o, \sim\neg\ell + 1:p=o$	for $p \in P$ and $o \in \{\text{TRUE}, \text{FALSE}\}$
$\ell + 1:v=d \leftarrow \ell:t, \sim\neg\ell + 1:v=d$	for all $t \in T, (v, d) \in \text{wr}(t)$
$\ell + 1:p=\text{FALSE} \leftarrow \ell:t$	for all $t \in T, p \in \bullet t \setminus t^\bullet$
$\ell + 1:p=\text{TRUE} \leftarrow \ell:t$	for all $t \in T, p \in t^\bullet \setminus \bullet t$
$\ell + 1:v=\text{null} \leftarrow \ell:t$	for all $v \in \mathcal{V}', t \in T$, s.t. $\text{wr}(t)(v) = \emptyset$
$\ell + 1:\text{trans}=\text{TRUE} \leftarrow \ell:t$	for all $t \in T$
$\neg\ell + 1:p=d' \leftarrow \ell + 1:p=d$	for $p \in P \cup \{\text{trans}\}, d, d' \in \{\text{TRUE}, \text{FALSE}\}$ and $d \neq d'$
$\neg\ell + 1:v=d' \leftarrow \ell + 1:v=d$	for $v \in \mathcal{V}' d, d' \in \text{adm}(v) \cup \{\text{null}\}$ and $d \neq d'$

the constraints

$\leftarrow \ell:t, \ell:s$	for $t, s \in T$, s.t. $t \neq s$
$\leftarrow \ell:t, \sim\neg\ell + 1:v=d$	for $t \in T, v \in \mathcal{V}'$ s.t. $\text{wr}(t)(v) \neq \emptyset$, $d \in \{\text{null}\} \cup \text{adm}(v) \setminus \text{wr}(t)(v)$
$\leftarrow \ell:t, \ell:p=\text{FALSE}$	for $t \in T, p \in \bullet t$
$\leftarrow \ell + 1:p=d, \ell + 1:p=d$	for $p \in P \cup \{\text{trans}\}, d \in \{\text{TRUE}, \text{FALSE}\}$
$\leftarrow \sim\ell + 1:p=\text{TRUE}, \sim\ell + 1:p=\text{FALSE}$	for $p \in P \cup \{\text{trans}\}$
$\leftarrow \ell + 1:v=d, \ell + 1:v=d$	for $v \in \mathcal{V}', d \in \text{adm}(v) \cup \{\text{null}\}$
$\leftarrow \sim\ell + 1:v=d_1, \dots, \sim\ell + 1:v=d_k, \sim\ell + 1:v=\text{null}$	for $v \in \mathcal{V}', \{d_1, \dots, d_k\} = \text{adm}(v)$

and the constraints corresponding to the guards; that is, for each transition t s.t. $\text{gd}(t) \neq \text{true}$ we consider $\overline{\text{gd}(t)} = \bigvee_{i=1}^k t_1^i \wedge \dots \wedge t_{\ell_i}^i$:

$$\begin{aligned} &\leftarrow \ell:t, \ell:t_1^{1\text{BC}}, \dots, \ell:t_{n_1}^{1\text{BC}} \\ &\dots \\ &\leftarrow \ell:t, \ell:t_1^{k\text{BC}}, \dots, \ell:t_{n_k}^{k\text{BC}} \end{aligned}$$

To show that $\bigcup_{i=0}^{\ell+1} \Phi_i(\rho \xrightarrow{t_\ell} (M_{\ell+1}, \eta_{\ell+1}))$ is a stable model of $P_{\ell+1}(\text{BC}(W))$ we use the *Splitting Sets* technique as introduced in (Lifschitz & Turner, 1994); a generalisation of stratification which enables the splitting of a program in two parts on the basis of atoms in the head of the rules, and provides a way of characterising the stable models in terms of the stable modes of the two parts. More specifically, A splitting set for a program P is any set of atoms U such that, for every rule $r \in P$, if $\text{head}(r) \cap U \neq \emptyset$, then $\text{atoms}(r) \subseteq U$. The set of rules $r \in P$ such that $\text{atoms}(r) \subseteq U$ is the bottom of P relative to U , denoted by $\text{bot}_U(P)$. The set $\text{top}_U(P) = P \setminus \text{bot}_U(P)$ is the top of P relative to U .

Let's consider the set

$$\begin{aligned} U = \{ &i:f=d \mid i \leq \ell, f \text{ fluent constant and } d \text{ constant} \} \cup \\ &\{i:a \mid i < \ell, a \text{ action constant} \} \end{aligned}$$

then U is a splitting set for $P_{\ell+1}(\text{BC}(W))$ and $\text{bot}_U(P_{\ell+1}(\text{BC}(W))) = P_\ell(\text{BC}(W))$; therefore we just need to show that $\Phi_\ell^\tau(\rho \xrightarrow{t_\ell} (M_{\ell+1}, \eta_{\ell+1})) \cup \Phi_{\ell+1}^\nu(\rho \xrightarrow{t_\ell} (M_{\ell+1}, \eta_{\ell+1}))$ is a stable model of $e_U(P_{\ell+1}(\text{BC}(W)) \setminus P_\ell(\text{BC}(W)), \bigcup_{i=0}^\ell \Phi_i(\rho))$ where $e_U(P, X)$ is defined in (Lifschitz & Turner, 1994) as:

$$e_U(P, X) = \{r \mid \text{exists } r' \in P \text{ s.t. } \text{body}^+(r') \cap U \subseteq X, \text{body}^-(r') \cap U \cap X = \emptyset \\ \text{head}(r) = \text{head}(r'), \text{body}^+(r) = \text{body}^+(r') \setminus U, \text{body}^-(r) = \text{body}^-(r') \setminus U\}$$

that is, only rules whose bodies are not falsified by X , and the remaining literal in U are removed.

Following Definition 19:

$$\begin{aligned} \Phi_{\ell+1}^\nu(\rho) = & \{\ell + 1 : p = \text{TRUE}, \neg(\ell + 1 : p = \text{FALSE}) \mid p \in P, M_{\ell+1}(p) > 0\} \cup \\ & \{\ell + 1 : p = \text{FALSE}, \neg(\ell + 1 : p = \text{TRUE}) \mid p \in P, M_{\ell+1}(p) = 0\} \cup \\ & \{\ell + 1 : v = o, \neg(\ell + 1 : v = \text{null}) \mid v \in \mathcal{V}', \eta_{\ell+1}(v) = o\} \cup \\ & \{\ell + 1 : v = \text{null} \mid v \in \mathcal{V}', \eta_{\ell+1}(v) \text{ is undefined}\} \cup \\ & \{\neg(\ell + 1 : v = o) \mid v \in \mathcal{V}', o \in \text{adm}(v), \eta_{\ell+1}(v) \neq o \text{ or } \eta_{\ell+1}(v) \text{ is undefined}\} \cup \\ & \{\ell + 1 : \text{trans} = \text{TRUE}, \neg(\ell + 1 : \text{trans} = \text{FALSE})\} \cup \\ \Phi_\ell^\tau(\rho) = & \{\ell : t_\ell\} \cup \{\neg(\ell : t) \mid t \in T, t \neq t_\ell\} \end{aligned}$$

We can discard the actual constraints, since they are satisfied by construction of $\Phi_\ell^\tau(\rho \xrightarrow{t_\ell} (M_{\ell+1}, \eta_{\ell+1})) \cup \Phi_{\ell+1}^\nu(\rho \xrightarrow{t_\ell} (M_{\ell+1}, \eta_{\ell+1}))$ because we assume that $(M_\ell, \eta_\ell) \xrightarrow{t_\ell} (M_{\ell+1}, \eta_{\ell+1})$ is a valid firing. The constraints referring to guards are satisfied because either $\neg(\ell : t \in \Phi_\ell^\tau(\rho \xrightarrow{t_\ell} (M_{\ell+1}, \eta_{\ell+1})))$ or at least one of the $\neg(\ell : t_i^{j\text{BC}} \in \Phi_\ell^\nu(\rho \xrightarrow{t_\ell} (M_{\ell+1}, \eta_{\ell+1})))$ for each of the constraint, because t_i^j holds in η_ℓ iff $\ell : t_i^{k\text{BC}} \in \Phi_\ell^\nu(\rho)$ (Lemma 3).

We then focus on the rules in $e_U(P_{\ell+1}(\text{BC}(W)) \setminus P_\ell(\text{BC}(W)), \bigcup_{i=0}^\ell \Phi_i(\rho))$ with terms in the head. The positive ones:

$$\begin{array}{ll} \ell : t \vee \neg \ell : t & \text{for all } t \in T \\ \ell + 1 : p = \text{FALSE} \leftarrow \ell : t & \text{for all } t \in T, p \in \bullet t \setminus t^\bullet \\ \ell + 1 : p = \text{TRUE} \leftarrow \ell : t & \text{for all } t \in T, p \in t^\bullet \setminus \bullet t \\ \ell + 1 : v = \text{null} \leftarrow \ell : t & \text{for all } v \in \mathcal{V}', t \in T, \text{s.t. } \text{wr}(t)(v) = \emptyset \\ \ell + 1 : \text{trans} = \text{TRUE} \leftarrow \ell : t & \text{for all } t \in T \\ \neg \ell + 1 : \text{trans} = d' \leftarrow \ell + 1 : \text{trans} = d & \text{for } d, d' \in \{\text{TRUE}, \text{FALSE}\} \text{ and } d \neq d' \\ \neg \ell + 1 : p = d' \leftarrow \ell + 1 : p = d & \text{for } p \in P, d, d' \in \{\text{TRUE}, \text{FALSE}\} \text{ and } d \neq d' \\ \neg \ell + 1 : v = d' \leftarrow \ell + 1 : v = d & \text{for } v \in \mathcal{V}', d, d' \in \text{adm}(v) \cup \{\text{null}\} \text{ and } d \neq d' \end{array}$$

and the ones with weak negation in the body:

$$\begin{array}{ll} \ell + 1 : v = o \leftarrow \sim \neg \ell + 1 : v = o & \text{for } v \in \mathcal{V}' \text{ and } o = \eta_\ell(v) \\ \ell + 1 : v = \text{null} \leftarrow \sim \neg \ell + 1 : v = \text{null} & \text{for } v \in \mathcal{V}' \text{ and } v \notin \eta_\ell \\ \ell + 1 : p = \text{TRUE} \leftarrow \sim \neg \ell + 1 : p = \text{TRUE} & \text{for } p \in P \text{ and } M_\ell(p) > 0 \\ \ell + 1 : p = \text{FALSE} \leftarrow \sim \neg \ell + 1 : p = \text{FALSE} & \text{for } p \in P \text{ and } M_\ell(p) = 0 \\ \ell + 1 : v = d \leftarrow \ell : t, \sim \neg \ell + 1 : v = d & \text{for all } t \in T, (v, d) \in \text{wr}(t) \end{array}$$

whose reduct w.r.t. $\Phi_\ell^\tau(\rho \xrightarrow{t_\ell} (M_{\ell+1}, \eta_{\ell+1})) \cup \Phi_{\ell+1}^\nu(\rho \xrightarrow{t_\ell} (M_{\ell+1}, \eta_{\ell+1}))$ are:

$$\begin{array}{ll}
 \ell + 1:v=o & \text{for } v \in \mathcal{V}' \text{ and } o = \eta_{\ell+1}(v) \\
 \ell + 1:v=\text{null} & \text{for } v \in \mathcal{V}' \text{ and } v \notin \eta_{\ell+1} \\
 \ell + 1:p=\text{TRUE} & \text{for } p \in P \text{ and } M_{\ell+1}(p) > 0 \\
 \ell + 1:p=\text{FALSE} & \text{for } p \in P \text{ and } M_{\ell+1}(p) = 0 \\
 \ell + 1:v=d \leftarrow \ell:t & \text{for all } t \in T, (v, d) \in \mathbf{wr}(t), d = \eta_{\ell+1}(v)
 \end{array}$$

We show that $\Phi_\ell^\tau(\rho \xrightarrow{t_\ell} (M_{\ell+1}, \eta_{\ell+1})) \cup \Phi_{\ell+1}^\nu(\rho \xrightarrow{t_\ell} (M_{\ell+1}, \eta_{\ell+1}))$ is also a minimal model for the reduct. It's easy to see that it satisfy all the rules by construction, so we show minimality.

Removing any term from $\Phi_\ell^\tau(\rho)$ would contradict one of the rules

$$\ell:t \vee \neg\ell:t$$

therefore $\ell:t_\ell$ must be in $\Phi_\ell^\tau(\rho)$ and it's the only positive action constant term. Therefore $\ell + 1:trans=\text{TRUE}$ cannot be removed as well, otherwise the rule

$$\ell + 1:trans=\text{TRUE} \leftarrow \ell:t_\ell$$

would be violated as well.

Positive literals $\ell + 1:v=o$, $\ell + 1:p=o$ in $\Phi_{\ell+1}^\nu(\rho \xrightarrow{t_\ell} (M_{\ell+1}, \eta_{\ell+1}))$ for $v \in \mathcal{V}'$, $p \in P$ cannot be removed because it'd contradict the corresponding rule among the ones:

$$\begin{array}{ll}
 \ell + 1:v=o & \text{for } v \in \mathcal{V}' \text{ and } o = \eta_{\ell+1}(v) \\
 \ell + 1:v=\text{null} & \text{for } v \in \mathcal{V}' \text{ and } v \notin \eta_{\ell+1} \\
 \ell + 1:p=\text{TRUE} & \text{for } p \in P \text{ and } M_{\ell+1}(p) > 0 \\
 \ell + 1:p=\text{FALSE} & \text{for } p \in P \text{ and } M_{\ell+1}(p) = 0
 \end{array}$$

On the other hand, the negated literals must be included otherwise some of the rules

$$\begin{array}{ll}
 \neg\ell + 1:p=d' \leftarrow \ell + 1:p=d & \text{for } p \in P, d, d' \in \{\text{TRUE}, \text{FALSE}\} \text{ and } d \neq d' \\
 \neg\ell + 1:v=d' \leftarrow \ell + 1:v=d & \text{for } v \in \mathcal{V}' \text{ and } d, d' \in \mathbf{adm}(v) \cup \{\text{null}\} \text{ and } d \neq d'
 \end{array}$$

would be unsatisfied. □

Lemma 5. *Let $TS_{\text{BC}(W)} = (\mathcal{A}, S, S_0, \delta)$, then*

1. $S_0 = \{s_0\}$;
2. if $(s, A, s') \in \delta$ then $|A| = 1$.

Proof. To prove the first statement we should note that $P_0(\text{BC}(W))$ includes the following facts derived from the **initially** statements:

$$\begin{array}{ll}
 0 : \text{start} = \text{TRUE} & \\
 0 : p = \text{FALSE} & \text{for } p \in P \text{ and } p \neq \text{start} \\
 0 : v = \text{null} & \text{for } v \in \mathcal{V}' \\
 0 : \text{trans} = \text{TRUE} &
 \end{array}$$

which fixes a value for each fluent, therefore for any pair of stable models M, M' of $P_0(\text{BC}(W))$ $\nu_{\text{BC}}^0(M) = \nu_{\text{BC}}^0(M') = s_0$.

Regarding the second property, we first note that for all $s \in S$, $s(\text{trans}) = \text{TRUE}$ because $P_\ell(\text{BC}(W))$ only the literals $i:\text{trans}=\text{TRUE}$ appear in the head of any rule for all $i \leq \ell$, and the constraint

$$\leftarrow \sim i:\text{trans}=\text{TRUE}, \sim i:\text{trans}=\text{FALSE}$$

forces either $i:\text{trans}=\text{TRUE}$ or $i:\text{trans}=\text{FALSE}$ to be in any stable model. Therefore $i:\text{trans}=\text{TRUE}$ must be in any stable model of $P_\ell(\text{BC}(W))$ for any $i \leq \ell$.

Then we recall that $(s, A, s') \in \delta$ iff there is a stable model $M_{\ell+1}$ of $P_{\ell+1}(\text{BC}(W))$ for some $\ell \geq 0$, s.t. $\nu_{\text{BC}}^\ell(M_{\ell+1}) = s'$, $\nu_{\text{BC}}^\ell(M_{\ell+1}) = s$, and $\{a \in T \mid \ell:a \in M_{\ell+1}\} = A$. First A cannot be empty because the only rules with $\ell+1:\text{trans}=\text{TRUE}$ in the head are

$$\ell+1:\text{trans}=\text{TRUE} \leftarrow \ell:t \quad \text{for all } t \in T$$

therefore there must be at least a $a \in T$ s.t. $\ell:t \in M_{\ell+1}$. Moreover, the dynamic laws

$$\text{false after } t, s \quad \text{for } (t, s) \in T \times T \text{ and } t \neq s$$

correspond to the constraints

$$\leftarrow i : t, i : s \quad \text{for } (t, s) \in T \times T \text{ and } t \neq s$$

which prevent two parallel actions. □

Lemma 6 (Correctness of BC encoding). *Let W be a DAW-net and $\text{BC}(W)$ its encoding, then for any $\ell \geq 0$ if M_ℓ is a stable model of $P_\ell(\text{BC}(W))$, then*

$$\rho = \Phi_0^-(M_\ell) \xrightarrow{\tau_0(M_\ell)} \Phi_1^-(M_\ell) \xrightarrow{\tau_1(M_\ell)} \dots \xrightarrow{\tau_{\ell-1}(M_\ell)} \Phi_\ell^-(M_\ell)$$

is a sequence of valid firings of W of length ℓ .

Proof. We prove the lemma by induction on ℓ . First we consider $P_0(\text{BC}(W))$, which contains the facts

$$\begin{aligned} 0:\text{start}=\text{TRUE} \\ 0:p=\text{FALSE} & \quad \text{for } p \in P \text{ and } p \neq \text{start} \\ 0:v=\text{null} & \quad \text{for } v \in \mathcal{V}' \\ 0:\text{trans}=\text{TRUE} \end{aligned}$$

that, together with Lemma 5, ensures that $\Phi_0^-(M_0) = (M_0, \eta_0)$ satisfies the conditions

$$\begin{aligned} M_0 &= \{(start, 1)\} \cup \{(p, 0) \mid p \in P \setminus \{start\}\} \\ \eta_0 &= \emptyset \end{aligned}$$

corresponding to the initial state (M_s, η_s) .

For the inductive step, let $M_{\ell+1}$ be a stable model for $P_{\ell+1}(\text{BC}(W))$, then we show that $\Phi_\ell^-(M_{\ell+1}) \xrightarrow{\tau_\ell(M_{\ell+1})} \Phi_{\ell+1}^-(M_{\ell+1})$ is a valid firing according to Def. 11. Since Lemma 5 ensures that the mappings are well defined, we need to show that

1. $a = \tau_\ell(M_{\ell+1})$ is enabled in $\Phi_\ell^-(M_{\ell+1})$:

$$\ell:p=\text{TRUE} \in M_{\ell+1} \quad \text{for } p \in \bullet a$$

and that's the case because if there's $p' \in \bullet a$ s.t. $\ell:p=\text{TRUE} \notin M_{\ell+1}$ then $\ell:p=\text{FALSE} \in M_{\ell+1}$ and the constraint

$$\leftarrow \ell:a, \ell:p'=\text{FALSE}$$

deriving from 17 would be contradicted.

2. the guard of a is satisfied in $\Phi_\ell^-(M_{\ell+1})$; i.e. the formula $\bigvee_{i=1}^k t_1^i \wedge \dots \wedge t_{n_i}^i \equiv_{\mathcal{D}'} \neg \text{gd}(t)$ is not satisfied in $\Phi_\ell^-(M_{\ell+1})$:²⁴

$$\begin{aligned} \{\ell:t_1^{1\text{BC}}, \dots, \ell:t_{n_1}^{1\text{BC}}\} &\not\subseteq M_{\ell+1} \\ &\dots \\ \{\ell:t_1^{k\text{BC}}, \dots, \ell:t_{n_k}^{k\text{BC}}\} &\not\subseteq M_{\ell+1} \end{aligned}$$

if for any r $\{\ell:t_1^r, \dots, \ell:t_{n_r}^r\} \subseteq M_{\ell+1}$ then the corresponding constraint

$$\leftarrow \ell:a, \ell:t_r^{1\text{BC}}, \dots, \ell:t_{n_r}^{\text{BC}}$$

would be contradicted.

3. the marking of $\Phi_{\ell+1}^-(M_{\ell+1})$ is updated according to Def. 3:

$$\begin{aligned} \ell + 1:p=\text{FALSE} \in M_{\ell+1} & \quad \text{for } p \in \bullet a \setminus a^\bullet \\ \ell + 1:p=\text{TRUE} \in M_{\ell+1} & \quad \text{for } p \in a^\bullet \setminus \bullet a \\ \ell + 1:p=d \in M_{\ell+1} & \quad \ell:p=d \in M_{\ell+1}, \text{ for } p \in (P \setminus a^\bullet \cup \setminus \bullet a) \cup (a^\bullet \cap \bullet a) \end{aligned}$$

the first two conditions are ensured by the rules generated by laws 11 and 12

$$\begin{aligned} \ell + 1:p=\text{FALSE} \leftarrow \ell:a & \quad \text{for all } p \in \bullet a \setminus a^\bullet \\ \ell + 1:p=\text{TRUE} \leftarrow \ell:a & \quad \text{for all } p \in a^\bullet \setminus \bullet a \end{aligned}$$

while the ‘‘inertial’’ laws 10

$$\ell + 1:p=o \leftarrow \ell:p=o, \sim \neg \ell + 1:p=o \quad \text{for } p \in P \text{ and } o \in \{\text{TRUE}, \text{FALSE}\}$$

ensure the third condition; because those are the only kind of rules where a term like $\ell + 1:p=b$ appears in the head.

4. the assignment η' of $\Phi_{\ell+1}^-(M_{\ell+1})$ is updated according to $\text{wr}(a)$:

$$\begin{aligned} \ell + 1:v=\text{null} \in M_{\ell+1} & \quad \text{for } v \in \mathcal{V}' \text{ s.t. } \text{wr}(a)(v) = \emptyset \\ \ell + 1:v=d \in M_{\ell+1} & \quad \text{for } v \in \mathcal{V}' \text{ s.t. } \text{wr}(a)(v) \neq \emptyset, \text{ and some } d \in \text{wr}(a)(v) \\ \ell + 1:v=d \in M_{\ell+1} & \quad \ell:v=d \in M_{\ell+1}, \text{ for } v \in \mathcal{V}' \setminus \text{dom}(\text{wr}(a)) \end{aligned}$$

24. Each term t_j^i is either of the form $v = o$ or $\neg \text{def}(v)$, and $v = o^{\text{BC}} \mapsto v = o$, $\neg \text{def}(v^{\text{BC}}) \mapsto v = \text{null}$.

The first condition is ensured by the rule

$$\ell + 1:v=\text{null} \leftarrow \ell:a \quad \text{for all } v \in \mathcal{V}', \text{ s.t. } \text{wr}(a)(v) = \emptyset$$

derived from law 14. The second from the constraints

$$\leftarrow \ell:a, \sim \neg \ell + 1:v=d \quad \text{for } v \in \mathcal{V}' \text{ s.t. } \text{wr}(a)(v) \neq \emptyset, d \in \{\text{null}\} \cup \text{adm}(v) \setminus \text{wr}(a)(v)$$

derived from the law 15 since if $\neg \ell + 1:v=o \notin M_{\ell+1}$ then $\ell + 1:v=o \in M_{\ell+1}$, because a value must be associated to any variable and this will cause the inclusion of all the “negated” assignments for the other values different from o :

$$\begin{aligned} \leftarrow \sim \ell + 1:v=d_1, \dots, \sim \ell + 1:v=d_k, \sim \ell + 1:v=\text{null} & \quad \text{for } v \in \mathcal{V}', \{d_1, \dots, d_k\} = \text{adm}(v) \\ \neg \ell + 1:v=d' \leftarrow \ell + 1:v=d & \quad \text{for } v \in \mathcal{V}' \text{ } d, d' \in \text{adm}(v) \cup \{\text{null}\} \\ & \quad \text{and } d \neq d' \end{aligned}$$

The last property derives from the “inertial” constraint

$$\neg \ell + 1:v=d' \leftarrow \ell + 1:v=d \quad \text{for } v \in \mathcal{V}' \text{ } d, d' \in \text{adm}(v) \cup \{\text{null}\} \text{ and } d \neq d'$$

derived from law 9 and the fact that any other rule with terms like $\ell + 1:v=o$ in the head have $\ell:t$ in the body for some $t \in T$ therefore would be covered by the previous two conditions (if $t = a$) or trivially satisfied by a false body. \square

Appendix C. Further technical details of Section 7

Lemma 8. *Let W a DAW-net model and $\text{PDDL}(W)$ its encoding into a planning domain:*

1. *if $(M, \eta) \xrightarrow{t} (M', \eta')$ is a valid firing of W , then there is a ground action a_t of α_t s.t. $(\Psi(M, \eta), a_t, \Psi(M', \eta')) \in \gamma$;*
2. *if there is a ground action a_t of α_t s.t. $(s, a_t, s') \in \gamma$ then $\Psi^-(s) \xrightarrow{t} \Psi^-(s')$ is a valid firing of W .*

Proof. 1. Let $(M, \eta) \xrightarrow{t} (M', \eta')$ be a valid firing of W , then by Definition 11:

- (a) $\{p \in P \mid M(p) > 0\} \supseteq \bullet t$
- (b) for every $p \in P$:²⁵

$$M'(p) = \begin{cases} 0 & \text{if } p \in \bullet t \setminus t \bullet \\ 1 & \text{if } p \in t \bullet \\ M(p) & \text{otherwise} \end{cases}$$

- (c) $\mathcal{D}, \eta \models \text{gd}(t)$,
- (d) assignment η' is such that, if $\text{WR} = \{v \mid \text{wr}(t)(v) \neq \emptyset\}$, $\text{DEL} = \{v \mid \text{wr}(t)(v) = \emptyset\}$:
 - its domain $\text{dom}(\eta') = \text{dom}(\eta) \cup \text{WR} \setminus \text{DEL}$;

25. The simplification derives from the 1-safe assumption.

- for each $v \in \text{dom}(\eta')$:

$$\eta'(v) = \begin{cases} d \in \text{wr}(t)(v) & \text{if } v \in \text{WR} \\ \eta(v) & \text{otherwise.} \end{cases}$$

Let's consider a grounding a_t of $\alpha_t = (t(z_{v_1}, \dots, z_{v_k}), \text{pre}(\alpha_t), \text{eff}(\alpha_t))$ as in Equation (28) where $z_{v_i} = \eta'(v_i)$ if $v_i \in \text{WR}$ and $z_{v_i} = \text{null}$ otherwise. We show that $(\Psi(M, \eta), a_t, \Psi(M', \eta')) \in \gamma$ by proving that $\text{pre}(a)$ is satisfied in $\Psi(M, \eta)$ and that $\gamma(\Psi(M, \eta), a_t) = \Psi(M', \eta')$.

By definition

$$\text{pre}(\alpha_t) = \text{gd}(t)^{\text{PDDL}} \wedge \bigwedge_{v \in \{v_1, \dots, v_k\}} \text{wr}_{t,v}(z_v) \wedge \bigwedge_{p \in \bullet t} (p = \text{TRUE})$$

$\text{gd}(t)^{\text{PDDL}}$ is satisfied because $\mathcal{D}, \eta \models \text{gd}(t)$ and Lemma 7; for all $v_i \in \{v_1, \dots, v_k\}$ $z_{v_i} = \eta'(v_i) \in \text{wr}(t)(v_i)$ or $z_{v_i} = \text{null}$, therefore $\eta'(v_i) \in \text{wr}_{t,v_i}$ as defined in Eq. 26; finally for all $p \in \bullet t$ $p = \text{TRUE}$ because $\{p \in P \mid M(p) > 0\} \supseteq \bullet t$.

By Definition 21 and given that the effects of the action are

$$\text{eff}(\alpha_t) = \bigwedge_{v \in \{v_1, \dots, v_k\}} (v = z_v) \wedge \bigwedge_{p \in \bullet t \setminus t^\bullet} (p = \text{FALSE}) \wedge \bigwedge_{p \in t^\bullet} (p = \text{TRUE})$$

the planning state $\gamma(\Psi(M, \eta), a_t)$ is defined as

$$\begin{aligned} \gamma(\Psi(M, \eta), a_t) = & \{(v \mapsto \eta'(v)) \mid v \in \text{WR} \setminus \text{DEL}\} \cup \\ & \{(v \mapsto \text{null}) \mid v \in \text{DEL}\} \cup \\ & \{(v \mapsto o) \mid v \in \mathcal{V} \setminus \text{WR}, (v \mapsto o) \in \Psi(M, \eta)\} \cup \\ & \{(p \mapsto \text{FALSE}) \mid p \in \bullet t \setminus t^\bullet\} \cup \\ & \{(p \mapsto \text{TRUE}) \mid p \in t^\bullet\} \cup \\ & \{(p \mapsto b) \mid p \in P \setminus (\bullet t \cup t^\bullet), (p, b) \in \Psi(M, \eta)\} \end{aligned}$$

therefore $\gamma(\Psi(M, \eta), a_t) = \Psi(M', \eta')$.

2. Let a_t be a ground action of $\alpha_t = (t(z_{v_1}, \dots, z_{v_k}), \text{pre}(\alpha_t), \text{eff}(\alpha_t))$ and $(s, a_t, s') \in \gamma$, then by Equation (28) and Definition 21

- s satisfies $\text{gd}(t)^{\text{PDDL}}$;
- for all z_v where $v \in \text{WR}$, $z_v \in \text{wr}(t)(v)$ if $v \notin \text{DEL}$ and $z_v = \text{null}$ otherwise;
- $(p \mapsto \text{TRUE}) \in s$ for all $p \in \bullet t$

because s satisfies $\text{pre}(a_t)$, therefore the DAW-net state $(M, \eta) = \Psi(s)$ satisfies

$$\begin{aligned} \{p \in P \mid M(p) > 0\} & \supseteq \bullet t \\ \mathcal{D}, \eta & \models \text{gd}(t) \end{aligned}$$

Moreover the state s' is such that

$$\begin{aligned}
 s' = \gamma(s, a_t) = & \{(v \mapsto z_v) \mid v \in \text{WR} \setminus \text{DEL}\} \cup \\
 & \{(v \mapsto \text{null}) \mid v \in \text{DEL}\} \cup \\
 & \{(v \mapsto o) \mid v \in \mathcal{V} \setminus \text{WR}, (v \mapsto o) \in s\} \cup \\
 & \{(p \mapsto \text{FALSE}) \mid p \in \bullet t \setminus t^\bullet\} \cup \\
 & \{(p \mapsto \text{TRUE}) \mid p \in t^\bullet\} \cup \\
 & \{(p \mapsto b) \mid p \in P \setminus (\bullet t \cup t^\bullet), (p \mapsto b) \in s\}
 \end{aligned}$$

so the DAW-net state $(M', \eta') = \Psi(s')$ satisfies

$$\begin{aligned}
 \eta'(v) &= \begin{cases} z_v \in \text{wr}(t)(v) & \text{if } v \in \text{WR} \setminus \text{DEL} \\ \text{undefined} & \text{if } v \in \text{DEL} \\ \eta(v) & \text{otherwise} \end{cases} \\
 M'(p) &= \begin{cases} 0 & \text{if } p \in \bullet t \setminus t^\bullet \\ 1 & \text{if } p \in t^\bullet \\ M(p) & \text{otherwise} \end{cases}
 \end{aligned}$$

Therefore $\Psi^-(s) \xrightarrow{t} \Psi^-(s')$ is a valid firing.

□