

# The Complexity of Data-Free Nfer<sup>\*</sup>

Sean Kauffman<sup>1</sup>[0000-0001-6341-3898], Kim Guldstrand  
Larsen<sup>2</sup>[0000-0002-5953-3384], and Martin Zimmermann<sup>2</sup>[0000-0002-8038-2453]

<sup>1</sup> Queen’s University, Kingston, Canada  
`sean.k@queensu.ca`

<sup>2</sup> Aalborg University, Aalborg, Denmark  
`{kgl,mzi}@cs.aau.dk`

**Abstract.** Nfer is a Runtime Verification language for the analysis of event traces that applies rules to create hierarchies of time intervals. This work examines the complexity of the evaluation and satisfiability problems for the data-free fragment of nfer. The evaluation problem asks whether a given interval is generated by applying rules to a known input, while the satisfiability problem asks if an input exists that will generate a given interval.

By excluding data from the language, we obtain polynomial-time algorithms for the evaluation problem and for satisfiability when only considering inclusive rules. Furthermore, we show decidability for the satisfiability problem for cycle-free specifications and undecidability for satisfiability of full data-free nfer.

**Keywords:** Interval Logic · Complexity · Runtime Verification

## 1 Introduction

Nfer is an interval logic for analyzing and comprehending event traces that has been used in a wide range of applications, from anomaly detection in autonomous vehicles [14] to spacecraft telemetry analysis [17]. However, its high complexity demands that users restrict the features they incorporate into their applications to ensure tractability. Despite this, no work exists that examines the runtime complexity of nfer without data; an obvious restraint on the power of the language that more closely resembles propositional interval logics like Halpern and Shoham’s logic of intervals (HS) and Duration Calculus [26]. These languages still tend to be undecidable in the general case, however, so it is unclear if this restriction on nfer helps with tractability. In this paper, we show that evaluation of the data-free variant of nfer is tractable and, furthermore, satisfiability for this variant is tractable with additional restrictions.

Nfer was developed by scientists from NASA’s Jet Propulsion Laboratory (JPL) in collaboration with other researchers to analyze event traces from remote systems like spacecraft [18, 16, 17]. In nfer, specifications consist of rules that

---

<sup>\*</sup> Supported by Digital Research Centre Denmark (DIREC) and the Natural Sciences and Engineering Research Council of Canada (NSERC) [ref. RGPIN-2024-04280].

describe and label relationships between time periods called intervals. Applying **nfer** rules to an event trace yields a hierarchy of these intervals that is easier for humans and machines to comprehend than the raw events.

**Nfer** typically operates on intervals with data, but here we define a data-free fragment of the language. Data-free **nfer** is expressive enough for many use cases, having appeared, for example, to analyze the Sequential Sense-Process-Send (SSPS) dataset [17]. Data-free **nfer** is also the target for an algorithm to mine rules from real-time embedded systems [15].

Recent work analyzing the evaluation complexity of **nfer** has shown that it is undecidable for the full language, but with various decidable fragments [20, 21]. These fragments mostly remain intractable, however, with PTIME complexity only possible by employing a meta-constraint on the size of the results that may not be practical in many cases. Those works did not examine data-free **nfer** as a fragment, however, despite it being an obvious restriction with precedent in the literature. A major advantage of restricting **nfer** to the data-free fragment is that the satisfiability problem becomes interesting. With data, it is trivial to show that satisfiability for **nfer** is undecidable. This follows from the results in [21] where **nfer** is shown to have undecidable evaluation. One can encode a Turing machine using **nfer** rules and satisfiability asks if there is an initial tape such that the machine terminates.

Without data, however, it is much less obvious if satisfiability is undecidable. In fact, we show that satisfiability for the full data-free **nfer** language is still undecidable, but we achieve decidability by restricting to a cycle-free or inclusive-only fragment, the latter of which we demonstrate is decidable in PTIME. That the satisfiability of inclusive, data-free **nfer** is decidable in PTIME has exciting implications for practitioners, since these checks can be implemented in event-trace analysis tools [13]. We also show that the evaluation problem for data-free **nfer** is in PTIME without any artificial restrictions on the size of the result from meta-constraints.

All proofs omitted due to space restrictions can be found in [19].

*Related Work* Other works have examined the complexity of interval-based logics. Halpern et al. introduced an interval temporal logic and examined its decidability in [11]. Chaochen et al. found decidable and undecidable fragments of an extension of that work, Duration Calculus [7], over both discrete and dense time [6]. Bolander et al. later introduced Hybrid Duration Calculus (HDC) that added the ability to name an interval and refer to it in a formula [4]. They showed that HDC can express Allen’s relations and is decidable over discrete and dense time domains with non-elementary complexity.

Other works have investigated the complexity of HS [12], a modal logic based on Allen’s Temporal Logic (ATL). Montanari et al. examined the satisfiability problem for the subset of HS over the natural numbers with only *begins/begun by* and *meets* operators and proved it to be EXPSpace-complete [23]. The same authors later proved that adding the *met by* operator increases the complexity of the language to be decidable only over finite total orders [22]. Aceto et al. later examined the expressive power of all fragments of HS over total orders [1].

## 2 Data-Free Nfer

We denote the set of nonnegative integers as  $\mathbb{N}$ . The set of Booleans is given as  $\mathbb{B} = \{true, false\}$ . We fix a finite alphabet  $\Sigma$  of event identifiers and a finite alphabet  $\mathcal{I}$  of interval identifiers such that  $\Sigma \subseteq \mathcal{I}$ . A word is a sequence of identifiers  $\sigma = \sigma_0\sigma_1 \cdots \sigma_{|\sigma|-1}$  where  $\sigma_i \in \Sigma$ . Given a word  $\sigma$ , we define the non-empty subsequence  $\sigma_{[s,e]} = \sigma_s \cdots \sigma_e$ , where  $0 \leq s \leq e \leq |\sigma| - 1$ .

An event represents a named state change in an observed system. An event is a pair  $(\eta, t)$  where  $\eta \in \Sigma$  is its identifier and  $t \in \mathbb{N}$  is the timestamp when it occurred. The set of all events is  $\mathbb{E} = \Sigma \times \mathbb{N}$ . A trace is a sequence of events  $\tau = (\eta_0, t_0)(\eta_1, t_1) \cdots (\eta_{n-1}, t_{n-1})$  where  $n = |\tau|$  and  $t_i \leq t_j$  for all  $i < j$ .

Intervals represent a named period of state in an observed system. An interval is a triple  $(\eta, s, e)$  where  $\eta \in \mathcal{I}$  is its identifier, and  $s, e \in \mathbb{N}$  are the starting and ending timestamps where  $s \leq e$ . We denote the set of all intervals by  $\mathbb{I}$ . A set of intervals is called a *pool* and the set of all pools is  $\mathbb{P} = 2^{\mathbb{I}}$ . We say that an interval  $i = (\eta, s, e)$  is labeled by  $\eta$  and define the accessor functions  $id(i) = \eta$ ,  $start(i) = s$ , and  $end(i) = e$ . An interval of duration zero is an *atomic* interval.

*Syntax.* The data-free **nfer** syntax consists of *rules*. There are two forms of rules: inclusive and exclusive. Inclusive rules test for the existence of two intervals matching a temporal constraint. Exclusive rules test for the existence of one interval and the absence of another interval matching a temporal constraint. When such a pair is found, a new interval is produced with an identifier specified by the rule and timestamps taken from the matched intervals. We define the syntax of these rules as follows:

- Inclusive rules have the form  $\eta \leftarrow \eta_1 \oplus \eta_2$  and
- exclusive rules have the form  $\eta \leftarrow \eta_1 \text{ unless } \ominus \eta_2$

where  $\eta, \eta_1, \eta_2 \in \mathcal{I}$  are identifiers,  $\oplus \in \{\text{before, meet, during, coincide, start, finish, overlap, slice}\}$  is a *clock predicate* on three intervals (one for each of  $\eta, \eta_1$ , and  $\eta_2$ ), and  $\ominus \in \{\text{after, follow, contain}\}$  is a clock predicate on two intervals (one for each of  $\eta_1$  and  $\eta_2$ ). For a rule  $\eta \leftarrow \eta_1 \oplus \eta_2$  or  $\eta \leftarrow \eta_1 \text{ unless } \ominus \eta_2$  we say that  $\eta$  appears on the left-hand and the  $\eta_i$  appear on the right-hand side.

*Semantics.* The semantics of the **nfer** language is defined in three steps: the semantics  $R$  of individual rules on pools, the semantics  $S$  of a specification (a list of rules) on pools, and the semantics  $T$  of a specification on traces of events.

We first define the semantics of inclusive rules with the interpretation function  $R$ . Let  $\Delta$  be the set of all rules. Semantic functions are defined using the brackets  $\llbracket \_ \rrbracket$  around syntax being given semantics.

$$\begin{aligned}
R \llbracket \_ \rrbracket &: \Delta \rightarrow \mathbb{P} \rightarrow \mathbb{P} \\
R \llbracket \eta \leftarrow \eta_1 \oplus \eta_2 \rrbracket \pi &= \\
&\{ i \in \mathbb{I} : i_1, i_2 \in \pi \cdot id(i) = \eta \wedge id(i_1) = \eta_1 \wedge id(i_2) = \eta_2 \wedge \oplus(i, i_1, i_2) \}
\end{aligned}$$

In the definition, an interval  $i$  is a member of the produced pool when two existing intervals in  $\pi$  match the identifiers  $\eta_1$  and  $\eta_2$  and the temporal constraint  $\oplus$ . The identifier of  $i$  is given in the rule and  $\oplus$  defines its start and end timestamps.

The clock predicates referenced by  $\oplus$  are shown in Table 1. These relate two intervals using the familiar ATL temporal operators [2] and also specify the start and end timestamps of the produced intervals. For the example **before**( $i, i_1, i_2$ ),  $i_1$  and  $i_2$  are matched when  $i_1$  ends **before**  $i_2$  begins. The generated interval  $i$  has start and end timestamps inherited from the intervals  $i_1$  and  $i_2$ , i.e., no new timestamps are generated by applying **before**( $i, i_1, i_2$ ). This is true for all other rules as well.

We now define the semantics of exclusive rules with the function  $R$ .

$$\begin{aligned}
R \llbracket \eta \leftarrow \eta_1 \text{ unless } \ominus \eta_2 \rrbracket \pi &= \\
&\{ i \in \mathbb{I} : i_1 \in \pi \cdot id(i) = \eta \wedge id(i_1) = \eta_1 \wedge \\
&\quad start(i) = start(i_1) \wedge end(i) = end(i_1) \wedge \\
&\quad \neg (\exists i_2 \in \pi \cdot i_2 \neq i_1 \wedge id(i_2) = \eta_2 \wedge \ominus(i_1, i_2)) \}
\end{aligned}$$

Like with inclusive rules, exclusive rules match intervals in the input pool  $\pi$  to produce a pool of new intervals. The difference is that exclusive rules produce new intervals where one existing interval in  $\pi$  matches the identifier  $\eta_1$  and no intervals exist in  $\pi$  that match the identifier  $\eta_2$  such that the clock predicate  $\ominus$  holds for the  $\eta_1$ -labeled and the  $\eta_2$ -labeled interval.

The three possibilities referenced by  $\ominus$  are shown in Table 2. These clock predicates relate two intervals using familiar ATL temporal operators while the timestamps of the produced interval are copied from the included interval rather than being defined by the clock predicate. For the example **after**( $i_1, i_2$ ),  $i_1$  and  $i_2$  would be matched (if  $i_2$  existed) if  $i_1$  begins after  $i_2$  ends, and this match would result in *no* interval being produced. If such an interval  $i_2$  is absent, an interval is produced with timestamps matching  $i_1$ .

**Table 1.** Formal definition of **nfer** clock predicates for inclusive rules

$\oplus$	Constraints on $i, i_1$ , and $i_2$
<b>before</b>	$end(i_1) < start(i_2) \wedge start(i) = start(i_1) \wedge end(i) = end(i_2)$
<b>meet</b>	$end(i_1) = start(i_2) \wedge start(i) = start(i_1) \wedge end(i) = end(i_2)$
<b>during</b>	$start(i_2) = start(i) \leq start(i_1) \wedge end(i_1) \leq end(i_2) = end(i)$
<b>coincide</b>	$start(i_1) = start(i_2) = start(i) \wedge end(i_1) = end(i_2) = end(i)$
<b>start</b>	$start(i_1) = start(i_2) = start(i) \wedge end(i) = \max(end(i_1), end(i_2))$
<b>finish</b>	$end(i) = end(i_1) = end(i_2) \wedge start(i) = \min(start(i_1), start(i_2))$
<b>overlap</b>	$start(i_1) < end(i_2) \wedge start(i_2) < end(i_1) \wedge$ $start(i) = \min(start(i_1), start(i_2)) \wedge end(i) = \max(end(i_1), end(i_2))$
<b>slice</b>	$start(i_1) < end(i_2) \wedge start(i_2) < end(i_1) \wedge$ $start(i) = \max(start(i_1), start(i_2)) \wedge end(i) = \min(end(i_1), end(i_2))$

**Table 2.** Formal definition of **nfer** clock predicates for exclusive rules

$\ominus$	Constraints on $i_1$ and $i_2$
<b>after</b>	$start(i_1) > end(i_2)$
<b>follow</b>	$start(i_1) = end(i_2)$
<b>contain</b>	$start(i_2) \geq start(i_1) \wedge end(i_2) \leq end(i_1)$

The interpretation function  $S$  defines the semantics of a finite list of rules, called a specification. Given a specification  $\delta_1 \cdots \delta_n \in \Delta^*$  and a pool  $\pi \in \mathbb{P}$ ,  $S[\_]$  recursively applies  $R[\_]$  to the rules in order, passing each instance the union of  $\pi$  with the intervals returned by already completed calls.

$$S[\_] : \Delta^* \rightarrow \mathbb{P} \rightarrow \mathbb{P}$$

$$S[\delta_1 \cdots \delta_n] \pi = \begin{cases} S[\delta_2 \cdots \delta_n](\pi \cup R[\delta_1] \pi) & \text{if } n > 0 \\ \pi & \text{otherwise} \end{cases}$$

An **nfer** specification  $D \in \Delta^*$  forms a directed graph  $G(D)$  with vertices for the rules in  $D$  connected by edges representing identifier dependencies. An edge exists in  $G(D)$  from  $\delta$  to  $\delta'$  iff there is an identifier  $\eta$  that appears on the left-hand side of  $\delta$  and the right-hand side of  $\delta'$ . We say that  $D$  contains a cycle if  $G(D)$  contains one; otherwise  $D$  is cycle-free.

The rules in a cycle in an **nfer** specification must be iteratively evaluated until a fixed point is reached. As intervals may never be destroyed by rule evaluation, inclusive rules may be repeatedly evaluated, safely. However, exclusive rules may not be evaluated until the intervals on which they depend are known to be present or absent.

For example, suppose a specification with the two rules  $\delta_1 = c \leftarrow a$  **meet**  $b$  and  $\delta_2 = a \leftarrow c$  **meet**  $b$ . Given  $\pi = \{(a, 0, 1), (b, 1, 2), (b, 2, 3), (b, 3, 4), (d, 4, 5)\}$ , we have  $R[\delta_1] \pi = \{(c, 0, 2)\}$  and  $R[\delta_2](\pi \cup R[\delta_1] \pi) = \{(a, 0, 3)\}$ . The rules must be applied a second time to reach a fixed point that includes the interval  $(c, 0, 4)$ . Now consider the consequences if the specification also contained the exclusive rule  $\delta_3 = b \leftarrow d$  **unless follow**  $c$ . After the first evaluation,  $(c, 0, 4)$  is not yet produced, so evaluating  $\delta_3$  would generate  $(b, 4, 5)$ , an incorrect result. As such, *exclusive rules may not appear in cycles* but may appear in a specification that contains cycles among inclusive rules.

To find the cycles in a specification, we compute the strongly-connected components of the directed graph  $G(D)$  formed by the rules in  $D$ . Each strongly connected component represents either a cycle or an individual rule outside of a cycle. We then sort the components in topological order and iterate over each component until a fixed point is reached.

The interpretation function  $T[\_]$  defines the semantics of a specification applied to a trace of events. To ensure consistency with prior work and to simplify our presentation, we overload  $T[\_]$  to operate on an event trace  $\tau \in \mathbb{E}^*$  by first converting  $\tau$  to the pool  $\{init(e) : e \text{ is an element of } \tau\}$  where  $init(\eta, t) = (\eta, t, t)$ .

$$\begin{aligned}
T \llbracket \_ \rrbracket &: \Delta^* \rightarrow \mathbb{P} \rightarrow \mathbb{P} \\
T \llbracket \delta_1 \cdots \delta_n \rrbracket \pi &= \pi_{\ell+1,1}, \pi_{1,1} = \pi \wedge \\
&\mathcal{D} = \text{SCC}(\delta_1 \cdots \delta_n) \wedge (D_1 \cdots D_\ell) = \text{topsort}(\mathcal{D}) . \\
&\pi_{i+1,1} = \bigcup_{j>0} \pi_{i,j}, \pi_{i,j+1} = S \llbracket D_i \rrbracket (\pi_{i,j})
\end{aligned}$$

where  $\text{SCC}(\delta_1 \cdots \delta_n)$  is the set  $\mathcal{D}$  of strongly connected components of the graph  $G(\delta_1 \cdots \delta_n)$  and  $\text{topsort}(\mathcal{D})$  is a topological sort of these components.

### 3 Satisfiability

We are interested in the existential **nfer** satisfiability problem: Given a specification  $D$ , a set of identifiers  $\Sigma$ , and a target identifier  $\eta_T$ , is there an input trace of events  $\tau \in \mathbb{E}^+$  with identifiers in  $\Sigma$  such that an  $\eta_T$ -labeled interval is in  $T \llbracket D \rrbracket \tau$ ? The **nfer** satisfiability problem is interesting in part because of the restriction of input identifiers to  $\Sigma \subseteq \mathcal{I}$ . If  $\eta_T \in \Sigma$ , then any specification is trivially satisfiable. When  $\eta_T \notin \Sigma$ , however, then a  $\eta_T$ -labeled interval must be derived. This problem is non-trivial and, as we shall see, undecidable in general.

To see how data-free **nfer** specifications may be satisfiable or not, consider the following two example specifications for the target identifier  $\eta_T$  and input identifiers  $\Sigma = \{a, b\}$ :

$$D_{\text{sat}} = \left\{ \begin{array}{l} A \leftarrow a \text{ before } b \\ B \leftarrow A \text{ meet } b \\ \eta_T \leftarrow A \text{ overlap } B \end{array} \right. \quad D_{\text{unsat}} = \left\{ \begin{array}{l} A \leftarrow b \text{ before } X \\ B \leftarrow a \text{ meet } b \\ \eta_T \leftarrow a \text{ overlap } B \end{array} \right.$$

A satisfying event trace for  $D_{\text{sat}}$  is  $\tau_1 = (a, 1), (b, 2)$ , since  $T \llbracket D_{\text{sat}} \rrbracket \tau_1 = \{(a, 1, 1), (b, 2, 2), (A, 1, 2), (B, 1, 2), (\eta_T, 1, 2)\}$ . For  $D_{\text{unsat}}$ , no  $\eta_T$ -labeled interval can be produced because **overlap** requires one of the two matched intervals to have positive duration: for an interval  $i$ ,  $\text{end}(i) - \text{start}(i) > 0$ . Since  $a$ -labeled intervals must be initial, they are atomic (zero duration). That leaves  $B$ -labeled intervals produced by another rule. The rule that produces  $B$ -labeled intervals, however, only matches initial intervals with the same timestamps. As such, any  $B$ -labeled interval will also have zero duration, and the **overlap** rule will never be matched. Finally, the **before** rule can never be applied, as no  $X$ -labeled interval can be generated.

#### 3.1 Data-Free nfer Satisfiability is Undecidable

In this section, we show the undecidability of the data-free **nfer** satisfiability problem by a reduction from the emptiness problem for the intersection of two Context-Free Grammars (CFGs). The undecidability result relies on the recursive nature of **nfer**, i.e., an  $\eta$ -labeled interval can be produced from another  $\eta$ -labeled interval, and on its negation capabilities, i.e., via exclusive rules.

**Theorem 1.** *The data-free nfer satisfiability problem is undecidable.*

We now show how to simulate a CFG  $G$  with data-free **nfer** rules with a designated identifier  $\eta_G$  so that a word  $w$  is accepted by the CFG iff applying the rules to events that correspond to  $w$  generates an interval over the same period with identifier  $\eta_G$ . Then the intersection of two CFGs  $G_1$  and  $G_2$  is nonempty if and only if applying the corresponding rules generates, for some sequence of events corresponding to a word, two intervals with the same starting and ending timestamps, one with identifier  $\eta_{G_1}$  and one with  $\eta_{G_2}$ . The existence of two such intervals can again be captured by a data-free **nfer** rule producing an interval with a target identifier.

Formally, a CFG is a four-tuple  $(V, \Sigma, P, S)$ , where  $V$  is a finite set of non-terminals (or variables),  $\Sigma$  is the finite set of terminals that are disjoint from  $V$ ,  $P$  is a finite set of productions of the form  $v \rightarrow w$  where  $v \in V$  and  $w \in (V \cup \Sigma)^*$ , and  $S$  is the initial non-terminal. We assume, without loss of generality, that a CFG is in Chomsky-normal form [8].<sup>3</sup> This means that all productions are in one of two forms:  $A \rightarrow BC$  or  $A \rightarrow a$  where  $A, B, C \in V$ ,  $a \in \Sigma$ , and  $S \notin \{B, C\}$ .

Given a grammar  $(V, \Sigma, P, S)$ , where  $A \in V$ ,  $w, x, y \in (V \cup \Sigma)^*$ , and  $(A \rightarrow x) \in P$ , then we say that  $wAy$  yields  $wxy$ , written  $wAy \Rightarrow wxy$ . We write  $w \xRightarrow{*} y$  if  $w = y$  or there exists a sequence of strings  $x_1, x_2, \dots, x_n$  for  $n \geq 0$  such that  $w \Rightarrow x_1 \Rightarrow x_2 \Rightarrow \dots \Rightarrow x_n \Rightarrow y$ . The language of the grammar is  $\{w \in \Sigma^* : S \xRightarrow{*} w\}$ . If a word is in the language of the grammar we say it has a derivation in the grammar. Deciding if the intersection of the languages of two CFGs is empty is undecidable [3].

We use grammar  $G_x = (V_x, \Sigma_x, P_x, S_x)$  accepting  $\{a(a^n)(b^n) : n > 0\}$  as a running example, where  $V_x = \{A, B, M, M', S_x\}$ ,  $\Sigma_x = \{a, b\}$ , and  $P_x = \{A \rightarrow a, B \rightarrow b, S_x \rightarrow AM, M \rightarrow AB, M \rightarrow AM', M' \rightarrow MB\}$ . The string  $aab$  has the derivation  $S_x \Rightarrow AM \Rightarrow AAB \Rightarrow AAb \Rightarrow Aab \Rightarrow aab$ .

We present six types of data-free **nfer** rules to simulate the intersection of two CFGs  $G = (V, \Sigma, P, S)$  and  $G' = (V', \Sigma, P', S')$ , where  $V \cap V' = \emptyset$  and  $P \cap P' = \emptyset$ . The first four steps are necessary to account for events with coincident timestamps and because simulating a CFG requires the sequential composition of non-terminals, while **nfer** rules cannot perform sequential composition directly on atomic intervals. Then the final two steps map the productions of a CFG and their intersection directly to data-free **nfer** rules. The six types of rules are:

1. Rules to label non-unique timestamps in an event trace so that they can be filtered out. We do so, because event traces in **nfer** are allowed to have events with the same timestamps while there is only one letter at each position of a word. So, to simplify our translation between event traces and words, we just filter out events with non-unique timestamps.
2. Rules that then perform the actual event filtering to only include events with unique timestamps.
3. Rules that label every interval in a trace by its starting event, i.e. where some event occurs at the start and some other event occurs at the end, we label the interval by the starting event.

<sup>3</sup> Note that we, w.l.o.g., disregard the empty word.

4. Rules that select the minimal starting-event-labeled intervals, i.e. the intervals where no other interval is subsumed by that interval. The result of this step is a set of contiguous intervals labeled by their starting event. These minimal intervals are totally ordered and in one-to-one correspondence with the original events with unique timestamps. Thus, we have transformed the event trace into an *equivalent* pool of intervals.
5. Rules that simulate the productions of the CFGs on the pool of minimal starting-event-labeled intervals. The generated intervals encode a derivation tree. The word corresponding to the event trace is accepted by a CFG if an interval is generated that is labeled by that grammar's initial non-terminal.
6. A rule that labels an interval by a given target label if the simulation of the two CFGs labeled the same interval by their initial non-terminals. The interval is generated if the word corresponding to the same event trace is accepted by both CFGs.

We begin by relating event traces to words. Event traces form a total preorder as some timestamps may be equal, while the symbols in a word are totally ordered by their index (no two symbols have the same index). To convert a word to an event trace, we add a timestamp equal to the index of the event. Given a word  $\sigma \in \Sigma^*$ ,  $TRACE(\sigma) = (\sigma_0, 0), (\sigma_1, 1), \dots, (\sigma_{n-1}, n-1)$  where  $n = |\sigma|$ . For example,  $TRACE(abc) = (a, 0), (a, 1), (b, 2)$ .

When converting an event trace to a word, however, we must only consider events with unique timestamps. The following example trace illustrates the reasoning at this step: consider  $\tau_x = (a, 0), (a, 1), (a, 2), (b, 2), (b, 3), (b, 4)$  where both an  $a$ -labeled and a  $b$ -labeled event occur at timestamp 2. To convert  $\tau_x$  to a word, we want to order its identifiers using only their timestamps, and, consequently, the two events with timestamp 2 cannot be ordered. As such, we ensure that only events with unique timestamps affect the generation of intervals involved in simulating a CFG. Before we show how to generate those intervals we define formally what we mean by unique timestamps.

Given a trace  $\tau = (\eta_0, t_0), (\eta_1, t_1), \dots, (\eta_{n-1}, t_{n-1})$ ,  $t_i$  is unique in  $\tau$  if for all  $j \neq i$  we have  $t_j \neq t_i$ . Let  $UNIQ(\tau) = \{t_{i_0}, t_{i_1}, \dots, t_{i_{k-1}}\}$  be the set of unique timestamps in  $\tau$  such that  $t_{i_j} < t_{i_{j'}}$  for all  $j < j'$ , i.e., we enumerate the unique timestamps of  $\tau$  in increasing order. Then,  $WORD(\tau) = \eta_{i_0}, \eta_{i_1}, \dots, \eta_{i_{k-1}}$ .

We now define data-free **nfer** rules that capture the definition of  $WORD$ . The rules first generate atomic **SPOIL** intervals where multiple events share timestamps and then filter the events to only those that do not share timestamps with those **SPOIL** intervals.

Given the alphabet  $\Sigma$ , we define rules to generate **SPOIL** intervals for non-unique timestamps. We let **SPOIL** be a new identifier ( $\text{SPOIL} \notin \Sigma$ ).

$$D_1 = \{\text{SPOIL} \leftarrow a \text{ coincide } b : (a, b) \in \Sigma \times \Sigma \wedge a \neq b\} \quad (1)$$

For example, applying  $D_1$  to the example trace  $\tau_x$  defined above results in the following intervals:  $T[D_1](\tau_x) = \{(a, 0, 0), (a, 1, 1), (a, 2, 2), (b, 2, 2), (\text{SPOIL}, 2, 2), (b, 3, 3), (b, 4, 4)\}$ . Figure 1 shows this example on a timeline. In the top of the figure, the solid line shows time progressing from left to right, with the identifiers



appearing in the trace given below their associated timestamps. The new SPOIL-labeled interval is shown below the timeline, having been generated by the rules in  $D_1$  shown on the right. The remainder of the figure relates to steps 2, 3, and 4.




Fig. 1. Example of applying steps 1-4 from the proof of Theorem 1

**Proposition 1.** Given a trace  $\tau$ ,  $T[[D_1]](\tau)$  characterizes  $WORD(\tau)$  in the following sense:  $\{t : (SPOIL, t, t) \in T[[D_1]](\tau)\}$  is the difference between the set of timestamps in  $\tau$  and  $UNIQ(\tau)$ .

Now, we can define rules that filter the events to only those with unique timestamps by excluding any that coincide with SPOIL-labeled intervals. These rules ensure that the **nfer** simulation of a CFG uses exactly the same events that are used in  $WORD$ : those with unique timestamps. Note that the intervals generated by the rules in  $D_i$  for steps  $i \in \{2, 3, 4\}$  are labeled by identifiers annotated by the step number ( $a_i$ ) where  $a \in \Sigma$  and  $a_i \notin V \cup V' \cup \Sigma$ .

$$D_2 = \{a_2 \leftarrow a \text{ unless contain SPOIL} : a \in \Sigma\} \quad (2)$$

Figure 1 shows the result of applying  $D_2$  to the result of  $T[[D_1]](\tau_x)$ . The intervals  $(a_2, 0, 0)$  and  $(a_2, 1, 1)$  annotate  $a$ -labeled intervals that do not coincide with a SPOIL-labeled interval, while  $(b_2, 3, 3)$  and  $(b_2, 4, 4)$  annotate the  $b$ -labeled intervals that do not coincide with a SPOIL-labeled interval. No such annotated intervals are produced at timestamp 2, where the rules in  $D_1$  generated a SPOIL-labeled interval.

Recall that the intervals that result from the rules in  $D_2$  are still atomic, i.e. they are effectively events and have a duration of zero. The next step is to

use those atomic intervals to generate every interval in the trace with a positive duration (restricted to those with unique starting and ending timestamps in the original trace  $\tau$ ). We label every such interval with a label derived from its start.

$$D_3 = \{a_3 \leftarrow a_2 \text{ before } b_2 : (a, b) \in \Sigma \times \Sigma\} \quad (3)$$

As shown in Figure 1, the intervals generated by applying  $D_3$  in our example are  $(a_3, 0, 1)$  from the rule  $a_3 \leftarrow a_2$  **before**  $a_2$ ,  $(b_3, 3, 4)$  from  $b_3 \leftarrow b_2$  **before**  $b_2$ , and  $(a_3, 1, 3)$ ,  $(a_3, 0, 3)$ ,  $(a_3, 1, 4)$ ,  $(a_3, 0, 4)$  from  $a_3 \leftarrow a_2$  **before**  $b_2$ .

Now, we introduce rules that filter the intervals produced by the rules in  $D_3$  so that only the *minimal* intervals remain. A minimal interval is one where no other interval (with the same label) is subsumed by it. The resulting intervals form a contiguous sequence covering all unique timestamps in  $\tau$  where their meeting points are the atomic intervals produced by  $D_2$ .


$$D_4 = \{a_4 \leftarrow a_3 \text{ unless contain } a_3 : a \in \Sigma\} \quad (4)$$

The reason for generating this contiguous sequence of intervals is that we need to transform the input into elements that are *sequentially composable* using data-free **nfer** rules. To understand why, recall our example event trace:  $\tau_x = (a, 0), (a, 1), (a, 2), (b, 2), (b, 3), (b, 4)$ . As seen in Figure 1, the atomic intervals that result from applying  $D_1$  and  $D_2$  to this trace are  $\{(a_2, 0, 0), (a_2, 1, 1), (b_2, 3, 3), (b_2, 4, 4)\}$ . Because these intervals do not overlap (they are atomic and have unique timestamps) we can see from Table 1 that the only clock predicate that can match two subsequent intervals (i.e., no labeled interval exists between the end of the first and beginning of the second) is **before**. The rules in  $D_3$ , then, do that (match intervals using **before** rules) but these match both subsequent and non-subsequent intervals. Applying the rule  $a_3 \leftarrow a_2$  **before**  $b_2$ , for example, produces  $(a_3, 1, 3)$ ,  $(a_3, 0, 3)$ ,  $(a_3, 1, 4)$ , and  $(a_3, 0, 4)$ . To match only subsequent intervals requires the rules from step four (in  $D_4$ ). Applying  $a_4 \leftarrow a_3$  **unless contain**  $a_3$  only generates  $(a_4, 0, 1), (a_4, 1, 3)$  because they do not *contain* another  $a_3$ -labeled interval, while  $(a_3, 0, 3)$ ,  $(a_3, 1, 4)$ , and  $(a_3, 0, 4)$  do contain  $(a_3, 0, 1)$  and  $(a_3, 1, 2)$ .

At this point, we must discuss what we call the *Ancillary Event Phenomenon*. Because we must generate sequentially composable intervals to simulate a CFG, and because these intervals must label the time *between* events, inevitably one event per trace must be unrepresented by such intervals. Since we choose to label the intervals by their starting event, the final event in the trace with a unique timestamp does not label an interval. We call this the *ancillary event* in a trace. In  $\tau_x$ , the ancillary event is  $(b, 4)$ .

After applying  $D_1 \cup D_2 \cup D_3 \cup D_4$ , we have the intervals  $(a_4, 0, 1), (a_4, 1, 3)$ , and  $(b_4, 3, 4)$ . These intervals are now *sequentially composable* because they (uniquely) **meet** at timestamps 1 and 3, meaning we can use the **meet** clock predicate to match only the contiguous intervals and no others.

With the sequentially composable intervals produced by the rules in  $D_4$ , we now can simulate the productions of the two CFGs. Recall that  $P$  and  $P'$  are



**Fig. 2.** Example of applying step 5 from the proof of Theorem 1

the disjoint sets of these productions.

$$D_5 = \{ A \leftarrow a_4 \text{ coincide } a_4 : (A \rightarrow a) \in P \cup P' \} \cup \{ A \leftarrow B \text{ meet } C : (A \rightarrow BC) \in P \cup P' \} \quad (5)$$

Unlike the rules from  $D_1 \cup D_2 \cup D_3 \cup D_4$ , the rules in  $D_5$  may contain cycles and must be iterated over until a fixed point is reached.

Figure 2 shows the result of applying  $D_5$  to the running example. Each rule in  $D_5$ , shown on the right side of the figure, maps to a production in  $P_x$  and the intervals they produce simulate a derivation for  $\tau_x$  in  $G_x$ . Applying  $A \leftarrow a_4$  **before**  $a_4$  produces  $(A, 0, 1)$  and  $(A, 1, 3)$ , while  $B \leftarrow b_4$  **before**  $b_4$  produces  $(B, 3, 4)$ . Then, applying  $M \leftarrow A$  **meet**  $B$  produces  $(M, 1, 4)$  and applying  $S_x \leftarrow A$  **meet**  $M$  produces  $(S_x, 0, 4)$ . As  $S_x$  is the initial non-terminal for  $G_x$ , an  $S_x$ -labeled interval in the fixed point indicates that the trace  $\tau_x$  during that interval is in the language of  $G_x$ .

Next, we show that the data-free **nfer** simulation has the desired properties. We begin by showing correctness for a single grammar, starting with soundness.

**Lemma 1.** *Given a CFG  $G = (V, \Sigma, P, S)$  and a word  $\sigma \in \Sigma^*$ , fix an identifier  $a \in \Sigma$  for the ancillary event. Then,  $\sigma \in \mathcal{L}(G) \Leftrightarrow (S, 0, |\sigma| - 1) \in T[\bigcup_{i=1}^5 D_i] ( \text{TRACE}(\sigma \cdot a) )$ .*

*Proof.* The proof is by induction over  $j - i$ , showing that for a non-terminal  $A \in V$ ,  $A \xrightarrow{*} \sigma_{[i,j]} \Leftrightarrow (A, i, j) \in T[\bigcup_{i=1}^5 D_i] ( \text{TRACE}(\sigma \cdot a) )$ .  $\square$

We now show completeness for a single grammar.

**Lemma 2.** *Given a CFG  $G = (V, \Sigma, P, S)$  and a trace  $\tau \in \mathbb{E}^*$  such that  $|\tau| \geq 2$ , let  $t$  be the second largest timestamp in  $\text{UNIQ}(\tau)$  and let  $\sigma = \text{WORD}(\tau)$ . Then,  $(S, 0, t) \in T[\bigcup_{i=1}^5 D_i] (\tau) \Leftrightarrow \sigma_{[0, |\sigma| - 2]} \in \mathcal{L}(G)$ .*

*Proof.* By induction over  $j - i$ , showing that for a non-terminal  $A \in V$  and  $\tau = (\eta_0, t_0) \cdots (\eta_{n-1}, t_{n-1})$ ,  $A \xrightarrow{*} \text{WORD}(\tau_{[i,j]}) \Leftrightarrow (A, t_i, t_j) \in T[\bigcup_{i=1}^5 D_i] (\tau)$ .  $\square$

Finally, we check that a word is accepted by both grammars by labeling as  $\eta_T$  where the timestamps of any  $S$ -and- $S'$ -labeled intervals are the same. If any word

has a derivation in both  $G$  and  $G'$ , then applying  $\bigcup_{i=1}^6 D_i$  to the corresponding trace will result in a  $\eta_T$ -labeled interval in the fixed point.

$$D_6 = \{\eta_T \leftarrow S \text{ coincide } S'\} \quad (6)$$

For example, suppose a second grammar  $G'_x$  was introduced accepting the language  $a^+b^+$ , where its initial non-terminal was  $S'_x$ . The word  $aab$  is in the language of  $G'_x$  and so applying  $\bigcup_{i=1}^5 D_i$  for  $G'_x$  to the trace  $\tau_x$  would yield a fixed point containing the interval  $(S'_x, 0, 4)$ . Since  $(S_x, 0, 4)$  coincides with this interval, applying the rule in  $D_6$  will yield  $(\eta_T, 0, 4)$ .

**Lemma 3.** *Given CFGs  $G$  and  $G'$  and a word  $\sigma \in \Sigma^*$ , fix an identifier  $a \in \Sigma$  for the ancillary event. Then,  $\sigma \in \mathcal{L}(G) \cap \mathcal{L}(G') \Leftrightarrow (\eta_T, 0, |\sigma| - 1) \in T[\bigcup_{i=1}^6 D_i](\text{TRACE}(\sigma \cdot a))$ .*

*Proof.* Lemma 1 implies that the CFGs  $G$  and  $G'$  are simulated by  $\bigcup_{i=1}^5 D_i$  for a word  $\sigma$  and applying  $D_6$  finds words in the language of both grammars.  $\square$

**Lemma 4.** *Given CFGs  $G$  and  $G'$  and a trace  $\tau \in \mathbb{E}^*$  such that  $|\tau| \geq 2$ , let  $t$  be the second largest timestamp in  $\text{UNIQ}(\tau)$  and let  $\sigma = \text{WORD}(\tau)$ . Then,  $(\eta_T, 0, t) \in T[\bigcup_{i=1}^6 D_i](\tau) \Leftrightarrow \sigma_{[0, |\sigma| - 2]} \in \mathcal{L}(G) \cap \mathcal{L}(G')$*

*Proof.* Lemma 2 implies that the CFGs  $G$  and  $G'$  are simulated by  $\bigcup_{i=1}^5 D_i$  for a trace  $\tau$  and applying  $D_6$  finds words in the language of both grammars.  $\square$

Now, we can prove Theorem 1.

*Proof.* Applying Lemmas 3 and 4 we obtain that  $\mathcal{L}(G) \cap \mathcal{L}(G')$  is non-empty if and only if there is a  $\tau \in \mathbb{E}^+$  such that  $(\eta_T, \_ , \_ ) \in T[\bigcup_{i=1}^6 D_i](\tau)$ . This shows that the undecidable non-emptiness problem for the intersection of two CFGs can be reduced to the data-free **nfer** satisfiability problem.  $\square$

As satisfiability of data-free **nfer** is undecidable, we now turn our attention to examining fragments with decidable satisfiability. We identify two such fragments: Inclusive **nfer**, where only inclusive rules are permitted, and Cycle-free **nfer**, where specifications can be evaluated without a fixed-point computation.

### 3.2 Inclusive Data-Free nfer Satisfiability is in PTime

We begin our study with the case where an **nfer** specification may contain cycles but only contains inclusive rules.

**Theorem 2.** *The data-free, inclusive nfer satisfiability problem is in PTIME.*

We show that there is a polynomial-time algorithm that determines if an input trace  $\tau$  exists such that an  $\eta_T$ -labeled interval is in  $T[D]\tau$  for a given specification  $D$ . To do this, we show how the satisfiability of a data-free Inclusive-**nfer** specification can be proven through an analysis of the rules without guessing

a witnessing trace. This is due to the monotone nature of inclusive **nfer** rules: new events added to an input trace only add intervals and cannot invalidate existing ones. We leverage this fact to show how only two factors influence the satisfiability of cycle-free, inclusive **nfer** specifications: producibility from events in  $\Sigma$  and the requirement of positive duration for some intervals.

To begin, observe that inclusive **nfer** rules are monotone in nature. The interpretation functions  $R$ ,  $S$ , and  $T$  only add intervals; they never remove them. Furthermore, if the rule is inclusive,  $R$  only tests for the existence of intervals; it only tests for non-existence in the case of exclusive rules. This means that we may always introduce new events into an input trace without needing to keep track of prior results. The consequence is that ensuring that a  $\eta_T$ -labeled interval appears in a fixed-point of  $T[[D]]$   $\tau$  only requires showing that a rule  $\delta_T$  exists in  $D$  with  $\eta_T$  on its left-hand side and that  $\delta_T$  may be matched by intervals resulting from  $\Sigma$ -labeled events. This concept is very similar to graph reachability and we define it here inductively.

**Definition 1.** *Let  $\Sigma$  be a set of input identifiers and  $D$  an inclusive **nfer** specification. An identifier  $\eta$  is producible by  $D$  iff  $\eta \in \Sigma$  or if there exists a rule  $(\eta \leftarrow \eta_1 \oplus \eta_2) \in D$  and both  $\eta_1$  and  $\eta_2$  are producible by  $D$ .*

We now prove that satisfiability for an **nfer** specification using only rules with the **before** operator is equivalent to producibility. We discuss specifications with only **before**-rules here because they allow us to ignore the interaction between events, which have zero duration, and **nfer** operators which require positive duration. We address this complication after proving Proposition 2.

**Proposition 2.** *Given a set of input identifiers  $\Sigma$  and a target identifier  $\eta_T$ , an **nfer** specification  $D_b$  containing only **before** rules is satisfiable iff  $\eta_T$  is producible by  $D_b$ .*

*Proof.* If  $\eta_T \in \Sigma$ , then  $D_b$  is satisfied by the trace  $(\eta_T, 0)$ . If  $\eta_T \notin \Sigma$ , then for  $D_b$  to be satisfiable there must be a rule  $\delta_T$  in  $D_b$  with  $\eta_T$  on its left-hand side. Next, observe from the definition of **before** in Table 1 that the only requirement of  $\delta_T$  to produce a  $\eta_T$ -labeled interval is that there exist intervals  $i_1$  and  $i_2$  such that  $end(i_1) < start(i_2)$ . Clearly, if  $id(i_1) \in \Sigma$  and  $id(i_2) \in \Sigma$  we can create an input trace that satisfies this. If either identifier on the right-hand side of  $\delta_T$  is not in  $\Sigma$ , then apply the same logic inductively for that identifier.

The reverse follows by a similarly straightforward induction: if a  $\eta_T$ -labeled interval is producible then  $D_b$  is satisfiable.  $\square$

To permit inclusive operators beyond **before**, we must address the requirement of *positive duration*. To see why we need to address positive duration, take, for example, the **overlap** operator. Again from Table 1, we see that **overlap** requires that  $start(i_1) < end(i_2)$  and  $start(i_2) < end(i_1)$ . If we assume a zero duration for  $i_2$ , we still must have positive duration for  $i_1$ :  $start(i_1) < start(i_2) = end(i_2) < end(i_1)$ , and the same holds for  $i_2$  if we assume zero duration for  $i_1$ . This means that, for an **overlap**-rule to match, at least one interval it matches

must have positive duration. As such, **overlap**-rules cannot match two initial intervals (events), as they have zero duration.

Thus, producibility is insufficient to show satisfiability for inclusive-**nfer** specifications. We must augment our definition of what is producible to account for what intervals may be produced with positive duration.

Table 3 defines two functions from rules to sets of subsets of identifiers,  $\text{match}_+ : \Delta \rightarrow 2^{2^X}$  and  $\text{add}_+ : \Delta \rightarrow 2^{2^X}$ . The  $\text{match}_+$  function returns the identifiers that must appear in intervals with positive duration for a given rule to match (produce an interval). The  $\text{add}_+$  function returns the identifiers that must appear in intervals with positive duration for a given rule to produce an interval with a positive duration. Both functions return values in Conjunctive-Normal Form (CNF), meaning that at least one element of each set must have positive duration. For example,  $\text{match}_+(\eta \leftarrow \eta_1 \text{meet } \eta_2) = \emptyset$  because **meet** can match two intervals with zero duration, but  $\text{add}_+(\eta \leftarrow \eta_1 \text{meet } \eta_2) = \{\eta_1, \eta_2\}$  because at least one of the two intervals it matches must have positive duration for the result to have positive duration.

**Table 3.** Positive duration requirements on  $\delta = (\eta \leftarrow \eta_1 \oplus \eta_2)$  in CNF (for the sake of readability, we identify a set of sets by a list of sets)

$\oplus$	<i>before</i>	<i>meet</i>	<i>during</i>	<i>coincide</i>	<i>start</i>	<i>finish</i>	<i>overlap</i>	<i>slice</i>
$\text{match}_+(\delta)$	$\emptyset$	$\emptyset$	$\emptyset$	$\emptyset$	$\emptyset$	$\emptyset$	$\{\eta_1, \eta_2\}$	$\{\eta_1, \eta_2\}$
$\text{add}_+(\delta)$	$\emptyset$	$\{\eta_1, \eta_2\}$	$\{\eta_2\}$	$\{\eta_1\}, \{\eta_2\}$	$\{\eta_1, \eta_2\}$	$\{\eta_1, \eta_2\}$	$\{\eta_1, \eta_2\}$	$\{\eta_1, \eta_2\}$

**Definition 2.** Let  $\Sigma$  be a set of input identifiers and  $D$  an inclusive **nfer** specification. An identifier  $\eta$  is positive-duration capable in  $D$  iff there exists a rule  $(\eta \leftarrow \eta_1 \oplus \eta_2) = \delta \in D$  such that for all  $A \in \text{add}_+(\delta)$ , there exists an identifier in  $A$  that is positive-duration capable in  $D$ .

Note that **before**-rules  $\delta$  are always positive-duration capable, as we have  $\text{add}_+(\delta) = \emptyset$  for such  $\delta$ . We now define duration-sensitive producibility for identifiers in an inclusive **nfer** specification using  $\text{match}_+$  and the definition of positive-duration capable identifiers.

**Definition 3.** Let  $\Sigma$  be a set of input identifiers and  $D$  an inclusive **nfer** specification. An identifier  $\eta$  is duration-sensitive producible by  $D$  iff  $\eta \in \Sigma$  or if there exists a rule  $(\eta \leftarrow \eta_1 \oplus \eta_2) = \delta \in D$  such that  $\eta_1$  and  $\eta_2$  are both duration-sensitive producible by  $D$  and for all  $M \in \text{match}_+(\delta)$ , there exists an identifier in  $M$  that is positive-duration capable in  $D$ .

It should be clear that the consequence of Definition 3 is the following lemma.

**Lemma 5.** Given a target identifier  $\eta_T$ , an inclusive **nfer** specification  $D$  is satisfiable iff  $\eta_T$  is duration-sensitive producible by  $D$ .

Now, we can prove Theorem 2.

*Proof.* By Lemma 5, satisfiability for data-free inclusive **nfer** is equivalent to duration-sensitive producibility. As duration-sensitive producibility is defined inductively, there is a straightforward polynomial-time satisfiability algorithm implementing Definition 3.  $\square$

### 3.3 Cycle-Free Data-Free nfer Satisfiability is Decidable

Next, we consider data-free **nfer** with inclusive and exclusive rules, but without cycles. Here, our decidability result is obtained by a transformation to *monadic first order logic* (MFO) (see, e.g., [10] for details) over strings.

**Theorem 3.** *The cycle-free data-free nfer satisfiability problem is decidable.*

*Proof.* Given a cycle-free specification  $D$ , we will, for each identifier  $\eta$ , construct an MFO formula  $\varphi_{D,\eta}(t_0, t_1)$  with free variables  $t_0, t_1$ , such that  $\eta$  is satisfiable with respect to  $D$  iff  $\varphi_{D,\eta}(t_0, t_1)$  is satisfiable, i.e. there exists a string (word)  $w$  over  $2^\Sigma$  and an assignment  $v : \{t_0, t_1\} \rightarrow \{0, \dots, |w| - 1\}$  such that  $w, v \models \varphi_{D,\eta}(t_0, t_1)$ . Note, that for an event identifier  $\eta \in \Sigma$ ,  $\eta(\cdot)$  is a monadic predicate, where  $\eta(t)$  evaluates to true in a string  $w$  over  $2^\Sigma$  at position  $t$  if and only if the set of  $w$  at position  $t$  contains  $\eta$ . That is  $w, v \models \eta(t)$  if and only if  $\eta \in w_{v(t)}$ .

Now from  $w, v$ , where  $w, v \models \varphi_{D,\eta}(t_0, t_1)$ , a satisfying input trace for  $\eta$  may be obtained as the concatenation  $\tau_{w,v} = (w_{t_0}, t_0)(w_{t_0+1}, t_0+1) \dots (w_{t_1}, t_1)$ , where for a set  $\sigma = \{\eta_1, \dots, \eta_k\} \subseteq \Sigma$  and  $t \in \mathbb{N}$ , we denote by  $(\sigma, t)$  the (any) string  $(\eta_1, t)(\eta_2, t) \dots (\eta_k, t)$ .

Since  $D$  is cycle-free, we may order the identifiers by a topological sort of the directed graph formed by the rules. The construction of  $\varphi_{D,\eta}(t_0, t_1)$  now proceeds by induction on this order. In the base case,  $\eta$  is an event identifier. Here  $\varphi_{D,\eta}(t_0, t_1) = (t_0 = t_1 \wedge \eta(t_0))$ . For the inductive case,  $\varphi_{D,\eta}(t_0, t_1)$  is obtained by a disjunction of all the rules for  $\eta$  in  $D$ , i.e.:

$$\varphi_{D,\eta}(t_0, t_1) = \bigvee_{\eta \leftarrow \eta_1 \text{ op } \eta_2 \in D} \psi_{D,\eta_1 \text{ op } \eta_2}(t_0, t_1)$$

where  $\text{op} \in \{\oplus, \text{unless } \ominus\}$ . Here the definition  $\psi_{D,\eta_1 \text{ op } \eta_2}(t_0, t_1)$  is obtained using the MFO formulas for  $\eta_1$  and  $\eta_2$  from the induction hypothesis. Here we just give the definition for two rules, one inclusive and one exclusive rule, leaving the remaining rules for the reader to provide.

$$\begin{aligned} \psi_{D,\eta_1 \text{ before } \eta_2}(t_0, t_1) &= \exists t'_0, t'_1. t_0 \leq t'_1 < t'_0 \leq t_1 \wedge \varphi_{D,\eta_1}(t_0, t'_1) \wedge \varphi_{D,\eta_2}(t'_0, t_1) \\ \psi_{D,\eta_1 \text{ unless after } \eta_2}(t_0, t_1) &= \varphi_{D,\eta_1}(t_0, t_1) \wedge \forall t'_0, t'_1. (t'_0 \leq t'_1 < t_0 \leq t_1) \Rightarrow \\ &\quad \neg \varphi_{D,\eta_2}(t'_0, t'_1) \end{aligned}$$

Thus, decidability of MFO satisfiability over finite strings [5, 9, 25] yields decidability of cycle-free data-free **nfer** satisfiability.  $\square$

Though the reduction to MFO in Theorem 3 yields the desired decidability result, it comes with a non-elementary complexity [24]. We leave it open whether the problem has elementary complexity.

## 4 Evaluation of Data-free nfer

The evaluation problem for **nfer** asks, given a specification  $D$ , a trace  $\tau$  of events, and a target identifier  $\eta_T$ , is there an  $\eta_T$ -labeled interval in  $T[[D]]\tau$ ? The problem has been extensively studied in the presence of data, with complexities ranging from undecidable (for arbitrary data and cycles in the rules) to PTIME (for finite data under the minimality constraint). We refer to [21] for an overview of the results. One case that has not been considered thus far is the complexity of the evaluation problem for data-free **nfer**.

Obviously, the result from [21] for finite-data covers the case of data-free specifications, but without the “minimality” meta-constraint that artificially limits the size of the result, evaluation with only inclusive rules is PSPACE-complete (without cycles) and respectively EXPTIME-complete (with cycles). Here, we show that these results depend on the availability of (finite) data: data-free **nfer** can be evaluated in polynomial time (even without minimality).

**Theorem 4.** *The evaluation problem for data-free nfer is in PTIME.*

*Proof.* Consider an input consisting of a specification  $D$ , a trace  $\tau$  of events, and a target identifier  $\eta_T$ , and let  $k$  be the number of unique timestamps in  $\tau$ .

Recall that an interval is completely specified by its identifier in  $\mathcal{I}$  and its starting and ending timestamp. Hence, as the application of rules does not create new timestamps (cf. Table 1 and Table 2), the number of intervals in  $T[[D]]\tau$  is bounded by  $k^2|\mathcal{I}|$ . Furthermore, whether a rule is applicable to two intervals can be checked in constant time. Thus, one can compute  $T[[D]]\tau$  in polynomial time and then check whether it contains an  $\eta_T$ -labeled interval.  $\square$

## 5 Conclusion and Future Work

We have studied the complexity of the satisfiability and evaluation problems for Data-free **nfer**. We proved that the evaluation problem is in PTIME and the satisfiability problem is undecidable in the general case, but decidable for cycle-free specifications and in PTIME for specifications with only inclusive rules.

There are still open questions around the complexity of **nfer** that may be interesting. We showed that satisfiability for data-free **nfer** is decidable for cycle-free specifications, but we do not prove a tight bound and we suspect it may be possible to achieve improvements on the non-elementary upper-bound we give. Another open question is if satisfiability is decidable for restricted cases of **nfer with data**, for example if specifications are cycle-free and data is finite. We are also interested in the complexity of *monitoring nfer*. Here and in other works, **nfer** is presented with an offline semantics. A naïve monitoring algorithm might simply recompute produced intervals each time a new event arrives, but we suspect that better monitoring complexity can be achieved without requiring assumptions beyond temporal ordering. We hope that this work inspires others to examine the complexity of other modern Runtime Verification (RV) languages.



## References


1. Aceto, L., Della Monica, D., Goranko, V., Ingólfssdóttir, A., Montanari, A., Sciavicco, G.: A complete classification of the expressiveness of interval logics of Allen's relations: the general and the dense cases. *Acta Informatica* **53**(3), 207–246 (Apr 2016). <https://doi.org/10.1007/s00236-015-0231-4>
2. Allen, J.F.: Maintaining knowledge about temporal intervals. *Communications of the ACM* **26**(11), 832–843 (1983)
3. Bar-Hillel, Y., Perles, M., Shamir, E.: On formal properties of simple phrase structure grammars. *STUF - Language Typology and Universals* **14**(1-4), 143–172 (1961). <https://doi.org/doi:10.1524/stuf.1961.14.14.143>, <https://doi.org/10.1524/stuf.1961.14.14.143>
4. Bolander, T., Hansen, J.U., Hansen, M.R.: Decidability of a hybrid duration calculus. *Electronic Notes in Theoretical Computer Science* **174**(6), 113–133 (2007). <https://doi.org/10.1016/j.entcs.2006.11.029>, proceedings of the International Workshop on Hybrid Logic (HyLo 2006)
5. Büchi, J.R.: Weak second-order arithmetic and finite automata. *Mathematical Logic Quarterly* **6**(1-6), 66–92 (1960). <https://doi.org/https://doi.org/10.1002/malq.19600060105>, <https://onlinelibrary.wiley.com/doi/abs/10.1002/malq.19600060105>
6. Chaochen, Z., Hansen, M.R., Sestoft, P.: Decidability and undecidability results for duration calculus. In: *Theoretical Aspects of Computer Science*. pp. 58–68. STACS '93, Springer-Verlag (1993). <https://doi.org/10.5555/646509.694831>
7. Chaochen, Z., Hoare, C., Ravn, A.P.: A calculus of durations. *Information Processing Letters* **40**(5), 269–276 (1991). [https://doi.org/10.1016/0020-0190\(91\)90122-X](https://doi.org/10.1016/0020-0190(91)90122-X)
8. Chomsky, N.: On certain formal properties of grammars. *Information and Control* **2**(2), 137–167 (1959). [https://doi.org/10.1016/S0019-9958\(59\)90362-6](https://doi.org/10.1016/S0019-9958(59)90362-6)
9. Elgot, C.C.: Decision problems of finite automata design and related arithmetics. *Transactions of the American Mathematical Society* **98**(1), 21–51 (1961), <http://www.jstor.org/stable/1993511>
10. Grädel, E., Thomas, W., Wilke, T. (eds.): *Automata, Logics, and Infinite Games: A Guide to Current Research*, LNCS, vol. 2500. Springer (2002). <https://doi.org/10.1007/3-540-36387-4>, <https://doi.org/10.1007/3-540-36387-4>
11. Halpern, J., Manna, Z., Moszkowski, B.: A hardware semantics based on temporal intervals. In: *Automata, Languages and Programming*. pp. 278–291. Springer (1983). <https://doi.org/10.1007/BFb0036915>
12. Halpern, J.Y., Shoham, Y.: A propositional modal logic of time intervals. *J. ACM* **38**(4), 935–962 (1991). <https://doi.org/10.1145/115234.115351>
13. Kauffman, S.: Log analysis and system monitoring with nfer. *Science of Computer Programming* **225** (01 2023). <https://doi.org/10.1016/j.scico.2022.102909>
14. Kauffman, S., Dunne, M., Gracioli, G., Khan, W., Benann, N., Fischmeister, S.: Palisade: A framework for anomaly detection in embedded systems. *Journal of Systems Architecture* **113**, 101876 (2021). <https://doi.org/10.1016/j.sysarc.2020.101876>
15. Kauffman, S., Fischmeister, S.: Mining temporal intervals from real-time system traces. In: *International Workshop on Software Mining (SoftwareMining'17)*. pp. 1–8. IEEE (2017). <https://doi.org/10.1109/SOFTWAREMINING.2017.8100847>
16. Kauffman, S., Havelund, K., Joshi, R.: nfer—a notation and system for inferring event stream abstractions. In: *International Conference on Runtime Verification*

- (RV'16). LNCS, vol. 10012, pp. 235–250. Springer (2016). [https://doi.org/10.1007/978-3-319-46982-9\\_15](https://doi.org/10.1007/978-3-319-46982-9_15)
17. Kauffman, S., Havelund, K., Joshi, R., Fischmeister, S.: Inferring event stream abstractions. *Formal Methods in System Design* **53**, 54–82 (2018). <https://doi.org/10.1007/s10703-018-0317-z>
  18. Kauffman, S., Joshi, R., Havelund, K.: Towards a logic for inferring properties of event streams. In: International Symposium on Leveraging Applications of Formal Methods (ISoLA'16). LNCS, vol. 9953, pp. 394–399. Springer (2016). [https://doi.org/10.1007/978-3-319-47169-3\\_31](https://doi.org/10.1007/978-3-319-47169-3_31)
  19. Kauffman, S., Larsen, K.G., Zimmermann, M.: The complexity of data-free nfer (2024), <https://arxiv.org/abs/2407.03155>
  20. Kauffman, S., Zimmermann, M.: The complexity of evaluating nfer. In: International Symposium on Theoretical Aspects of Software Engineering (TASE'22). LNCS, vol. 13299, pp. 388–405. Springer (07 2022). [https://doi.org/10.1007/978-3-031-10363-6\\_26](https://doi.org/10.1007/978-3-031-10363-6_26)
  21. Kauffman, S., Zimmermann, M.: The complexity of evaluating nfer. *Science of Computer Programming* **231** (2024). <https://doi.org/10.1016/j.scico.2023.103012>
  22. Montanari, A., Puppis, G., Sala, P.: Maximal decidable fragments of halpern and shoham's modal logic of intervals. In: Automata, Languages and Programming. pp. 345–356. Springer (2010). [https://doi.org/10.1007/978-3-642-14162-1\\_29](https://doi.org/10.1007/978-3-642-14162-1_29)
  23. Montanari, A., Puppis, G., Sala, P., Sciavicco, G.: Decidability of the interval temporal logic ABB over the natural numbers. In: Theoretical Aspects of Computer Science. pp. 597–608. STACS '10, Schloss Dagstuhl - Leibniz-Zentrum für Informatik (2010), <https://hal.archives-ouvertes.fr/hal-00717798>
  24. Stockmeyer, L.J.: The Complexity of Decision Problems in Automata Theory and Logic. Ph.D. thesis, MIT (1974)
  25. Trakhtenbrot, B.A.: Finite automata and the logic of one-place predicates. *Siberskii Matematicheskii Zhurnal* **3**(1), 103–131 (1962)
  26. Valentin Goranko, A.M., Sciavicco, G.: A road map of interval temporal logics and duration calculi. *Journal of Applied Non-Classical Logics* **14**(1-2), 9–54 (2004). <https://doi.org/10.3166/jancl.14.9-54>

## Appendix


### A A Graphical Explanation of Clock Predicates

In Figure 3, an interval is represented by a line with its label written above. Here, time moves from left to right and the  $B$ -labeled interval is drawn only once and shared by each example. For the example  $A \leftarrow B$  **before**  $C$ , the  $B$ -labeled and  $C$ -labeled intervals are matched because the  $B$ -labeled interval ends **before** the  $C$ -labeled interval begins.



**Fig. 3.**  $nfer$  clock predicates ( $\oplus$ ) for inclusive rules

In Figure 4, an interval is represented by a line with its label written above and *the absence* of an interval is represented by the line being *dashed*. Like above, time moves from left to right and the  $B$ -labeled interval is drawn only once and shared by each example. For the example  $A \leftarrow B$  **unless after**  $C$ , the  $B$ -labeled and  $C$ -labeled intervals would be matched (if the  $C$ -labeled interval existed), which would result in the  $A$ -labeled interval *not* being produced.



**Fig. 4.**  $nfer$  clock predicates ( $\ominus$ ) for exclusive rules

### B An Algorithm for Duration-Sensitive Producibility

Algorithm 1 implements Definition 3 to determine the satisfiability of an inclusive  $nfer$  specification. It works by computing sets of positive-duration capable identifiers ( $I_+$ ) and duration-sensitive producible identifiers ( $I_\Sigma$ ). Recall that the topological sort of the strongly-connected components of the graph formed by the rules of an  $nfer$  specification gives a dependency ordering for the evaluation of

those rules. This order is computed on Line 2 and the strongly-connected components  $\mathcal{D}$  are iterated over. This set is then looped over  $|\mathcal{D}|$  times on Lines 3 and 4. In the worst case, we must loop  $|\mathcal{D}|$  times to reach a fixed point. For each rule, then, Line 5 tests that both the identifiers on the right-hand side are duration-sensitive producible for that rule. If they are duration-sensitive producible, it then tests on Lines 6 and 8 that the rule's duration requirements are met using  $\text{match}_+$  and  $\text{add}_+$ . If, after iterating over the rules,  $\eta_T$  is labeled as duration-sensitive producible, then the  $D$  is satisfiable.

---

**Algorithm 1** Algorithm checking Data-Free Inclusive-nfer satisfiability

---

**Require:** Specification  $D$ , identifiers  $\Sigma$ , target identifier  $\eta_T$

```

1:  $I_+ \leftarrow \emptyset$ ,  $I_\Sigma \leftarrow \Sigma$ 
2: for  $\mathcal{D} \in \text{topsort}(\text{SCC}(D))$  do
3:   for  $1 \dots |\mathcal{D}|$  do
4:     for  $\eta \leftarrow \eta_1 \oplus \eta_2 \in \mathcal{D}$  do
5:       if  $\{\eta_1, \eta_2\} \cap I_\Sigma = \{\eta_1, \eta_2\}$  then
6:         if  $\forall M \in \text{match}_+(\eta \leftarrow \eta_1 \oplus \eta_2)$ .  $M \cap I_+ \neq \emptyset$  then
7:            $I_\Sigma \leftarrow I_\Sigma \cup \{\eta\}$ 
8:         if  $\forall A \in \text{add}_+(\eta \leftarrow \eta_1 \oplus \eta_2)$ .  $A \cap I_+ \neq \emptyset$  then
9:            $I_+ \leftarrow I_+ \cup \{\eta\}$ 
10:        end for looping over rules in  $\mathcal{D}$ 
11:       end for repeating  $|\mathcal{D}|$  times
12:     end for looping over strongly-connected components of  $D$ 
13:   if  $\eta_T \in I_\Sigma$  then return SAT
14: else return UNSAT

```

---

*Example 1.* Suppose a target identifier  $\eta_T$  and a set of input identifiers  $\Sigma = \{a, b\}$ . We now consider the result of applying Algorithm 1 to a satisfiable list of inclusive rules with no cycles.

1.  $A \leftarrow b$  **before**  $X$ .  $X$  is not in  $I_\Sigma$  ( $I_+ = \emptyset$ ,  $I_\Sigma = \{a, b\}$ )
2.  $A \leftarrow a$  **before**  $b$ .  $a$  and  $b$  are both in  $I_\Sigma$  and **before** matches and adds duration without any other requirements ( $I_+ = \{A\}$ ,  $I_\Sigma = \{a, b, A\}$ )
3.  $B \leftarrow a$  **meet**  $b$ .  $a$  and  $b$  are both in  $I_\Sigma$  and  $\text{match}_+$  for **meet** does not require positive duration, but  $\text{add}_+$  does ( $I_+ = \{A\}$ ,  $I_\Sigma = \{a, b, A, B\}$ )
4.  $\eta_T \leftarrow a$  **overlap**  $B$ .  $a$  and  $B$  are both in  $I_\Sigma$ , but  $\text{match}_+$  and  $\text{add}_+$  require one of  $a, B$  to be positive-duration capable ( $I_+ = \{A\}$ ,  $I_\Sigma = \{a, b, A, B\}$ )
5.  $\eta_T \leftarrow A$  **overlap**  $B$ .  $A$  and  $B$  are both in  $I_\Sigma$ , and  $A$  is positive-duration capable, meeting the requirement ( $I_+ = \{A, \eta_T\}$ ,  $I_\Sigma = \{a, b, A, B, \eta_T\}$ )

**Lemma 6.** *Algorithm 1 runs in  $O(n^2)$  where  $n$  is the number of rules.*

*Proof.* The topological sort of the strongly-connected components of the graph of  $D$  can be computed in linear time. Then each rule is visited at most  $|\mathcal{D}|$  times, if all rules are in the same strongly-connected component. All other operations are sub-linear.  $\square$