

Monitoring First-Order Interval Logic^{*}

Klaus Havelund¹, Moran Omer², and Doron Peled²

¹Jet Propulsion Laboratory, California Institute of Technology, USA

²Department of Computer Science, Bar Ilan University, Israel

Abstract. Runtime verification is used for monitoring the execution of systems, e.g. checking sequences of reported events against formal specifications. Typically the specification refers to the individual monitored events. In this work we perceive the events as defining intervals, each defined by a *begin* and a subsequent *end* event. Allen’s logic allows assertions about the relationship between such named intervals. We suggest a formalism that extends Allen’s logic into a first-order logic that allows quantification over intervals; in addition, intervals can carry data. We provide a monitoring algorithm and describe an implementation and experiments performed with it. We furthermore describe an alternative method for monitoring properties in this logic, by translating them into first-order past-time temporal logic, monitored with the tool DeJaVu.

1 Introduction

Runtime verification allows monitoring of system executions, represented as execution traces, against a specification, either online as traces are generated, or offline after their generation. The monitored trace consists typically of events that can also carry data. The specification is often given using a temporal logic or as a state machine. The runtime algorithm checks for compatibility with the execution in an incremental way, where some summary of the reported execution prefix is updated upon the arrival of newly occurring event. This practice is aimed at both providing an early verdict, and at managing the incremental computational effort between consecutive events. Keeping pace with the speed of the reported events is a challenge to the online monitoring.

While runtime verification, as described above, is concerned with monitoring specifications that refer to single observed *events*, we study here monitoring specifications that refer to observed *intervals*. We consider an interval as being generated from a pair of observed *begin* and *end* events, with appropriate parameters. The focus on intervals is motivated by our experience [14], that engineers, as a way of comprehending complexity, tend to perceive large traces as being partitioned into overlapping sections (intervals), each concerned with a particular task. Temporal logic does not capture this sectional view well, since the formulas get overly complex.

^{*} The research performed by the first author was carried out at Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration. The research performed by the second and third authors was partially funded by Israeli Science Foundation grant 1464/18: “Efficient Runtime Verification for Systems with Lots of Data and its Applications”.

Allen's (temporal) logic [1], also referred to as *Allen's interval algebra*, is a popular formalism for reasoning about the relation between intervals that occur on a timeline. It is often used for planning in AI. Allen's logic deals with a finite set of named intervals, referring directly to the interval names, e.g., $A < B$ means that the interval A must end before the interval B begins. This can be quite restrictive for describing the behavior of systems, where many intervals with the same characteristics can occur, and where distinguishing specific intervals directly by name in the specification is inconvenient or even impossible.

We look at the more general problem of monitoring properties where we can *quantify over intervals*, as e.g., in the formula $\exists A \exists B (A < B)$, stating that there exist (at least) two intervals A and B such that A ends before B begins. We also consider the problem where intervals may contain data. Consequently, the logic allows expressing cases that involve relations between intervals that are embedded in the trace with many, sometimes irrelevant, intervals in between. The runtime verification allows “pattern matching” against these cases in a monitored trace.

We present a matching runtime verification algorithm. The algorithm decides whether any prefix of the execution (the currently observed trace) satisfies the specification. The runtime verification is based on updating a summary of the observed prefixes upon the arrival of each new interval *begin* and *end* event. The trick we employ is to maintain several sets of interval identifiers, and tuples of such, corresponding to the different Allen operators. These variables record those intervals and relations that have begun and not completed yet, as well as those intervals and relations that have been completed. For example, a *begin* event for one interval A followed by a *begin* event for another interval B , is stored (in some variable containing a set of such pairs) as a potential for an A interval, as well as (in a different variable) a potential for an A interval overlapping with a B interval, where A starts, then B starts, then A ends and then B ends. An occurrence of an *end* event for A and then an *end* event for B will complete the picture to decide that A overlapped with B , as well as, of course, having seen completed A and B intervals.

Our logic and runtime verification algorithm is implemented in the tool MonAmi¹. The implementation encodes interval identifiers and data as bit vectors, which are then represented as BDDs. The bit vectors are obtained by a simple enumeration scheme. Such BDDs are useful for compacting interval identifiers and data when storing them in sets, and also makes negation (set complement) non-problematic. We provide an alternative monitoring algorithm by translating into past first-order logic and using the tool DeJaVu. We experiment and compare the two methods.

Related Work The use of BDDs in runtime verification has been explored in [12] for the first-order past time temporal logic DeJaVu, which is an event logic, in contrast to the interval logic explored here. However, the enumeration scheme for creating bit vectors from data and then converting them to BDDs is similar. Numerous event logics have been developed during the past two decades, including [15, 18, 4, 23, 10, 9, 5, 3, 12, 29] to mention just a few.

¹ Monitoring Allen logic modal intervals.

Monitoring of Allen logic is explored in [24]. In that logic, however, intervals are referred to by explicit names, such as $A < B$. This means that one can only specify static patterns, one instance of a particular pattern: that there is one A and one B , such that $A < B$. This is in contrast to MonAmi, where we can quantify over such intervals. Specifically this means that we can specify repeated patterns in the trace e.g., that every interval A with some specific data d is always followed by some other interval B with some data d' .

The most closely related monitoring system is nfer [14, 21, 22], also influenced by Allen's logic. Its specification formalism consists of Prolog-like interval-generating rules (see, e.g., Figure 1). The objective of nfer is to *generate* intervals from a trace of events, as an abstraction of the trace, to e.g. support trace comprehension by humans. Generated intervals can, for example, be visualized. In contrast, the objective of MonAmi is to *verify* intervals, provided as input. nfer only allows a limited form of negation, referred to as *exclusive rules* in [21], making property specification harder, and it is unknown what the limitations are wrt. expressiveness. Our logic allows free negation, and consequently implication. nfer supports Boolean conditions over data as well as computations on data, resulting in new data being stored in the generated intervals. In order to reduce computational complexity, nfer operates in its default mode with a minimality principle, where the before-operator (MonAmi's $<$ operator) only matches the smallest intervals, whereas MonAmi matches all candidate intervals. Section 6 compares MonAmi with nfer further.

A different kind of extension to Allen's logic, where the various relations between operators are promoted into modalities was suggested by Halpern and Shoham [11].

2 Preliminaries

To motivate the study of interval-based specification, we first present the original *Allen Temporal Logic* (ATL).

Syntax. In its basic form, ATL has the following syntax:

$$\varphi ::= (\varphi \wedge \varphi) \mid \neg\varphi \mid A < B \mid AmB \mid AoB \mid AsB \mid AdB \mid AfB \mid A = B$$

where A and B are *intervals* from a finite set of intervals \mathfrak{I} , m stands for *meets*, o for *overlaps*, s for *starts*, d for *during*, and f for *finishes*. The original definition of the logic also includes the symmetric versions of these operators, e.g., an operator for $AmiB$ for BmA , etc., which does not add to the expressive power.

Semantics. A *model* $M = \langle E, \prec, \simeq \rangle$ for Allen's logic, consists of a finite set of events $E = \{begin(A) \mid A \in \mathfrak{I}\} \cup \{end(A) \mid A \in \mathfrak{I}\}$, a linear order $\prec \subseteq E \times E$, and an equivalence relation $\simeq \subseteq E \times E$, where $\preceq = (\prec \cup \simeq)^*$ (the transitive closure of the union of the two relations), such that:

- For each $A \in \mathfrak{I}$, $begin(A) \prec end(A)$.
- \simeq is a partition of the set E into equivalence classes.
- $(\prec \cap \simeq) = \emptyset$.
- For every $a, b \in E$, either $a \preceq b$ or $b \preceq a$.

Thus, M is a linear order between equivalence classes. We call the relation \prec *before*, and \simeq *coincides*. The semantics is given as follows.

- $M \models (\varphi \wedge \psi)$ if $M \models \varphi$ and $M \models \psi$.
- $M \models \neg\varphi$ if $M \not\models \varphi$.
- $M \models A < B$ if $end(A) \prec begin(B)$.
- $M \models AmB$ if $end(A) \simeq begin(B)$.
- $M \models AoB$ if $begin(A) \prec begin(B) \prec end(A) \prec end(B)$.
- $M \models AsB$ if $begin(A) \simeq begin(B)$ and $end(A) \prec end(B)$.
- $M \models AdB$ if $begin(B) \prec begin(A)$ and $end(A) \prec end(B)$.
- $M \models AfB$ if $begin(B) \prec begin(A)$ and $end(A) \simeq end(B)$.
- $M \models A = B$ if $begin(A) \simeq begin(B)$ and $end(A) \simeq end(B)$.

As usual, we can define additional operators, in particular, $(\varphi \vee \psi) = \neg(\neg\varphi \wedge \neg\psi)$ and $(\varphi \rightarrow \psi) = (\neg\varphi \vee \psi)$. As an example, consider then the ATL formula:

$$((B_1 dL \wedge B_2 dL) \wedge B_1 < B_2) \tag{1}$$

It asserts about three intervals B_1 , B_2 and L , that B_1 appears before B_2 and both are embedded within L . Monitoring Allen's logic is described in [24].

3 A First-Order Interval Logic

We will explore now the monitoring of a first-order logic variant of Allen's temporal logic, which we term FoATL. While the original logic refers to a fixed set of intervals, our variant allows quantification over the intervals that occur in the trace, which can optionally carry data. The logic also allows to relate different intervals with respect to their data values. The formalism supports monitoring of behaviors consisting of a large, perhaps unbounded, number of intervals, where patterns of behavior that consist of intervals are related in ways expressed using the specification. For example, a relationship such as in formula (1) can refer to any embedding within a sequence of intervals, matching this pattern, rather than referring to three particular intervals that appear in the input.

The setting. We monitor a sequence of events of the form $begin(z)$ and $end(z)$, where z is a sequence of parameters. The first parameter is an *interval enumeration*, also referred to as *interval id*, used to identify matching *begin* and *end* events; the rest of the parameters, which can be of different types, is optional. An additional parameter can be e.g., a label representing the kind of interval, where a label *Boot* represents that it is a *boot* interval. For example, consider the sequence of events:

$$begin(1, Load), begin(2, Boot), end(2), begin(3, Boot), end(3), end(1)$$

These events form three intervals corresponding to the intervals L , B_1 , and B_2 appearing in ATL formula (1). Our logic alters Allen's logic by adding quantification over the intervals. Hence, instead of fixed intervals, which can be referred to in a formula by their explicit name as constants, we allow interval *variables* A , B , ... that can be instantiated

to any of the intervals that appear in the model (the observed trace). Moreover, the intervals can carry data, and we write in the logic $A(d)$ to denote that the data of the interval assigned to the variable A has the constant value d . We can also verify whether two intervals A and B carry the same value using $same(A, B)$.

We make a few simplifying assumptions in order to concentrate on the main challenges of runtime verification of a first-order interval logic. However, the presented approach is extensible and the restrictions can be easily removed:

- We assume a matching unique integer value per interval, an *enumeration*, though it does not have to appear in consecutive order, is given for each related pair of events, e.g., $begin(5)$ and $end(5)$.
- Events can contain additional parameters besides the enumeration. For simplicity, we assume that there is at most a *single* data value parameter, e.g., an integer or a string, and that it appears within the interval starting event, e.g., $begin(5, abc)$. In a more general setting, different numbers of parameters can appear for different intervals, and the parameters may appear only at the beginning, at the end or in both events defining the interval.
- The monitored events appear one at a time. As there is no co-incidence of events, the relations are restricted to $A < B$ (before), $A o B$ (overlaps) and $A i B$ (for *includes*, which is the symmetric operator of Allen's *d during*). Hence, there is a total order between the events. It reflects the implementation where observed events occur one at a time. It furthermore simplifies the presentation and *incurs no real restriction on the theory involved*.
- Quantification is applied to the (completed) intervals that have occurred. Thus, as in Allen's logic, the specification does not refer to intervals that were opened with $begin(A)$ and were not closed yet with $end(A)$. The logic can of course be extended to deal with unfinished intervals.
- We assume that as part of the monitoring, the restrictions on well formedness of the enumerations are checked. Multiple $begin(A)$ or $end(A)$ events cannot occur for the same interval A , and an $end(A)$ event cannot precede a $begin(A)$ event.
- We allow referring to the data elements in intervals, and also compare them. We offer in the syntax (and our implementation) the predicate *same* that relates intervals with the same data value. This can be extended to other relations that compare values.

Syntax of FoATL. The syntax is as follows.

$$\varphi ::= (\varphi \wedge \varphi) \mid \neg\varphi \mid A(d) \mid (A < B) \mid (A o B) \mid (A i B) \mid \exists A\varphi \mid same(A, B)$$

where A and B are variables (representing intervals) from a set of *interval* variables \mathfrak{I} , and d is a value from some fixed domain D of data values. Parentheses can be removed when clear from the context. A specification does not include free variables. Consider for example the following formula:

$$\exists A \exists B \exists C (A(Load) \wedge B(Boot) \wedge C(Boot) \wedge A i B \wedge A i C \wedge B < C).$$

This specification describes the existence of three intervals with the same relations between them as the intervals L , B_1 , and B_2 appearing in the ATL formula (1).

Semantics of FoATL. Let \mathfrak{I} be the finite set of *interval* variables over the enumerations in the observed execution prefix. We assume the following semantic components:

- $\sigma = e(1)e(2) \dots e(n)$ is a sequence of events of the form $begin(i)$ or $begin(i, d)$, and $end(i)$ as described above.
- $\rho : \mathfrak{I} \mapsto \mathcal{U}$ is a mapping from the interval variables \mathfrak{I} to a domain \mathcal{U} , which can be, e.g., the natural numbers, representing interval enumerations. We denote by $\rho[A \mapsto j]$ the mapping that is identical to ρ but returns the value j for the variable A .
- $data(j)$ is the data value associated with the interval whose enumeration is j .
- $start(j)$ is the number (position in the trace) of the event that starts the interval with enumeration j , i.e., the event $begin(j)$ (with an optional additional data value d).
- $finish(j)$ is the number (position in the trace) of the event that ends the interval with enumeration j , i.e., the event $end(j)$.

We can now define the semantics of the logic inductively on the structure of the formula.

- $(\rho, \sigma) \models (\varphi \wedge \psi)$ if $(\rho, \sigma) \models \varphi$ and $(\rho, \sigma) \models \psi$
- $(\rho, \sigma) \models \neg\varphi$ if $(\rho, \sigma) \not\models \varphi$.
- $(\rho, \sigma) \models A(d)$ if $\rho(A) = j$ and $data(j) = d$.
- $(\rho, \sigma) \models (A < B)$ if $\rho(A) = j$ and $\rho(B) = k$ and $finish(j) < start(k)$.
- $(\rho, \sigma) \models (A \circ B)$ if $\rho(A) = j$ and $\rho(B) = k$ and $start(j) < start(k) < finish(j) < finish(k)$.
- $(\rho, \sigma) \models (A i B)$ if $\rho(A) = j$ and $\rho(B) = k$ and $start(j) < start(k) < finish(k) < finish(j)$.
- $(\rho, \sigma) \models \exists A\varphi$ if there exist events $begin(j)$ (or $begin(j, d)$ for some d) and $end(j)$ in σ such that $\rho' = \rho[A \mapsto j]$ and $(\rho', \sigma) \models \varphi$.
- $(\rho, \sigma) \models same(A, B)$ if $\rho(A) = j$ and $\rho(B) = k$ and $data(j) = data(k)$.

Example properties.

1. $\neg\exists A\exists B(A < B \wedge same(A, B))$.
Disjoint intervals cannot have the same data value.
2. $\neg\exists A\exists B\exists C((A i B \wedge B i C))$.
No double nesting of intervals.
3. $\forall A\forall B((A < B \wedge (\neg\exists C(A < C \wedge C < B))) \rightarrow \neg(A(2) \wedge B(2)))$.
No two adjacent intervals (one completely after the other without any interval in between) can have both the same value 2.
4. $\forall A\forall B\forall C(((A \circ B) \wedge (B \circ C)) \rightarrow \neg(A \circ C))$.
At no point there is an overlapping of three intervals.

Interpretation. One can interpret the semantics of a formula over finite or infinite sequences. As the logic is tailored with an application of runtime verification in mind, one typical use is to require that for a given trace, all prefixes will satisfy a given FoATL specification. This is similar to the common use of temporal specifications of the form $\Box\varphi$, where φ is restricted to past modalities, i.e., to *safety properties* [2], typically seen in

runtime verification, see, e.g., [12, 13]. Nevertheless, other uses are possible as well. Generally, our implementation returns a truth value for the inspected property for each prefix of the monitored trace. Note that satisfaction of a property over an infinite trace does not entail that it is satisfied by all finite prefixes, e.g., for $\varphi = \forall A \exists B (A < B)$, which asserts that there is no *rightmost* interval. Conversely, $\neg\varphi$ is satisfied by every finite trace that includes at least one interval, but will not hold for a trace with infinitely many linearly ordered intervals.

4 The Monitoring Algorithm

Calculating the Relations between Intervals. Recall that in our setting, we are restricted to three possible relations between intervals: $<$, o , and i . Let X and Y be different intervals, defined by *begin* and *end* events, that appeared in the current observed monitored prefix. We distinguish the following three sets of pairs (X, Y) of enumerations of intervals.

- $X < Y$ (*before*). Events appear in the order $begin(X), end(X), begin(Y), end(Y)$.
- $X o Y$ (*overlaps*). Events appear in the order $begin(X), begin(Y), end(X), end(Y)$.
- $X i Y$ (*includes*). Events appear in the order $begin(X), begin(Y), end(Y), end(X)$.

We maintain for each prefix of an execution three sets of pairs of enumerations, $XXYY$ for $X < Y$, $XYXY$ for $X o Y$ and $XYXX$ for $X i Y$. Further sets of pairs (X, Y) correspond to possible prefixes of the four events ($begin(X), end(X), begin(Y)$, and $end(Y)$) in the above three cases, namely XY , XYX , $XYXX$ and XXY . The names of the sets reflect the order of appearance of interval events. For example, XXY represents pairs of intervals where some events of the type $begin(X), end(X), begin(Y)$ have already appeared in this order, but not yet $end(Y)$. When $end(Y)$ subsequently appears, this pair of intervals is removed from XXY and is added to $XXYY$.

We further define the set X of enumerations for events $begin(X)$ where an $end(X)$ has not yet appeared and XX as the set of enumerations, where both $begin(X)$ and $end(X)$ have occurred; this latter is the set of completed intervals. Together, this defines two sets of enumerations, and seven sets of pairs. Note that the names of these variables reflect *patterns* and are not to be taken literally. For example, the set denoted by XX will contain any interval Z where the begin and end events have been observed. It does not only contain intervals specifically named X .

We define these sets inductively on the length i of the trace: for $i = 0$, all the sets are empty; then the update of these sets after the i th event is defined according to Table 1. The rows correspond to the sets that are updated, and the columns to the i th event. The entries in the table detail how the set is updated after the i th event based on the values of the prior values of the sets. For example, for the set X (containing the open intervals), if the i th event is a $begin(Z)$ (or $begin(Z, d)$), then $X_i = X_{i-1} \cup \{Z\}$, and if the i th event is an $end(Z)$ (or $end(Z, d)$), then $X_i = X_{i-1} \setminus \{Z\}$. Our algorithm follows the updates in Table 1 upon arrival of any new event. We denote by \mathcal{U} the universal set of enumerations. The empty set is denoted by \emptyset . We denote by \bar{S} the complement of S , i.e., the set $\mathcal{U} \setminus S$. We will describe later how to implement these sets and operations using BDDs. Note that even through \mathcal{U} , the set of enumerations, can be infinite, at any point

in time we have observed only a finite number of enumerations. Hence, both the current set of observed enumerations and its complement can be represented in a finitary way, as will be described later.

The following rules impose validity checks on the order of the $begin(Z, d)$ (d , the data value, is optional) and $end(Z, d)$ events, causing the system to halt when violated. Specifically, for any interval Z , we allow only one $begin(Z, d)$ respectively one $end(Z, d)$ to occur, and $begin(Z, d)$ must appear before $end(Z, d)$. That is, on observing:

- $begin(Z, d)$: If $\{Z\} \cap (X \cup XX) \neq \emptyset$ then output “multiple begin”.
- $end(Z, d)$: If $\{Z\} \cap XX \neq \emptyset$ then output “multiple end”.
- If $\{Z\} \cap X = \emptyset$ then output “intervals ends before it begins”.

Set \ Event	$begin(Z, d)$	$end(Z, d)$
X (opened)	$X \cup \{Z\}$	$X \cap \{Z\}$
XX (closed)		$XX \cup \{Z\}$
XY	$XY \cup ((X \times \{Z\}))$	$XY \cap (\mathcal{U} \times \{\overline{Z}\}) \cap (\{\overline{Z}\} \times \mathcal{U})$
XYX		$(XYX \cap \overline{XYX}) \cup (XY \cap (\mathcal{U} \times \{Z\}))$
$XYXX$ (X <i>i</i> Y , includes)		$XYXX \cup (XYX \cap (\{Z\} \times \mathcal{U}))$
XYX		$(XYX \cap \overline{XYX}) \cup (XY \cap (\{Z\} \times \mathcal{U}))$
$XYXY$ (X <i>o</i> Y , overlaps)		$XYXY \cup (XYX \cap (\mathcal{U} \times \{Z\}))$
XXY	$XXY \cup (XX \times \{Z\})$	$XXY \cap (\mathcal{U} \times \{\overline{Z}\})$
$XXYY$ ($X < Y$, before)		$XXYY \cup (XXY \cap (\mathcal{U} \times \{Z\}))$
XD (X has data d)		$XD \cup \{(Z, d)\}$

Table 1: The update table.

The order of updating the sets is important: a set that is a prefix of another set, e.g., XY is a prefix set of XYX , hence it is updated *after* the latter. Thus, upon arrival of a new event, the value of XYX is updated based on the *old value* of XY , *before* updating XY .

In order to handle intervals with data, we add another set, XD , of pairs of the form (Z, d) , where Z is an interval enumeration and d is a data element. Then, upon the arrival of an event of the form $end(Z, d)$, we update $XD := XD \cup \{(Z, d)\}$. This construction can be easily extended to capture a different number of parameters n by keeping sets of $n + 1$ tuples.

Using BDDs to represent relations. Our algorithm is based on representing relations between data elements using Ordered Binary Decision Diagrams (OBDD, although we write BDD) [6]. A BDD is a compact representation for a Boolean function (arguments as well as result are Booleans) as a directed acyclic graph (DAG).

A BDD is obtained from a binary tree that represents a Boolean formula with some Boolean variables $x_1 \dots x_k$ by gluing together isomorphic subtrees. Each non-leaf node is labeled with one of the Boolean variables. A non-leaf node x_i is the source of two

arrows leading to other nodes. A dotted arrow represents that x_i has the Boolean value *false* (i.e., 0), while a thick arrow represents that it has the value *true* (i.e., 1). The variables (nodes) in the DAG occur in the same order along all paths from the root (hence the letter ‘O’ in OBDD). Nodes may be absent along some paths, when the result of the Boolean function does not depend on the value of the corresponding Boolean variable. Each path leads to a leaf node that is marked by either *true* or *false*, corresponding to the Boolean value returned by the function for the Boolean values on the path.

A Boolean function, and consequently a BDD, can represent a set of integer values as follows. Each integer value is, in turn, represented using a bit vector: a vector of bits $x_1 \dots x_k$ represents the integer value $x_1 \times 1 + x_2 \times 2 + \dots + x_k \times 2^k$, where the bit value of x_i is 1 for *true* and 0 for *false* and where x_1 is the *least* significant bit, and x_k is the *most* significant. For example, the integer 6 can be represented as the bit vector 110 (here, the most significant bit appears to the left) using the bits $x_1 = 0$, $x_2 = 1$ and $x_3 = 1$. To represent a *set* of integers, the BDD returns *true* for any combination of bits that represent an integer in the set. For example, to represent the set $\{4, 6\}$, we first convert 4 and 6 into the bit vectors 100 and 110, respectively. The Boolean function over x_1, x_2, x_3 is $(\neg x_1 \wedge x_3)$, which returns *true* exactly for these two bit vector combinations.

This representation can be extended to relations, or, equivalently, a set of tuples over integers. Here the Boolean variables are partitioned into n bitstrings $x^1 = x_1^1, \dots, x_{k_1}^1, \dots, x^n = x_1^n, \dots, x_{k_n}^n$, each representing an integer number, forming the bit string²:

$$x_1^1, \dots, x_{k_1}^1, \dots, x_1^n, \dots, x_{k_n}^n.$$

Using BDDs over enumerations of values. Representing data values such as strings and integers, which appear within the observed trace of events, may not lead to a good compact representation. Instead, based on the limited ability to compare data values allowed by FoATL, we represent in the BDD *enumerations* (natural numbers) for these values, rather than the values themselves. When a value (associated with a variable in the specification) appears for the first time in an observed event, we assign to it a new *enumeration*. Values can be assigned consecutive enumeration values³. We use a hash table to point from the value to its enumeration so that in subsequent appearances of this value the same enumeration will be used. For example, if the runtime verifier sees the input events $begin(1, a)$, $begin(2, b)$, $begin(3, c)$, it may encode the data a , b , and c as the bit vectors 000, 001, and 010, respectively. The approach results in several advantages:

1. It allows a shorter representation of very big values in the BDDs; the values are compacted into a smaller number of bits.
2. It contributes to the compactness of the BDDs because enumerations of values that are not far apart often share large bit patterns.

² In the implementation the same number of bits are used for all variables: $k_1 = k_2 = \dots = k_n$.

³ A refined algorithm can reuse enumerations that were used for values that can no longer affect the verdict of the RV process, see [12].

3. The monitoring algorithm is simple; the Boolean operators over summary elements: conjunction, disjunction, and negation, are replaced by the same operators over BDDs.
4. Given an efficient BDD package, the implementation can be very efficient. One can also migrate between BDD packages.
5. It allows full use of negation.

For implementing negation, we keep at least one enumeration value that represents all the enumerations that *did not* occur yet in *begin* and *end* events. For that matter, we can reserve the bitstring 11...11. When the number of values represented by the BDDs grows so that the BDD bits are insufficient, we dynamically add one more bit to the representation, doubling the available number of enumerations.

BDD Operators. We list now the operators on BDDs representing sets of value tuples, used in evaluating the verdict of the specification on the currently inspected prefix. A value tuple represents an interval and its data values, each being elements of the tuple. Recall, however, that we represent data by their enumerations (natural numbers), so we need to represent sets of tuples of enumerations. Recall furthermore that we can represent a tuple of data enumerations as a bit vector: $x_1^1, \dots, x_{k_1}^1, \dots, x_1^n, \dots, x_{k_n}^n$, being the concatenation of the bit vectors for the individual enumerations. A set of such is naturally represented by the BDD that returns true (1) for all the bit-vectors in the set. Useful operators on such BDDs are:

conj(\mathcal{B}, \mathcal{C}) The conjunction (intersection) of the BDDs \mathcal{B} and \mathcal{C} .

comp(\mathcal{B}) The complement of the BDD \mathcal{B} .

project(\mathcal{B}, X) Projects out the Boolean variables $x_1 \dots x_n$ that correspond to the parameter X of \mathcal{B} , obtaining $\exists x_1 \dots \exists x_n \mathcal{B}$.

restrict(z, \mathcal{B}) Restricts a BDD \mathcal{B} of the form XD relating intervals with their data i.e., with bits $x_1 \dots x_n d_1 \dots d_m$ to those sequences of bits where $x_1 \dots x_n$ encodes the interval and $d_1 \dots d_m$ encodes the data value z .

rename($\mathcal{B}, X \leftarrow X', Y \leftarrow Y', \dots$) Replaces the bits $x_1 x_2 \dots x_n$ with $x'_1 \dots x'_n$, the bits $y_1 \dots y_n$ by $y'_1 \dots y'_n$, etc. in the BDD \mathcal{B} .

Other operators, such as e.g. disjunction (union, or database co-join), can be defined in terms of the operators above in the standard way.

Completing the algorithm. The algorithm for the complete logic starts with setting all the sets in Table 1 to BDDs representing the empty sets of elements/pairs, according to their types. Upon the arrival of each new event of the type *begin*(z), (with or without an additional data parameter d) or *end*(z), two steps are executed.

Step 1: The sets of values/pairs are updated according to Table 1.

Step 2: BDDs of the form B_φ for the subformulas φ of the monitored property are updated recursively as follows:

$$- B_{(\varphi \wedge \psi)} = \text{conj}(B_\varphi, B_\psi)$$

- $\mathcal{B}_{\neg\varphi} = \text{comp}(\mathcal{B}_\varphi)$
- $\mathcal{B}_{A(d)} = \text{project}(\text{restrict}(d, \text{rename}(XD, X \leftarrow A)), D)$
- $\mathcal{B}_{A < B} = \text{rename}(XXYY, X \leftarrow A, Y \leftarrow B)$
- $\mathcal{B}_{A \circ B} = \text{rename}(XYXY, X \leftarrow A, Y \leftarrow B)$
- $\mathcal{B}_{A i B} = \text{rename}(XYYX, X \leftarrow A, Y \leftarrow B)$
- $\mathcal{B}_{\exists A\varphi} = \text{project}(\mathcal{B}_\varphi, A)$
- $\mathcal{B}_{\text{same}(A,B)} = \text{project}(\text{conj}(\text{rename}(XD, X \leftarrow A), \text{rename}(XD, X \leftarrow B)), D)$

5 Alternative Algorithm Translating to Past First-Order LTL

Given a representation of intervals as pairs of events of the form $\text{begin}(Z, d)$ and $\text{end}(Z)$, we can perform monitoring by translating the specification into past first-order LTL, referred to as QTL, as used by the tool DeJaVu [12, 20].

Syntax. The formulas of the core QTL logic are defined by the following grammar, where a is a constant representing a value in $\text{domain}(p)$. For simplicity of the presentation, we define here the logic with unary predicates, but this is not due to any principle limitation, and, in fact, DeJaVu supports predicates with multiple arguments, including zero arguments, which correspond to propositions.

$$\varphi ::= \text{true} \mid \text{false} \mid p(a) \mid p(x) \mid (\varphi \vee \varphi) \mid (\varphi \wedge \varphi) \mid \neg\varphi \mid (\varphi \mathcal{S} \varphi) \mid \ominus\varphi \mid \exists x \varphi \mid \forall x \varphi$$

The formulas have the following informal meaning. The formula $p(a)$ is true when the current (last observed) event is $p(a)$. The formula $p(x)$, for some variable $x \in V$, is true if x is bound to a constant a such that $p(a)$ appears as the current event. Variables get bound to constants with the quantifiers \exists and \forall . The formula $(\varphi_1 \mathcal{S} \varphi_2)$ (reads φ_1 *since* φ_2) means that φ_2 occurred in the past (including now) and since then (beyond that state) φ_1 has been true. This is the past dual of the common future time *until* modality. The property $\ominus\varphi$ means that φ is true in the previous step. This is the past dual of the common future time *next* modality. The formula $\exists x \varphi$ is true if there exists a constant a such that φ is true with x bound to a . The formula $\forall x \varphi$ is true if for all constants a , φ is true with x bound to a . We can also define the following additional temporal operators: $P\varphi = (\text{true} \mathcal{S} \varphi)$ (“previously”), and $H\varphi = \neg P\neg\varphi$ (“always in the past” or “historically”).

Semantics. Let σ be a sequence of events and i a natural number. Let γ be an assignment to the variables that appear free in a formula φ . Then $(\gamma, \sigma, i) \models \varphi$ if φ holds for the prefix $s_1 s_2 \dots s_i$ of the trace σ with the assignment γ . This is a standard definition, agreeing, e.g., with [5]. Note that by using past operators, the semantics is not affected by states s_j for $j > i$. Let $\text{free}(\varphi)$ be the set of free (i.e., unquantified) variables of a subformula φ . We denote by $\gamma|_{\text{free}(\varphi)}$ the restriction (projection) of an assignment γ to the free variables appearing in φ . Let ε be an empty assignment. In any of the following cases, $(\gamma, \sigma, i) \models \varphi$ is defined when γ is an assignment over $\text{free}(\varphi)$, and $i \geq 1$.

- $(\varepsilon, \sigma, i) \models \text{true}$.
- $(\varepsilon, \sigma, i) \models p(a)$ if $p(a) \in \sigma[i]$.
- $([v \mapsto a], \sigma, i) \models p(v)$ if $p(a) \in \sigma[i]$.
- $(\gamma, \sigma, i) \models (\varphi \wedge \psi)$ if $(\gamma|_{\text{free}(\varphi)}, \sigma, i) \models \varphi$ and $(\gamma|_{\text{free}(\psi)}, \sigma, i) \models \psi$.

- $(\gamma, \sigma, i) \models \neg\varphi$ if not $(\gamma, \sigma, i) \models \varphi$.
- $(\gamma, \sigma, i) \models (\varphi \mathcal{S} \psi)$ if for some $1 \leq j \leq i$, $(\gamma|_{free(\psi)}, \sigma, j) \models \psi$ and for all $j < k \leq i$, $(\gamma|_{free(\varphi)}, \sigma, k) \models \varphi$.
- $(\gamma, \sigma, i) \models \ominus\varphi$ if $i > 1$ and $(\gamma, \sigma, i-1) \models \varphi$.
- $(\gamma, \sigma, i) \models \exists x \varphi$ if there exists $a \in domain(x)$ such that $(\gamma[x \mapsto a], \sigma, i) \models \varphi$.

The translation from FoATL to QTL is as follows:

- $\mathcal{T}(\varphi \wedge \psi) = \mathcal{T}(\varphi) \wedge \mathcal{T}(\psi)$
- $\mathcal{T}(\neg\varphi) = \neg\mathcal{T}(\varphi)$
- $\mathcal{T}(A(d)) = P(end(A) \wedge \ominus(Pbegin(A, d)))$
- $\mathcal{T}(A < B) = P(end(B) \wedge \ominus(Pbegin(B, Bd)) \wedge \ominus P(end(A) \wedge \ominus Pbegin(A, Ad)))$
- $\mathcal{T}(A \circ B) = P(end(B) \wedge \ominus P(end(A) \wedge \ominus P(begin(B, Bd) \wedge \ominus Pbegin(A, Ad))))$
- $\mathcal{T}(A i B) = P(end(A) \wedge \ominus P(end(B) \wedge \ominus P(begin(B, Bd) \wedge \ominus Pbegin(A, Ad))))$
- $\mathcal{T}(\exists A \varphi) = \exists A \exists Ad \mathcal{T}(\varphi)$
- $\mathcal{T}(same(A, B)) = \exists d(P(end(A) \wedge \ominus Pbegin(A, d)) \wedge P(end(B) \wedge \ominus Pbegin(B, d)))$

It is interesting to note that the translation from FoATL to QTL does not make use of the operator \mathcal{S} , but only uses \ominus and P . The translation has been implemented in MonAmi. We can now monitor a FoATL formula by translating it to QTL using the above translation scheme, and monitor the generated QTL property with DejaVu using the algorithm described in [12]. We later compare the results of monitoring using an optimization of this translation with monitoring using MonAmi.

6 Implementation

We implemented a prototype monitoring tool [19] for our logic FoATL, called MonAmi. It is a Python-based tool for monitoring intervals, formed by events, by checking them against a FoATL property. The tool works with Python 3.6 and above. It uses the ‘dd’ Python package [8] for generating and manipulating BDDs, which itself uses the CUDD BDD package [7] in C. MonAmi uses several input files that define the configuration of the initial parameters, the property file, and the trace file when monitoring in offline mode (log analysis). A trace \mathcal{T} is a sequence of events $[begin, i, d]$ or $[end, i]$, where i is an interval enumeration, and d is the data. The tool can also be used for online monitoring, using the same algorithm, observing a trace dynamically generated by a program during its execution.

6.1 Experiments.

To evaluate MonAmi, we performed a comparison with the interval-based nfer tool [14], mentioned in the related work section on page 3. We expressed four properties using the formalisms of these two tools, all related to receiving data from a *planetary rover*, and evaluated tool performances (time and memory) on traces of different sizes. The planetary rover scenario is inspired by realistic properties of the *Curiosity Mars rover* [17]. The rover’s behavior is reported to ground via the following simplified intervals (amongst many): DL_IMAGE (downlink an image), DL_MOBPRM (downlink

mobility parameter values), DL_ARMPRM (downlink robotic arm parameter values), DL_FAIL (downlink fails), INS_ON (instrument power turned on), INS_FAIL (instrument powering fails), INS_RECOVER (instrument recovers), GET_CAMDATA (reading camera data), STARVE (thread starves), and BOOT (re-boot rover, e.g. after a failure).

The four properties expressed in the formalisms of MonAmi and nfer are shown in Figure 1. In nfer we state a property as a collection of Prolog-like interval-generating rules of the form $id \text{ :- body}$, where the rule body contains Allen’s operators applied to events and intervals generated by other rules. The result of a match of the body is a new interval with the name id , as specified by the rule head. Events and intervals can carry data, which can be used e.g. in **where**-conditions. The IVAL rule (used by all the four properties) generates intervals for all matching (same interval identifier) BEGIN and END events in the trace, and stores (**map**) their interval and data values in the generated IVAL event. The FOUND interval in each nfer property is generated when an error is detected. As mentioned previously, nfer allows negation, referred to as *exclusive rules* in [21]. The body of a rule can e.g. have the form ‘A unless after B’, meaning an A occurred and a B did not occur before. This form of negation has not been used in these properties.

<pre> 1. !exist B1, B2, D . B1('BOOT') & B2('BOOT') & D('DL_IMAGE') & B1 < B2 & (B1 i D B2 i D (B1 < D & D < B2) (B1 o D & ID i B2) (D o B2 & ID i B1)) 2. !exist D, F . (D('DL_MOBPRM') D('DL_ARMPRM')) & F('DL_FAIL') & D i F 3. !exist O, F, R . O('INS_ON') & F('INS_FAIL') & R('INS_RECOVER') & O < F & F < R & !exist X . (X('INS_ON') X('INS_RECOVER')) & O < X & X < R 4. !exist D, G, S . D('DL_IMAGE') & G('GET_CAMDATA') & S('STARVE') & D i S & G i S </pre>	<pre> IVAL :- BEGIN before END where BEGIN.interval = END.interval map { interval -> BEGIN.interval, data -> BEGIN.data } 1. BOOT :- IVAL where IVAL.data = "BOOT" DL :- IVAL where IVAL.data = "DL_IMAGE" DBOOT :- BOOT before BOOT FOUND :- DL during DBOOT 2. DL :- IVAL where IVAL.data = "DL_MOBPRM" IVAL.data = "DL_ARMPRM" FAIL :- IVAL where IVAL.data = "DL_FAIL" FOUND :- FAIL during DL 3. ON :- IVAL where IVAL.data = "INS_ON" FAIL :- IVAL where IVAL.data = "INS_FAIL" RECOVER :- IVAL where IVAL.data = "INS_RECOVER" EXEC :- ON before RECOVER FOUND :- FAIL during EXEC 4. DL :- IVAL where IVAL.data = "DL_IMAGE" GET :- IVAL where IVAL.data = "GET_CAMDATA" STARVE :- IVAL where IVAL.data = "STARVE" FOUND :- STARVE during (GET slice DL) </pre>
---	---

Fig. 1: Evaluated properties in MonAmi (left) and nfer (right).

The properties. Property 1 states that there is no DL_IMAGE during two BOOT intervals (after the start of the first and before the end of the second). Property 2 states that there is no DL_FAIL during a DL_MOBPRM or DL_ARMPRM interval. Property 3 states that there is no INS_FAIL in between an INS_ON and a subsequent closest INS_RECOVER. Note how in the MonAmi specification we need to express the concept of *closest* as an additional constraint (that there is no INS_ON or INS_RECOVER in between). In nfer this is the default semantics, also referred to as the *minimality* principle, see discussion below. Property 4 states that there is no STARVE during a period where

both an DL_IMAGE interval and a GET_CAMDATA interval are active. The nfer **slice** operator produces the intersection between two intervals. As mentioned, nfer’s default execution mode uses a principle of *minimality*, where nfer’s A **before** B operator (analog to MonAmi’s $A < B$ operator) searches the closest right-most B from a given A. The minimality principle, however, can be switched off; so it behaves like MonAmi. Properties 1, 2, and 4 are in nfer evaluated with minimality switched off. nfer was originally designed to run with minimality switched on. However, the C version of nfer offers the option of switching off minimality, while the Scala version was extended with this option in order to perform the experiment.

The traces. We created 5 trace files for each property of different sizes, with 1000, 2000, 4000, 8000, and 16000 events. The traces were generated to evaluate the natural execution mode of MonAmi (stop on first violation) for these properties, by creating the traces to be violated only at the last event. These were generated with a trace generator, guided by one rule for each property. The maximal number of overlapping intervals was also controlled by a parameter (we chose as limit of 3). To ensure that violation will not occur in the middle of the trace we set the data to be different from the ones that appears in the property, except for the violating events. MonAmi is compared to two versions of nfer, a first prototype version in Scala [22], and a later developed version in C [21].

The execution modes. In addition, MonAmi is run in two different modes. Recall from the section *Completing the algorithm* on page 10 that the complete algorithm executes in two steps. In Step 1 the variables in Table 1 are updated. In Step 2, the formula is evaluated based on the value of these variables. When run in *small step* mode (*S*), both steps are executed for each new event. When run in *big step* mode (*B*), only Step 1 is executed for each new event, whereas Step 2 is only executed at the end of monitoring. It corresponds to only observing the formula’s value after the final event, the semantics is unchanged. Small step mode will typically be used for online monitoring, whereas big step mode will typically be used for offline monitoring, e.g. analysis of log files. Obviously, only evaluating Step 2 once at the end provides an optimization. In our case, which is offline log analysis, we shall apply both modes for comparison. nfer evaluates its rules for each new event.

The results. Table 2 shows the results of the evaluation. The experiments were carried out on a Dell Latitude 5401 laptop (Intel Core I7-9850H 9th Gen, 32GB RAM, 512GB SSD) with Ubuntu 20.04.2 LTS OS. W.r.t. memory, nfer/C overall performs the best and nfer/Scala the worst. MonAmi/B (big step) and MonAmi/S (small step) both perform very close to the good performance of nfer. W.r.t. time, again nfer/C has the best performance. MonAmi/B, however, performs as well as or close to nfer/C. MonAmi/S generally performs least well w.r.t. time, except for the second property where nfer/Scala performs worse for larger traces. The first property requires more time than the second property, especially for MonAmi/S. This can be contributed to the higher complexity of the first formula. The better performance of nfer/C in general can potentially be attributed to the fact that it is implemented in C, whereas MonAmi is implemented in a mix of Python and C.

MonAmi and DejaVu. Table 3 shows results of evaluating MonAmi against DejaVu. We evaluated the FoATL properties 1-4 on page 6, monitored by MonAmi, against their translations to QTL, monitored by DejaVu, using a manual translation inspired by the

one presented in Section 5. The manual translation optimizes the resulting QTL formulas. In spite of this optimization, MonAmi clearly outperforms DeJaVu on the translated formulas, both w.r.t. memory use and time. DeJaVu’s evaluation strategy corresponds to MonAmi’s small step evaluation mode since the entire formula is evaluated in each step.

Property	Tool	1000	2000	4000	8000	16000
1	MonAmi/S	1.89 s 51.86 MB	9.46 s 52.43 MB	22.00 s 54.19 MB	72.93 s 78.02 MB	250.55 s 90.50 MB
	MonAmi/B	0.31 s 51.74 MB	0.60 s 52.48 MB	1.25 s 54.56 MB	3.82 s 58.94 MB	6.82 s 86.47 MB
	nfer/Scala	0.19 s 140.41 MB	0.35 s 164.09 MB	1.28 s 395.83 MB	4.42 s 365.73 MB	17.32 s 385.23 MB
	nfer/C	0.03 s 11.03 MB	0.05 s 11.48 MB	0.15 s 12.70 MB	0.52 s 15.15 MB	1.96 s 19.85 MB
2	MonAmi/S	0.37 s 51.71 MB	0.83 s 52.65 MB	2.88 s 54.35 MB	7.98 s 57.30 MB	10.65 s 63.39 MB
	MonAmi/B	0.17 s 51.67 MB	0.30 s 52.27 MB	0.61 s 54.34 MB	1.20 s 57.06 MB	2.47 s 64.27 MB
	nfer/Scala	0.25 s 147.85 MB	0.41 s 196.26 MB	1.19 s 352.84 MB	4.32 s 392.45 MB	18.73 s 662.18 MB
	nfer/C	0.02 s 11.00 MB	0.04 s 11.48 MB	0.14 s 12.75 MB	0.52 s 15.12 MB	1.98 s 19.89 MB
3	MonAmi/S	1.20 s 51.69 MB	3.89 s 52.62 MB	13.06 s 54.30 MB	61.25 s 59.08 MB	385.18 s 86.24 MB
	MonAmi/B	0.19 s 51.82 MB	0.36 s 52.48 MB	0.82 s 54.35 MB	1.69 s 57.09 MB	3.58 s 66.90 MB
	nfer/Scala	0.24 s 142.16 MB	0.44 s 191.50 MB	1.29 s 332.99 MB	4.78 s 391.98 MB	19.82 s 562.61 MB
	nfer/C	0.02 s 11.05 MB	0.05 s 11.49 MB	0.15 s 12.77 MB	0.54 s 15.18 MB	2.12 s 19.91 MB
4	MonAmi/S	0.51 s 51.85 MB	1.49 s 52.55 MB	4.74 s 53.91 MB	17.31 s 57.21 MB	54.80 s 64.79 MB
	MonAmi/B	0.18 s 51.70 MB	0.32 s 52.25 MB	0.72 s 53.88 MB	1.30 s 57.09 MB	2.74 s 65.87 MB
	nfer/Scala	0.20 s 150.56 MB	0.39 s 199.01 MB	1.23 s 402.66 MB	4.86 s 361.00 MB	18.29 s 531.94 MB
	nfer/C	0.02 s 11.10 MB	0.05 s 11.63 MB	0.15 s 13.01 MB	0.54 s 15.67 MB	2.16 s 21.08 MB

Table 2: MonAmi’s S and B modes versus nfer’s Scala and C versions.

7 Conclusion

We described an extension to Allen’s temporal logic, termed FoATL, that allows quantification over the intervals that occur in a monitored trace. We presented an efficient

algorithm for runtime-verification and implemented a prototype tool in Python. The implementation is based on representing sets of tuples of enumerations over the intervals and their data values as BDDs using the ‘dd’ package. We also presented a monitoring procedure that translates a FoATL formula into a first-order past-time temporal logic formula, monitored by the tool DejaVu. Experiments show that the direct implementation of our algorithm is far more efficient.

The closest tool related to MonAmi is nfer and we comment on the relation between these two tools and their capabilities. The FoATL logic allows for a very convenient form of quantification. nfer, in contrast, has the flavor of rule-based programming. FoATL allows free negation, and consequently implication, which is only allowed in a *limited sense* in the C version of nfer, and *not at all* in the Scala version. The limitation (if any) wrt. expressiveness of nfer’s notion of negation is unknown. MonAmi can be extended with time stamps, thereby allowing events to occur at the “same time”, and therefore allowing the Allen operators *meets*, *starts*, *finishes*, and *equals*. nfer relies as default on the minimal interpretation of the before-operator, choosing the closest rightmost interval. MonAmi can be easily extended to also allow this mode. Extending the logic to be first-order also w.r.t. data is considered for future work.

Property	Tool	1000	2000	4000	8000	16000
1	MonAmi/S	0.81 s 211.21 MB	2.14 s 216.38 MB	4.72 s 226.11 MB	13.94 s 248.01 MB	25.14 s 268.81 MB
	MonAmi/B	0.28 s 214.49 MB	0.52 s 217.48 MB	0.98 s 226.99 MB	2.08 s 245.93 MB	4.27 s 275.67 MB
	DejaVu	0.24 s 2.61 GB	0.73 s 2.61 GB	3.94 s 2.63 GB	21.12 s 2.63 GB	136.56 s 4.34 GB
2	MonAmi/S	0.69 s 214.19 MB	1.68 s 217.72 MB	3.52 s 224.88 MB	9.22 s 244.99 MB	26.14 s 272.22 MB
	MonAmi/B	0.27 s 216.28 MB	0.49 s 220.33 MB	1.07 s 224.12 MB	2.19 s 239.32 MB	4.42 s 284.65 MB
	DejaVu	21.82 s 6.09 GB	454.51 s 6.08 GB	∞ N/A	∞ N/A	∞ N/A
3	MonAmi/S	1.33 s 212.67 MB	4.28 s 219.07 MB	12.71 s 231.48 MB	46.47 s 261.21 MB	82.86 s 304.59 MB
	MonAmi/B	0.28 s 217.32 MB	0.57 s 221.24 MB	1.47 s 230.17 MB	2.26 s 236.92 MB	5.13 s 264.54 MB
	DejaVu	0.40 s 6.15 GB	1.36 s 6.14 GB	5.59 s 6.14 GB	38.96 s 6.12 GB	∞ N/A
4	MonAmi/S	0.95 s 210.78 MB	2.36 s 216.76 MB	6.61 s 225.45 MB	23.26 s 240.86 MB	79.95 s 287.96 MB
	MonAmi/B	0.2918 s 217.39 MB	0.54 s 219.58 MB	1.11 s 226.81 MB	2.13 s 248.91 MB	4.78 s 284.80 MB
	DejaVu	2.01 s 6.08 GB	13.67 s 6.08 GB	92.59 s 6.09 GB	∞ N/A	∞ N/A

Table 3: MonAmi’s S and B modes versus DejaVu
(∞ means more than 1000 seconds)

References

1. J. F. Allen, Maintaining Knowledge About Temporal Intervals, *Communications of the ACM*, 26 (11), 832–843.
2. B. Alpern, F. B. Schneider, Recognizing Safety and Liveness. *Distributed Computing* 2(3), 117-126, 1987.
3. B. D’Angelo, S. Sankaranarayanan, C. Sánchez, W. Robinson, B. Finkbeiner, H. B. Sipma, S. Mehrotra, Z. Manna: LOLA: Runtime Monitoring of Synchronous Systems, *TIME* 2005, 166-174.
4. H. Barringer, K. Havelund, TraceContract: A Scala DSL for Trace Analysis, *Proc. of the 17th International Symposium on Formal Methods (FM’11)*, LNCS Volume 6664, Springer, 2011.
5. D. A. Basin, F. Klaedtke, S. Müller, E. Zalinescu, Monitoring Metric First-Order Temporal Properties, *Journal of the ACM* 62(2), 45, 2015.
6. R. E. Bryant, Symbolic Boolean Manipulation with Ordered Binary-Decision Diagrams, *ACM Comput. Surv.* 24(3), 293-318, 1992.
7. CUDD BDD package [<https://davidkebo.com/cudd>]
8. The ‘dd’ Python package for manipulating Binary decision diagrams (BDDs) and Multi-valued decision diagrams (MDDs) [<https://github.com/tulip-control/dd>]
9. N. Decker, M. Leucker, D. Thoma, Monitoring Modulo Theories, *Journal of Software Tools for Technology Transfer*, Volume 18, Number 2, 2016.
10. S. Hallé, R. Villemaire, Runtime Enforcement of Web Service Message Contracts with Data, *IEEE Transactions on Services Computing*, Volume 5 Number 2, 2012.
11. J. Y. Halpern, Y. Shoham, A Propositional Modal Logic of Time Intervals, *Journal of ACM* 38(4), 935-962, 1991.
12. K. Havelund, D. Peled, D. Ulus, First-order Temporal Logic Monitoring with BDDs. *FM-CAD 2017*, 116-123
13. K. Havelund, G. Rosu, Synthesizing Monitors for Safety Properties, *TACAS’02*, LNCS Volume 2280, Springer, 342-356, 2002.
14. S. Kauffman, K. Havelund, R. Joshi, S. Fischmeister, Inferring Event Stream Abstractions. *Formal Methods System Design* 53(1): 54-82, 2018.
15. M. Kim, S. Kannan, I. Lee, O. Sokolsky, Java-MaC: a Run-time Assurance Tool for Java, *Proc. of the 1st Int. Workshop on Runtime Verification (RV’01)*, Elsevier, ENTCS 55(2), 2001.
16. O. Kupferman, M. Y. Vardi. Model Checking of Safety Properties. *Formal Methods System Design*, 19(3), 291-314, 2001.
17. Mars Curiosity Rover [<https://mars.nasa.gov/msl>]
18. P. O. Meredith, D. Jin, D. Griffith, F. Chen, G. Rosu, An Overview of the MOP Runtime Verification Framework, *J. Software Tools for Technology Transfer*, Springer, 2011.
19. MonAmi tool source code [<https://github.com/moraneus/MonAmI>]
20. DejaVu tool source code [<https://github.com/havelund/dejavu>]
21. nfer in C [<http://nfer.io>]
22. nfer in Scala [<https://github.com/rv-tools/nfer>].
23. G. Reger, H. Cruz, D. Rydeheard, MarQ: Monitoring at Runtime with QEA, *Proceedings of the 21st International Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS 2015)*, Springer, 2015.
24. G. Rosu, S. Bensalem, Allen Linear (Interval) Temporal Logic - Translation to LTL and Monitor Synthesis. *CAV 2006*: 263-277.
25. A. P. Sistla, Theoretical Issues in the Design and Analysis of Distributed Systems, Ph.D Thesis, Harvard University, 1983.
26. A. P. Sistla, M. Y. Vardi, P. Wolper, The Complementmentation Problem for Büchi Automata with Applications to Temporal Logic (Extended Abstract), *ICALP 1985*, 465-474, 1984.

27. L. J. Stockmeyer, A. R. Meyer, Word Problems Requiring Exponential Time: Preliminary Report, STOC, 1973: 1-9.
28. W. Thomas, Automata on Infinite Objects, Handbook of Theoretical Computer Science, Volume B: Formal Models and Semantics (B) 1990: 133-191.
29. D. Ulus, O. Maler, Specifying Timed Patterns using Temporal Logic, 21st International Conference on Hybrid Systems: Computation and Control, ACM, 2018, 167-176.