

Runtime Verification of Partially-Synchronous Distributed System

Ritam Ganguly^{1†}, Anik Momtaz^{1†} and Borzoo
Bonakdarpour^{1*}

¹Department of Computer Science and Engineering, Michigan
State University, East Lansing, 48823, sMichigan, USA.

*Corresponding author(s). E-mail(s): borzoo@msu.edu;
Contributing authors: gangulyr@msu.edu; momtazan@msu.edu;

[†]These authors contributed equally to this work.

Abstract

This paper focuses on runtime verification of distributed systems in the *partial synchronous* model, where a clock synchronization algorithm ensures a bound on maximum clock skew among all processes. We introduce two centralized monitoring technique where the specification in the linear temporal logic (LTL) is either represented by a deterministic finite automaton, or, we use a progression-based formula rewriting technique to reduce the distributed runtime verification problem to an SMT problem. We report on rigorous synthetic, as well as real-world case studies involving Cassandra (a non-SQL database management system) and data available from RACE (Runtime for Airspace Concept Evaluation) by NASA. We show that both our automata-based as well as our progression-based approached are effective in evaluating all the possible verdicts given a distributed computation with the progression-based approach having less overhead in general.

Keywords: Runtime Verification, Monitoring, Distributed Computation, Partially Synchronous, Cassandra

1 Introduction

Distributed monitoring consists of evaluating the execution of a distributed application with a centralized or decentralized monitor with respect to a formal specification. A distributed application typically comprises of multiple processes that do not share a global clock and memory, while attempting to accomplish a joint task. For example, in a distributed database, data is stored across different physical locations, possibly dispersed over a network of interconnected computers, and a monitor may want to ensure that queries to the database satisfy some type of consistency criteria.

10 *Motivation*

The main challenge with distributed monitoring lies within the fact that in the absence of a global clock, it is not always possible for the monitor to establish the correct order of occurrence of events across different processes. In fact, given the non-deterministic nature of distributed applications, it is perfectly foreseeable that a runtime monitor may produce different verdicts for the same distributed computation based on different ordering of events. In the case of complete asynchrony, this in turn results in a combinatorial blow-up of possibilities that the monitor must explore at run time, which in turn makes the problem computationally expensive. However, state-of-the-art networks, such as Google Spanner are augmented with clock synchronization techniques that result in partial-synchrony [1]. These clock synchronization techniques guarantee a maximum clock-skew of ε between any pair of processes. Having such a guarantee considerably limits the combinatorial blow-up, as events outside the window of ε can be ordered.

To give an example of the blow-up experienced by the monitor, consider Figure 1, where we have two processes P_1 and P_2 hosting two discrete variables x_1 and x_2 , respectively. Let us also consider the linear temporal logic (LTL) property $\varphi = \bigcirc(x_2 > x_1)$ and a maximum clock-skew, also known as clock-synchronization constant, to be $\varepsilon = 2$. Events $x_1 = 1$ and $x_2 = 0$, as well as $x_1 = 0$ and $x_2 = 2$, are not considered concurrent, as the events in these event pairs are more than ε time apart. However, events $x_1 = 1$ and $x_2 = 2$ are considered concurrent, as these events occurred within ε time from one another. Therefore, it is not possible to determine the exact ordering of these events, without a global clock. Thus, the formula gets evaluated to both true and false, as both possible ordering of events must be taken into account. The




Fig. 1 Distributed computation.

1 number of different possible ordering of events can increase dramatically as
 2 more events and processes are introduced.

3 Handling concurrent events generally results in combinatorial enumeration
 4 of all possibilities and, hence, intractability of distributed RV. Existing
 5 distributed RV techniques operate in two extremes: they either assume a
 6 global clock [2], which is unrealistic for large-scale distributed settings or
 7 assume complete asynchrony [3, 4], which do not scale well. To further elab-
 8 orate on our point, consider the processes P_1 and P_2 in Fig. 2, with events
 9 $\{e_0^1, e_1^1, e_2^1, e_3^1, e_4^1, e_5^1\}$ on process P_1 , and events $\{e_0^2, e_1^2, e_2^2, e_3^2, e_4^2\}$ on process P_2
 10 divided into two segments, seg_1 and seg_2 , and a LTL formula,

$$\varphi = \bigcirc \left(\diamond r \rightarrow (\neg p \mathcal{U} r) \right).$$

11 Observe that the predicate p (resp. r) is true at events e_0^2 and e_4^2 (resp. e_4^1), and
 12 in the rest of the events both predicates are false, denoted by \emptyset . The scenario
 13 where e_0^2 happens before e_0^1 and e_4^1 happens before e_4^2 , the LTL property, φ , is
 14 satisfied. However, the scenario where e_0^1 happens before e_0^2 and e_4^1 happens
 15 after e_4^2 , violates φ .

16 Thus, following the above example, the main research problem we aim to
 17 tackle in this paper is the following. Given a finite distributed computation and
 18 an LTL formula, our objective is to design efficient algorithms that determine
 19 whether or not the computation satisfies the formula. As shown above, the
 20 main obstacle is solving this problem is the explosion of interleavings at run
 21 time that need to be explored in order to monitor a computation.

22 Contributions

23 In order to address the combinatorial explosion of various interleavings intro-
 24 duced by the absence of a global clock, our first design choice is a practical
 25 assumption, namely, a bounded skew of ε between local clocks of each pair
 26 of processes, which is guaranteed by a clock synchronization mechanism (e.g.,
 27 NTP [5]).

28 Our first technique is based on constructing the LTL₃ [6] monitor automa-
 29 tion of an LTL formula and constructing multiple SMT queries to determine
 30 which states of the monitor automaton are reachable for a given distributed
 31 computation. For example, Fig. 3 shows the monitor automaton for formula φ
 32 mentioned earlier and one has to construct 4 different SMT queries to deter-
 33 mine the set of all possible reachable states at the end of the computation in




Fig. 2 Distributed computation.




Fig. 3 Monitor automaton for formula φ .

1 Fig. 2. We transform our monitoring decision problem into an SMT solving
 2 problem. The SMT instance includes constraints that encode (1) our monitoring
 3 algorithm based on the 3-valued semantics of LTL [6], (2) behavior of
 4 communicating processes and their local state changes in terms of a distributed
 5 computation, and (3) the happened-before relation subject to the ϵ clock skew
 6 assumption. Then, it attempts to concretize an uninterpreted function whose
 7 evaluation provides the possible verdicts of the monitor with respect to the
 8 given computation. In order to make the verification problem tractable, we
 9 chop a computation into multiple segments and effectively reduce the search
 10 space of each SMT query (see Fig. 4). Thus, the result of monitoring each
 11 segment (the possible LTL₃ states) should be carried to the next segment. Fur-
 12 thermore, given the fact that distributed applications nowadays run on massive
 13 cloud services, we extend our solution to a parallel monitoring algorithm to
 14 utilize the available computing infrastructure and achieve better scalability.

15 The intuition behind our second monitoring technique is that since (in the
 16 first approach) running SMT queries to test whether each state of the LTL₃
 17 monitor automaton is reachable is excessive, it should be sufficient to test
 18 whether temporal sub-formulas of an LTL formula hold in a distributed com-
 19 putation. Similar to the first approach, we utilize segmentation, to break down
 20 the problem size. In the second, approach to carry the result of monitoring
 21 from one segment to the next, we also develop a *formula progression* technique.
 22 Specifically, given a finite trace α , and an LTL formula φ , we define a function
 23 Pr , such that $\text{Pr}(\alpha, \varphi)$ characterizes the *progression* of φ and α . Progression is
 24 defined as the rewritten formula for future extensions of α depending on what
 25 has been observed thus far, which returns either **true**, **false**, or an LTL for-
 26 mula. We emphasize that the main difference between our technique and the
 27 classic rewriting technique [7] is that, function Pr takes a finite trace as input
 28 while the algorithm in [7] rewrites the input LTL formula in a state-by-state
 29 manner. This means that in our setting, rewriting based on the fixed point
 30 representation of temporal operators is not possible. Our motivation is due to
 31 the fact that when a given distributed computation is chopped into a number
 32 of segments then a state-by-state rewriting approach would incur too many




Fig. 4 Progression and segmentation.

1 SMT queries, making it unscalable. For example, in Fig. 4 (which is the com-
 2 putation in Fig. 2 chopped to two segments), our progression-based approach
 3 needs the same 4 SMT queries for seg_1 (2 for each of the sub-formulas $\diamond r$
 4 and $\square(\neg p)$) as compared to [8]. The evaluation yields formulas $\neg(\diamond r)$ and
 5 $\diamond r \rightarrow (\neg p \mathcal{U} r)$ as the possible formulas and as a result we only need to build
 6 4 SMT queries in seg_2 compared to 5 for the automata-based approach in [8].

7 Our method is fully implemented and the datasets generated during
 8 and/or analysed during the current study are available in [https://github.com/
 9 TART-MSU/dist-ltl-rv](https://github.com/TART-MSU/dist-ltl-rv). We make a detailed comparison between the proposed
 10 approaches in this paper through not only a set of vigorous synthetic
 11 experiments, but also monitoring the same set of consistency conditions in Cas-
 12 sandra. We also put our approach to test using a real-time airspace monitoring
 13 dataset (RACE) from NASA [9]. Our experiments show that the progression-
 14 based approach has 35% reduced overhead (See Section 6 as compared to the
 15 automata-based approach.

16 In summary, the main contributions of this paper is as follows:

- 17 • We transform our monitoring decision problem into an SMT¹ problem, to
 18 make for an efficient yet correct approach to consider different interleavings.
 19 Given an LTL formula, our solution provides all possible verdicts on a given
 20 computation.
- 21 • We present two monitoring approaches to address the challenges (men-
 22 tioned earlier) of distributed runtime verification with regard to LTL formulas
 23 under a partially synchronous setting. In our first approach, we keep track
 24 of the observed events and the possible future outcomes by employing
 25 an automata-based technique. In our second approach, we employ a more
 26 efficient progression-based technique, where we rewrite the given LTL speci-
 27 fications based on the current observations. For both of our approaches, we
 28 consider a fault-proof central monitor.
- 29 • We divide a given computation into multiple segments in order to make
 30 the verification problem tractable, and as a result, significantly reduce the
 31 search space of each SMT query. Furthermore, we parallelize our monitoring

¹Satisfiability modulo theories (SMT) is the problem of determining whether a formula involving Boolean expressions comprising of more complex formulas involving real numbers, integers, and/or various data structures is satisfiable.

1 technique in order to utilize the available computational resources and gain
2 greater scalability.

- 3 • Finally, we explore and report on extensive comparisons between our
4 automata-based approach and our progression-based approach in terms of
5 runtime and complexity.

6 *Comparison to the conference submission*

7 A preliminary version of this paper appeared in the 2020 International Con-
8 ference on Principles of Distributed Systems (OPODIS) [8]. The OPODIS’20
9 paper only included the automata-based approach. This paper extends the
10 OPODIS’20 paper in multiple fronts. First, the progression-based technique is
11 completely new. This includes the notion of progression functions for LTL and
12 its SMT formulation. This technique is also fully implemented and its perfor-
13 mance is rigorously analyzed and compared with the original automata-based
14 approach.

15 *Organization*

16 Section 2 presents the background concepts and the problem statement. We
17 go over the theory of formula progression in Section 3 and discuss SMT-based
18 approach both of the automata-based solution and the progression-based solu-
19 tion in Section 4. We introduce some optimization approaches to yield better
20 run time in Section 5, while we go over the analysis of experimental results in
21 Section 6. Related work is discussed in Section 7. Finally, we make concluding
22 remarks in Section 8.

23 **2 Preliminaries and Problem Statement**

24 **2.1 Linear Temporal Logics (LTL) for RV**

25 Let AP be a set of *atomic propositions* and $\Sigma = 2^{\text{AP}}$ be the set of all possible
26 *states*. A *trace* is a sequence $s_0s_1\dots$, where $s_i \in \Sigma$ for every $i \geq 0$. We denote
27 by Σ^* (resp., Σ^ω) the set of all finite (resp., infinite) traces. For a finite trace
28 $\alpha = s_0s_1\dots s_k$, $|\alpha|$ denotes its *length*, $k + 1$. Also, for $\alpha = s_0s_1\dots s_k$, by α^i ,
29 we mean trace $s_i s_{i+1} \dots s_k$ of α .

30 **2.1.1 Infinite-trace Semantics of LTL**

31 The syntax and semantics of the *linear temporal logic* (LTL) [10] are defined
32 for infinite traces. The syntax is defined by the following grammar:

$$\varphi ::= p \mid \neg\varphi \mid \varphi \vee \varphi \mid \bigcirc\varphi \mid \varphi \mathcal{U} \varphi$$

33 where $p \in \text{AP}$, and where \bigcirc and \mathcal{U} are the ‘next’ and ‘until’ temporal oper-
34 ators respectively. We view other propositional and temporal operators as
35 abbreviations, that is, $\text{true} = p \vee \neg p$, $\text{false} = \neg\text{true}$, $\varphi \rightarrow \psi = \neg\varphi \vee \psi$,




Fig. 5 LTL₃ monitor for $\varphi = a \mathcal{U} b$.

1 $\varphi \wedge \psi = \neg(\neg\varphi \vee \neg\psi)$, $\diamond\varphi = \text{true} \mathcal{U} \varphi$ (eventually φ), and $\square\varphi = \neg\diamond\neg\varphi$
 2 (always φ). We denote the set of all LTL formulas by Φ_{LTL} .

3 The infinite-trace semantics of LTL is defined as follows. Let $\sigma =$
 4 $s_0s_1s_2\cdots \in \Sigma^\omega$, $i \geq 0$, and let \models denote the *satisfaction* relation:

$$\begin{array}{ll}
 \sigma, i \models p & \text{iff } p \in s_i \\
 \sigma, i \models \neg\varphi & \text{iff } \sigma, i \not\models \varphi \\
 \sigma, i \models \varphi_1 \vee \varphi_2 & \text{iff } \sigma, i \models \varphi_1 \text{ or } \sigma, i \models \varphi_2 \\
 \sigma, i \models \bigcirc\varphi & \text{iff } \sigma, i+1 \models \varphi \\
 \sigma, i \models \varphi_1 \mathcal{U} \varphi_2 & \text{iff } \exists k \geq i : \sigma, k \models \varphi_2 \text{ and } \forall j \in [i, k) : \sigma, j \models \varphi_1
 \end{array}$$

5 Also, $\sigma \models \varphi$ holds if and only if $\sigma, 0 \models \varphi$ holds.

6 2.1.2 Finite-trace Semantics of LTL

7 In the context of RV, the 3-valued LTL (LTL₃ for short) [6] evaluates LTL
 8 formulas for *finite* traces, but with an eye on possible future extensions where
 9 as finite LTL, or FLTL [11] only takes into consideration the current trace with
 10 no eye towards the future. In LTL₃, the set of truth values is $\mathbb{B}_3 = \{\top, \perp, ?\}$,
 11 where \top (resp., \perp) denotes that the formula is *permanently* satisfied (resp.,
 12 violated), no matter how the current finite trace extends, and ‘?’ denotes an
 13 *unknown* verdict, i.e., there exists an extension that can violate the formula,
 14 and another extension that can satisfy the formula. Let $\alpha \in \Sigma^*$ be a non-empty
 15 finite trace. The truth value of an LTL₃ formula φ with respect to α , denoted
 16 by $[\alpha \models_3 \varphi]$, is defined as follows:

$$[\alpha \models_3 \varphi] = \begin{cases} \top & \text{if } \forall \sigma \in \Sigma^\omega : \alpha\sigma \models \varphi \\ \perp & \text{if } \forall \sigma \in \Sigma^\omega : \alpha\sigma \not\models \varphi \\ ? & \text{otherwise.} \end{cases}$$

17

18 **Definition 1** The LTL₃ *monitor* for a formula φ is the unique deterministic finite
 19 state machine $\mathcal{M}_\varphi = (\Sigma, Q, q_0, \delta, \lambda)$, where Q is the set of states, q_0 is the initial
 20 state, $\delta : Q \times \Sigma \rightarrow Q$ is the transition function, and $\lambda : Q \rightarrow \mathbb{B}_3$ is a function such
 21 that $\lambda(\delta(q_0, \alpha)) = [\alpha \models_3 \varphi]$, for every finite trace $\alpha \in \Sigma^*$. ■

8 *Runtime Verification of Partially-Synchronous Distributed System*

1 For example, Fig. 5, shows the monitor automaton for formula $\varphi = a \mathcal{U} b$.
 2 The syntax of FLTL is also identical to that of LTL, and the semantics is based
 3 on the truth values $\mathbb{B}_2 = \{\top, \perp\}$, where \top (resp., \perp) denotes that the formula is
 4 satisfied (resp., violated) given the current finite trace. For atomic propositions
 5 and Boolean operators, the semantics of FLTL is identical to those of LTL. Let
 6 φ , φ_1 , and φ_2 be LTL formulas, $\alpha = s_0 s_1 \dots s_n$ be a *non-empty* finite trace,
 7 and \models_F denote the satisfaction relation in FLTL. The semantics of FLTL for
 8 the temporal operators are as follows:

$$[\alpha \models_F \bigcirc \varphi] = \begin{cases} [\alpha^1 \models_F \varphi] & \text{if } \alpha^1 \neq \epsilon \\ \perp & \text{otherwise.} \end{cases}$$

9

$$[\alpha \models_F \varphi_1 \mathcal{U} \varphi_2] = \begin{cases} \top & \text{if } \exists k \in [0, n] : ([\alpha^k \models_F \varphi_2] = \top) \wedge \\ & \forall l \in [0, k] : ([\alpha^l \models_F \varphi_1] = \top) \\ \perp & \text{otherwise.} \end{cases}$$

10 *Remark 1* It is to be noted that the semantics of LTL, FLTL, and LTL₃ are not
 11 interchangeable as is demonstrated below:

12 Consider a formula $\varphi = \square p$, and a finite trace $\alpha = s_0 s_1 \dots s_n$. If $p \notin s_i$ for some
 13 $i \in [0, n]$, then $[\alpha \models_3 \varphi] = \perp$, that is, the formula is permanently violated and so is
 14 the case in FLTL where, $[\alpha \models_F \varphi] = \perp$. Now, consider formula $\varphi = \diamond p$. If $p \notin s_i$ for
 15 all $i \in [0, n]$, then $[\alpha \models_3 \varphi] = ?$. This is because there exist infinite extensions to α
 16 that can satisfy or violate φ in the infinite semantics of LTL. But, this is not the case
 17 in FLTL where $[\alpha \models_F \varphi] = \perp$ as it did not observe any p in the observed finite trace.

18 **2.2 Distributed Computations**

19 We assume a loosely coupled asynchronous message passing system, consisting
 20 of n reliable (that do not fail) processes, denoted by $\mathcal{P} = \{P_1, P_2, \dots, P_n\}$,
 21 without any shared memory or global clock. Channels are assumed to be FIFO,
 22 and lossless. In our model, each local state change is considered an event, and
 23 every message activity (send or receive) is also represented by a new event.
 24 Message transmission does not change the local state of processes and the
 25 content of a message is immaterial to our purposes. We will need to refer to
 26 some global clock which acts as a ‘*real*’ timekeeper. It is to be understood,
 27 however, that this global clock is a theoretical object used in definitions, and
 28 is *not* available to the processes.

29 We make a practical assumption, known as *partial synchrony*. The *local*
 30 *clock* (or time) of a process P_i , where $i \in [1, n]$, can be represented as an
 31 increasing function $c_i : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$, where $c_i(\chi)$ is the value of the local clock
 32 at global time χ . Then, for any two processes P_i and P_j , we have:

$$\forall \chi \in \mathbb{R}_{\geq 0}. |c_i(\chi) - c_j(\chi)| < \epsilon$$

with $\epsilon > 0$ being the maximum *clock skew*. The value ϵ is assumed to be fixed and known by the monitor in the rest of this paper. In the sequel, we make it explicit when we refer to ‘local’ or ‘global’ time. This assumption is met by using a clock synchronization algorithm, like NTP [5], to ensure bounded clock skew among all processes.

An *event* in process P_i is of the form $e_{\tau,\sigma}^i$, where σ is *logical time* (i.e., a natural number) and τ is the local time at global time χ , that is, $\tau = c_i(\chi)$. We assume that for every two events $e_{\tau,\sigma}^i$ and $e_{\tau',\sigma'}^i$, we have $(\tau < \tau') \Leftrightarrow (\sigma < \sigma')$.

Definition 2 A *distributed computation* on N processes is a tuple $(\mathcal{E}, \rightsquigarrow)$, where \mathcal{E} is a set of events partially ordered by Lamport’s *happened-before* (\rightsquigarrow) relation [12], subject to the partial synchrony assumption:

- In every process P_i , $1 \leq i \leq N$, all events are totally ordered, that is,

$$\forall \tau, \tau' \in \mathbb{R}_+. \forall \sigma, \sigma' \in \mathbb{Z}_{\geq 0}. (\sigma < \sigma') \rightarrow (e_{\tau,\sigma}^i \rightsquigarrow e_{\tau',\sigma'}^i).$$

- If e is a message send event in a process, and f is the corresponding receive event by another process, then we have $e \rightsquigarrow f$.
- For any two processes P_i and P_j , and any two events $e_{\tau,\sigma}^i, e_{\tau',\sigma'}^j \in \mathcal{E}$, if $\tau + \epsilon < \tau'$, then $e_{\tau,\sigma}^i \rightsquigarrow e_{\tau',\sigma'}^j$, where ϵ is the maximum clock skew.
- If $e \rightsquigarrow f$ and $f \rightsquigarrow g$, then $e \rightsquigarrow g$. ■

Definition 3 Given a distributed computation $(\mathcal{E}, \rightsquigarrow)$, a subset of events $C \subseteq \mathcal{E}$ is said to form a *consistent cut* iff when C contains an event e , then it contains all events that happened-before e . Formally, $\forall e \in \mathcal{E}. (e \in C) \wedge (f \rightsquigarrow e) \rightarrow f \in C$. ■

The *frontier* of a consistent cut C , denoted $\text{front}(C)$ is the set of events that happen last in the cut. $\text{front}(C)$ is a set of e_{last}^i for each $i \in [1, |\mathcal{P}|]$ and $e_{last}^i \in C$. We denote e_{last}^i as the last event in P_i such that $\forall e_{\tau,\sigma}^i \in \mathcal{E}. (e_{\tau,\sigma}^i \neq e_{last}^i) \rightarrow (e_{\tau,\sigma}^i \rightsquigarrow e_{last}^i)$.

2.3 Problem Statement

Given a distributed computation $(\mathcal{E}, \rightsquigarrow)$, a *valid* sequence of consistent cuts is of the form $C_0 C_1 C_2 \dots$, where for all $i \geq 0$, (1) C_i denotes a set of events included in the consistent cut, (2) C_i is a subset of its succeeding consistent cut, C_{i+1} , that is, $C_i \subset C_{i+1}$, and (3) C_{i+1} has one additional event compared to its preceding consistent cut C_i , that is, $|C_i| + 1 = |C_{i+1}|$. Let \mathcal{C} denote the set of all valid sequences of consistent cuts. We define the set of all traces of $(\mathcal{E}, \rightsquigarrow)$ as follows:

$$\text{Tr}(\mathcal{E}, \rightsquigarrow) = \left\{ \text{front}(C_0) \text{front}(C_1) \dots \mid C_0 C_1 C_2 \dots \in \mathcal{C} \right\}.$$

1 Now for our automata-based approach (resp. progression-based approach),
 2 the evaluation of the LTL formula φ with respect to $(\mathcal{E}, \rightsquigarrow)$ in the 3-valued
 3 semantics (resp. finite semantics) is the following:

$$[(\mathcal{E}, \rightsquigarrow) \models_3 \varphi] = \left\{ \alpha \models_3 \varphi \mid \alpha \in \text{Tr}(\mathcal{E}, \rightsquigarrow) \right\}$$

4 and

$$[(\mathcal{E}, \rightsquigarrow) \models_F \varphi] = \left\{ \alpha \models_F \varphi \mid \alpha \in \text{Tr}(\mathcal{E}, \rightsquigarrow) \right\}$$

5 respectively. This means evaluating a distributed computation with respect to
 6 a formula results in a *set* of verdicts, as a computation may involve several
 7 traces.

8 2.4 Hybrid Logical Clocks

9 A *hybrid logical clock* (HLC) [13] is a tuple (τ, σ, ω) for detecting one-way
 10 causality, where τ is the local time, σ ensures the order of send and receive
 11 events between two processes, and ω indicates causality between events. Thus,
 12 in the sequel, we denote an event by $e_{\tau, \sigma, \omega}^i$. More specifically, for a set \mathcal{E} of
 13 events:

- 14 • τ is the local clock value of events, where for any process P_i and two events
 15 $e_{\tau, \sigma, \omega}^i, e_{\tau', \sigma', \omega'}^i \in \mathcal{E}$, we have $\tau < \tau'$ iff $e_{\tau, \sigma, \omega}^i \rightsquigarrow e_{\tau', \sigma', \omega'}^i$.
- 16 • σ stipulates the logical time, where:
 - 17 – For any process P_i and any event $e_{\tau, \sigma, \omega}^i \in \mathcal{E}$, τ never exceeds σ , and their
 18 difference is bounded by ϵ (i.e. $\sigma - \tau \leq \epsilon$).
 - 19 – For any two processes P_i and P_j , and any two events $e_{\tau, \sigma, \omega}^i, e_{\tau', \sigma', \omega'}^j \in \mathcal{E}$,
 20 where event $e_{\tau, \sigma, \omega}^i$ receiving a message sent by event $e_{\tau', \sigma', \omega'}^j$, σ is updated
 21 to $\max\{\sigma, \sigma', \tau\}$. The maximum of the three values are chosen to ensure
 22 that σ remains updated with the largest τ observed so far. Observe that
 23 σ has similar behavior as τ , except the communication between processes
 24 has no impact on the value of τ for an event.
- 25 • $\omega : \mathcal{E} \rightarrow \mathbb{Z}_{\geq 0}$ is a function that maps each event in \mathcal{E} to the causality updates,
 26 where:
 - 27 – For any process P_i and a send or local event $e_{\tau, \sigma, \omega}^i \in \mathcal{E}$, if $\tau < \sigma$, then ω
 28 is incremented. Otherwise, ω is reset to 0.
 - 29 – For any two processes P_i and P_j and any two events $e_{\tau, \sigma, \omega}^i, e_{\tau', \sigma', \omega'}^j \in \mathcal{E}$,
 30 where event $e_{\tau, \sigma, \omega}^i$ receiving a message sent by event $e_{\tau', \sigma', \omega'}^j$, $\omega(e_{\tau, \sigma, \omega}^i)$ is
 31 updated based on $\max\{\sigma, \sigma', \tau\}$.
 - 32 – For any two processes P_i and P_j , and any two events $e_{\tau, \sigma, \omega}^i, e_{\tau', \sigma', \omega'}^j \in \mathcal{E}$,
 33 $(\tau = \tau') \wedge (\omega < \omega') \rightarrow e_{\tau, \sigma, \omega}^i \rightsquigarrow e_{\tau', \sigma', \omega'}^j$.

34 HLC is susceptible to faults that might be present in the system, such as
 35 missing clock-synchronization messages or processes changing their local clock




Fig. 6 HLC example.

1 values. We assume our system is free of such faults. Fig. 6 shows an HLC incor-
 2 porated partially synchronous concurrent timelines of three processes with
 3 $\varepsilon = 10$. Observe that the local times of all events in $\text{front}(C_1)$ are bounded by
 4 ε . Therefore, C_1 is a consistent cut, but C_0 and C_2 are not.

5 3 Formula Progression for LTL

6 In a synchronous system, verification on a computation can be performed in a
 7 state by state approach due to the existence of a total ordering of events [14].
 8 However, in a *partially* synchronous system, no such total ordering of events is
 9 possible. A distributed computation $(\mathcal{E}, \rightsquigarrow)$ may have different partial ordering
 10 of events dictated by different interleaving of events. Therefore, it is possible to
 11 obtain multiple verdicts on the same distributed computation $(\mathcal{E}, \rightsquigarrow)$. In order
 12 to explore these verdicts, we propose a monitoring approach based on *formula*
 13 *progression* that, if possible, *partially* evaluates a formula on the current compu-
 14 tation, and based on the verdict, provides a *rewritten* formula that is to be
 15 evaluated on the extensions of the computation. As an example, let us consider
 16 the formula to be monitored as, $\varphi = \diamond(a \rightarrow \diamond b)$. Now, if in some trace in a
 17 computation, the monitor observes a , then for the extensions of computations,
 18 it is enough to monitor the rewritten formula, $\varphi' = \diamond b$, as the final verdict is
 19 no longer dependent on the occurrence of a . We call this method of rewriting
 20 formula **Progression**, which we discuss in length in the following section.

21 **Definition 4** A *progression function* $\text{Pr} : \Sigma^* \times \Phi_{\text{LTL}} \rightarrow \Phi_{\text{LTL}}$ is one that for all finite
 22 traces $\alpha \in \Sigma^*$, infinite traces $\sigma \in \Sigma^\omega$, and formulas $\varphi \in \Phi_{\text{LTL}}$, we have: $\alpha\sigma \models \varphi$ iff
 23 and only if $\sigma \models \text{Pr}(\alpha, \varphi)$. ■

24 We emphasize that the main difference between our technique and the
 25 classic rewriting technique [7] is that, function Pr takes a finite traces as input,
 26 while the algorithm in [7] rewrite the input LTL formula in a state-by-state
 27 manner. This means that rewriting based on the fixed point representation of

temporal operators is not possible. The motivation for our approach comes from the fact the a given distributed computation is chopped into a number of *segments*, and verification of each segment is handled by an SMT query. A state by state approach would incur too many SMT queries, making it unscalable.

Remark 2 It is straightforward to see that for any $\alpha \in \Sigma^*$ and $\varphi \in \Phi$, if a progression function returns a non-trivial formula, which we denote by $\text{Pr}(\alpha, \varphi) = \varphi'$ for some $\varphi' \in \Phi_{\text{LTL}}$, then the verdict of monitoring is *unknown*.

Atomic propositions. Let $\varphi = p$ for some $p \in \text{AP}$. The verdict is provided depending upon whether or not $p \in \alpha(0)$. This is the only case where the output of Pr cannot be a rewritten formula; the possible verdicts are either **true** or **false**:

$$\text{Pr}(\alpha, \varphi) = \begin{cases} \mathbf{true} & \text{if } p \in \alpha(0) \\ \mathbf{false} & \text{if } p \notin \alpha(0) \end{cases}$$

Negation. Let $\varphi = \neg\phi$. We have $\text{Pr}(\alpha, \varphi) = \neg\text{Pr}(\alpha, \phi)$.

Disjunction. Let $\varphi = \varphi_1 \vee \varphi_2$. If either sub-formula φ_1 or φ_2 is evaluated to **false**, then the progression of φ becomes the other sub-formula φ_2 or φ_1 respectively, since that will be the only responsible sub-formula for the verdict of all future computations:

$$\text{Pr}(\alpha, \varphi) = \begin{cases} \mathbf{true} & \text{if } \text{Pr}(\alpha, \varphi_1) = \mathbf{true} \vee \text{Pr}(\alpha, \varphi_2) = \mathbf{true} \\ \mathbf{false} & \text{if } \text{Pr}(\alpha, \varphi_1) = \mathbf{false} \wedge \text{Pr}(\alpha, \varphi_2) = \mathbf{false} \\ \varphi'_2 & \text{if } \text{Pr}(\alpha, \varphi_1) = \mathbf{false} \wedge \text{Pr}(\alpha, \varphi_2) = \varphi'_2 \\ \varphi'_1 & \text{if } \text{Pr}(\alpha, \varphi_2) = \mathbf{false} \wedge \text{Pr}(\alpha, \varphi_1) = \varphi'_1 \\ \varphi'_1 \vee \varphi'_2 & \text{if } \text{Pr}(\alpha, \varphi_1) = \varphi'_1 \wedge \text{Pr}(\alpha, \varphi_2) = \varphi'_2 \end{cases}$$

Next operator. Let $\varphi = \bigcirc\phi$. The verdicts **true**, **false** and ϕ' can only be reached if α^1 is not an empty trace, that is, $|\alpha^1| \neq 0$. Otherwise, if we are at the last event in the trace, then the progression of φ becomes ϕ ; implying ϕ must hold at the beginning of the future extension:

$$\text{Pr}(\alpha, \varphi) = \begin{cases} \mathbf{true} & \text{if } \text{Pr}(\alpha^1, \phi) = \mathbf{true} \wedge |\alpha^1| \neq 0 \\ \mathbf{false} & \text{if } \text{Pr}(\alpha^1, \phi) = \mathbf{false} \wedge |\alpha^1| \neq 0 \\ \phi' & \text{if } \text{Pr}(\alpha^1, \phi) = \phi' \wedge |\alpha^1| \neq 0 \\ \phi & \text{if } |\alpha^1| = 0 \end{cases}$$

Always and eventually operators. Progression in the temporal operator ‘always’, \square (resp. ‘eventually’, \diamond) may yield **false** (resp. **true**) or remain

1 unchanged:

$$\Pr(\alpha, \varphi) = \begin{cases} \mathbf{false} & \text{if } [\alpha \models_F \varphi] = \perp \\ \Box \phi & \text{if otherwise} \end{cases}$$

$$\Pr(\alpha, \varphi) = \begin{cases} \mathbf{true} & \text{if } [\alpha \models_F \varphi] = \top \\ \Diamond \phi & \text{if otherwise} \end{cases}$$

3 Note that the semantics of FLTL is not frequently used, due to LTL₃ being
4 generally more expressive, as shown in [15]. However, LTL₃ cannot be used to
5 construct the progression rules. To be more precise, the ‘?’ (*unknown*) verdict
6 in LTL₃ semantics would raise additional and unnecessary complications in the
7 progression rules, as this verdict does not provide any additional information
8 as far as our progression-based approach is concerned. In fact, if progression
9 results in a formula, it represents the ‘?’ verdict in LTL₃. Therefore, we use FLTL
10 for specifying the progression rules without any loss of generality as shown
11 later in the proof of Lemma 1.

12 **Until operator.** Let $\varphi = \varphi_1 \mathcal{U} \varphi_2$. Recall that $\varphi_1 \mathcal{U} \varphi_2 = \varphi_2 \vee (\varphi_1 \wedge$
13 $\bigcirc(\varphi_1 \mathcal{U} \varphi_2))$. We divide the \mathcal{U} formula into two parts, one with globally ($\Box \varphi_1$)
14 and the other eventuality ($\Diamond \varphi_2$). These sub-formulas are evaluated separately
15 and the verdict of each of them is used to define the progression for the \mathcal{U}
16 operator. However, for the case when both φ_1 and φ_2 occur in the same
17 computation, we cannot come to a verdict without considering the order of
18 satisfaction of these sub-formulas. That is, on a given finite trace α , if φ_2 holds
19 in $\alpha(i)$ (denoted $\Diamond_i \varphi_2$) and φ_1 holds throughout in all states from $\alpha(0)$ to
20 $\alpha(i-1)$ (denoted $\Box_{i-1} \varphi_1$), then the progression of φ becomes **true**. If this is
21 not the case, and $\Box \varphi_1$ does not hold in α , the progression of φ becomes **false**,
22 since this signifies a break from the streak of φ_1 required for φ to hold. If it is
23 neither of the above two cases, and the evaluated verdict of $\Diamond \Pr(\alpha, \varphi_2)$ is \top ,
24 then this represents a case where we do not have enough information about
25 φ_1 to evaluate $\varphi_1 \mathcal{U} \varphi_2$. Thus, making the progression solely dependant on φ_1 .
26 The progression of φ remains unchanged if φ_1 holds throughout α , but φ_2 does
27 not hold anywhere:

$$\Pr(\alpha, \varphi) = \begin{cases} \mathbf{true} & \text{if } \exists i \in [0, |\alpha| - 1] . [\alpha \models_F \Diamond_i \Pr(\alpha, \varphi_2)] = \top \\ & \wedge [\alpha \models_F \Box_{i-1} \Pr(\alpha, \varphi_1)] = \top \\ \mathbf{false} & \text{if } [\alpha \models_F \Box \Pr(\alpha, \varphi_1)] = \perp \\ & \wedge \text{not the first case} \\ \Pr(\alpha, \varphi_1) & \text{if } [\alpha \models_F \Diamond \Pr(\alpha, \varphi_2)] = \top \\ & \wedge \text{not the second case} \\ \Pr(\alpha, \varphi_1) \mathcal{U} \Pr(\alpha, \varphi_2) & \text{if } [\alpha \models_F \Box \Pr(\alpha, \varphi_1)] = \top \\ & \wedge [\alpha \models_F \Diamond \Pr(\alpha, \varphi_2)] = \perp \end{cases}$$

28 **Example.** Consider the formula, $\varphi = \Diamond r \rightarrow (\neg p \mathcal{U} q)$ with sub-formulas
29 $\varphi_s = \{\Diamond r, q, \Diamond q, \Box p\}$, according to our progression rules. Consider the




Fig. 7 Progression example.

1 trace in Fig. 7 divided into three segments. In the first segment α , neither
 2 p , q nor r are present, and as far as the laws of the progression function
 3 defined above, φ remains unchanged for the next segment; i.e., $\text{Pr}(\alpha, \varphi) = \varphi$.
 4 In the second segment α' , proposition r is observed, this satisfies sub-formula
 5 $\diamond r$ the progressed formula becomes $\neg p \mathcal{U} q$; i.e., $\text{Pr}(\alpha', \varphi) = \neg p \mathcal{U} q$. In the
 6 next segment α'' , proposition q occurs before p . This falls under the first case
 7 of the until progression operator. Since q happens after a streak of $\neg p$, we
 8 arrive at the verdict **true**; i.e., $\text{Pr}(\alpha'', \neg p \mathcal{U} q) = \text{true}$. Put it another way,
 9 $\text{Pr}(\alpha\alpha'\alpha'', \varphi) = \text{true}$.

10 **Lemma 1** Given an LTL formula φ , and finite traces $\alpha, \sigma \in \Sigma^*$, trace $\alpha\sigma$ satisfies φ
 11 if and only if σ satisfies $\text{Pr}(\alpha, \varphi)$, as defined in above. Formally,

$$[\alpha\sigma \models_F \varphi] \iff [\sigma \models_F \text{Pr}(\alpha, \varphi)]$$

12 *Proof* We distinguish the following cases:

13
 14 **Case 1:** First, we consider the base case of this proof, where the formula is an
 15 atomic proposition, that is, $\varphi = p$.

16
 17 (\Rightarrow) Let us first consider that p is observed on the first state of $\alpha\sigma$. This implies,
 18 $[\alpha\sigma \models_F \varphi]$ yields **true**, and $\text{Pr}(\alpha, \varphi)$ yields \top . Therefore, $[\sigma \models_F \text{Pr}(\alpha, \varphi)]$ must also
 19 yield **true**.

20
 21 Now, let us consider that p is not observed on the first state of $\alpha\sigma$. This implies,
 22 $[\alpha\sigma \models_F \varphi]$ yields **false**, and $\text{Pr}(\alpha, \varphi)$ yields \perp . Therefore, $[\sigma \models_F \text{Pr}(\alpha, \varphi)]$ must
 23 also yield **false**.

24
 25 (\Leftarrow) Let us first consider that $[\sigma \models_F \text{Pr}(\alpha, \varphi)]$ yields **true**. This implies, $\text{Pr}(\alpha, \varphi)$
 26 yields \top , and $[\alpha\sigma \models_F \varphi]$ yields **true**. Therefore, p must have been observed on the
 27 first state of $\alpha\sigma$.

28
 29 Now, let us consider that $[\sigma \models_F \text{Pr}(\alpha, \varphi)]$ yields **false**. This implies, $\text{Pr}(\alpha, \varphi)$
 30 yields \perp , and $[\alpha\sigma \models_F \varphi]$ yields **false**. Therefore, p must not have been observed on
 31 the first state of $\alpha\sigma$.

32
 33 **Case 2:** Assume that the proof has been established for the case when the formula
 34 is $\varphi = \phi$. Now, we consider the case where the formula is $\varphi = \neg\phi$.

35
 36 We can say $[\alpha\sigma \models_F \neg\phi]$ is equivalent to $\neg[\alpha\sigma \models_F \phi]$ according to the
 37 finite-trace semantics of LTL. We can also say $[\sigma \models_F \text{Pr}(\alpha, \neg\phi)]$ is equivalent to

1 $[\sigma \models_F \neg \text{Pr}(\alpha, \phi)]$ since $\text{Pr}(\alpha, \neg\phi) = \neg \text{Pr}(\alpha, \phi)$ is defined as a progression rule.
 2 Furthermore, $[\sigma \models_F \neg \text{Pr}(\alpha, \phi)]$ is equivalent to $\neg[\sigma \models_F \text{Pr}(\alpha, \phi)]$ according to the
 3 finite-trace semantics of LTL.

4
 5 Based on our assumption, the proof has already been established for
 6 $[\alpha\sigma \models_F \phi] \iff [\sigma \models_F \text{Pr}(\alpha, \phi)]$. Therefore, $\neg[\alpha\sigma \models_F \phi] \iff \neg[\sigma \models_F \text{Pr}(\alpha, \phi)]$,
 7 and by extension, $[\alpha\sigma \models_F \neg\phi] \iff [\sigma \models_F \text{Pr}(\alpha, \neg\phi)]$

8
 9 **Case 3:** Assume that the proof has been established for the case when the formula
 10 is $\varphi = \phi$. Now, we consider the case where the formula is $\varphi = \bigcirc\phi$.

11
 12 Let us first consider the case where the length of the trace α is 1, that is, $|\alpha| = 1$
 13 and $|\alpha^1| = 0$. In this particular case, $[\alpha\sigma \models_F \bigcirc\phi]$ is equivalent to $[\sigma \models_F \phi]$.
 14 Furthermore, $\text{Pr}(\alpha, \bigcirc\phi) = \phi$; which implies, $[\sigma \models_F \text{Pr}(\alpha, \bigcirc\phi)]$ is equivalent to
 15 $[\sigma \models_F \phi]$. Therefore, $[\alpha\sigma \models_F \bigcirc\phi] \iff [\sigma \models_F \text{Pr}(\alpha, \bigcirc\phi)]$.

16
 17 Now, let us consider the case where the length of the trace α is longer than 1,
 18 that is, $|\alpha| \geq 1$ and $|\alpha^1| \geq 1$. In this case, $[\alpha\sigma \models_F \bigcirc\phi]$ is equivalent to $[\alpha^1\sigma \models_F \phi]$,
 19 and $[\sigma \models_F \text{Pr}(\alpha, \bigcirc\phi)]$ is equivalent to $[\sigma \models_F \text{Pr}(\alpha^1, \phi)]$.

20
 21 Based on our assumption, the proof has already been established for $[\alpha^1\sigma \models_F$
 22 $\phi] \iff [\sigma \models_F \text{Pr}(\alpha^1, \phi)]$. Therefore, $[\alpha\sigma \models_F \bigcirc\phi] \iff [\sigma \models_F \text{Pr}(\alpha, \bigcirc\phi)]$.

23
 24 **Case 4:** Assume that the proof has been established for the cases when the formulas
 25 are $\varphi = \varphi_1$ and $\varphi = \varphi_2$. Now, we consider the case where the formula is $\varphi = \varphi_1 \vee \varphi_2$.

26
 27 Based on our assumption, the proof has already been established for
 28 $[\alpha\sigma \models_F \varphi_1] \iff [\sigma \models_F \text{Pr}(\alpha, \varphi_1)]$ and $[\alpha\sigma \models_F \varphi_2] \iff [\sigma \models_F \text{Pr}(\alpha, \varphi_2)]$.
 29 Therefore, we can derive the following:

$$\begin{aligned} [\alpha\sigma \models_F (\varphi_1 \vee \varphi_2)] &\iff [\alpha\sigma \models_F \varphi_1] \vee [\alpha\sigma \models_F \varphi_2] \\ &\iff [\sigma \models_F \text{Pr}(\alpha, \varphi_1)] \vee [\sigma \models_F \text{Pr}(\alpha, \varphi_2)] \\ &\iff [\sigma \models_F \text{Pr}(\varphi_1 \vee \varphi_2)]. \end{aligned}$$

30
 31 **Case 5:** Now, we consider the case where the formula is $\varphi = \varphi_1 \mathcal{U} \varphi_2$.
 We prove this by induction on the length of α :

Base Case: $|\alpha| = 0$.

$$\begin{aligned} [\alpha\sigma \models_F \varphi] &\iff [\sigma \models_F \text{Pr}(\alpha, \varphi)] \\ &\iff [\sigma \models_F \varphi] \end{aligned}$$

32
 31 *Hypothesis Step:* $|\alpha| = k$.

$$[\alpha\sigma \models_F \varphi_1 \mathcal{U} \varphi_2]$$

$$\begin{aligned}
&\iff [\alpha\sigma \models_F (\varphi_2 \vee (\varphi_1 \wedge \mathcal{O}(\varphi_1 \mathcal{U} \varphi_2)))] \\
&\iff [\alpha\sigma \models_F \varphi_2] \vee [\alpha\sigma \models_F (\varphi_1 \wedge \mathcal{O}(\varphi_1 \mathcal{U} \varphi_2))] \\
&\iff [\alpha\sigma \models_F \varphi_2] \vee ([\alpha\sigma \models_F \varphi_1] \wedge [\alpha^1\sigma \models_F \varphi_1 \mathcal{U} \varphi_2]) \\
&\iff [\alpha\sigma \models_F \varphi_2] \vee ([\alpha\sigma \models_F \varphi_1] \wedge [\alpha^1\sigma \models_F (\varphi_2 \vee (\varphi_1 \wedge \mathcal{O}(\varphi_1 \mathcal{U} \varphi_2)))] \\
&\iff [\alpha\sigma \models_F \varphi_2] \vee ([\alpha\sigma \models_F \varphi_1] \wedge [\alpha^1\sigma \models_F \varphi_2]) \vee \dots \vee ([\alpha\sigma \models_F \varphi_1] \wedge \\
&\quad [\alpha^1\sigma \models_F \varphi_1] \wedge \dots \wedge [\alpha^{k-2}\sigma \models_F \varphi_1] \wedge [(\alpha^{k-1}\sigma \models_F \varphi_2)] \vee \\
&\quad ([\alpha\sigma \models_F \varphi_1] \wedge \dots \wedge [\alpha^{k-1}\sigma \models_F \varphi_1] \wedge [\alpha^k\sigma \models_F \varphi_1 \mathcal{U} \varphi_2]) \\
&\iff [\alpha\sigma \models_F \varphi_2] \vee ([\alpha\sigma \models_F \varphi_1] \wedge [\alpha^1\sigma \models_F \varphi_2]) \vee \dots \vee ([\alpha\sigma \models_F \varphi_1] \wedge \\
&\quad \dots \wedge [\alpha^{k-1}\sigma \models_F \varphi_1] \wedge [\sigma \models_F \varphi_1 \mathcal{U} \varphi_2])
\end{aligned}$$

Inductive Step: $|\alpha| = k + 1$

Trivially expanded from the above expansion.

$$\begin{aligned}
[\alpha\sigma \models_F \varphi_1 \mathcal{U} \varphi_2] &\iff [\alpha\sigma \models_F \varphi_2] \vee ([\alpha\sigma \models_F \varphi_1] \wedge [\alpha^1\sigma \models_F \varphi_2]) \vee \dots \vee \\
&\quad ([\alpha\sigma \models_F \varphi_1] \wedge \dots \wedge [\alpha^{k-1}\sigma \models_F \varphi_1] \wedge [\alpha^k\sigma \models_F \varphi_1] \wedge [\sigma \models_F \varphi_1 \mathcal{U} \varphi_2])
\end{aligned}$$

¹ Now, in order for $[\alpha\sigma \models_F \varphi_1 \mathcal{U} \varphi_2]$ to yield **true**, there must be a $k \geq 1$ such that
² $[\alpha\sigma \models_F \varphi_1] \wedge \dots \wedge [\alpha^{k-1}\sigma \models_F \varphi_1] \wedge [\alpha^k\sigma \models_F \varphi_2]$, that is

³

$$\begin{aligned}
[\alpha\sigma \models_F \varphi_1 \mathcal{U} \varphi_2] &\iff \exists k \geq 1 . [\alpha^0\sigma \models_F \varphi_1] \wedge \dots \wedge [\alpha^{k-1}\sigma \models_F \varphi_1] \wedge \\
&\quad [\alpha^k\sigma \models_F \varphi_2] \\
&\iff \exists k \geq 1 . [\alpha\sigma \models_F \blacklozenge_k \varphi_2] \wedge [\alpha\sigma \models_F \blackbox_{k-1} \varphi_1]
\end{aligned}$$

⁴ Note that the above recursive definition of Until allows us to evaluate any until
⁵ formula, and by extension, eventually ($\blacklozenge \varphi = \top \mathcal{U} \varphi$) and always formulas. Therefore,
⁶ we can evaluate any sub-formula using this fixed point representation of until.

⁷ ■

⁸ 4 SMT-based Solution

⁹ In this section, we elaborate on our solution for distributed monitoring
¹⁰ using the two monitoring techniques mentioned before: (1) automata-based
¹¹ approach, and (2) progression-based approach.

¹² 4.1 Overall Idea

¹³ **Automata-based approach.** Recall from Section 1 (Fig. 4) that monitor-
¹⁴ ing a distributed computation may result in multiple verdicts depending upon
¹⁵ different ordering of events. In other words, given a distributed computation

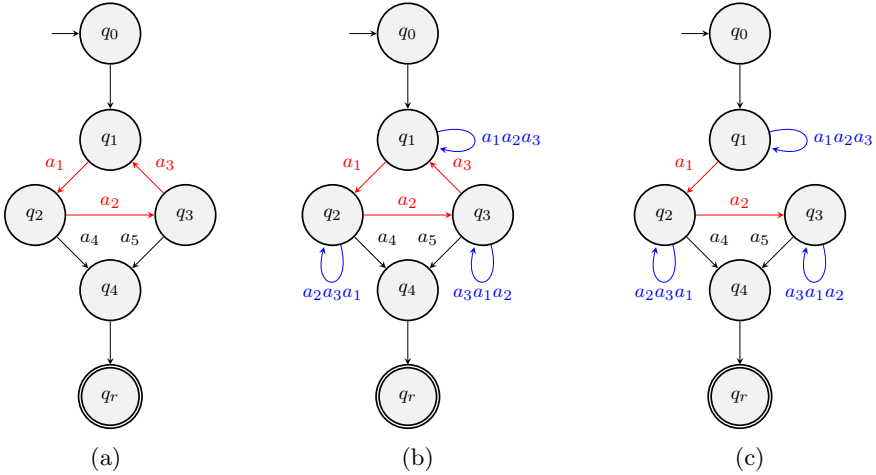


Fig. 8 Removing non-loop cycles in an LTL₃ Monitor.

1 ($\mathcal{E}, \rightsquigarrow$) and an LTL formula φ , different ordering of events may reach different
 2 states in the monitor automaton $\mathcal{M}_\varphi = (\Sigma, Q, q_0, \delta, \lambda)$ (as defined in Definition
 3 1). In order to ensure that all possible verdicts are explored, we generate
 4 an SMT instance for (1) the distributed computation $(\mathcal{E}, \rightsquigarrow)$, and (2) each
 5 possible path in the LTL₃ monitor. Thus, the corresponding decision problem
 6 is the following: given $(\mathcal{E}, \rightsquigarrow)$ and a monitor path $q_0 q_1 \cdots q_m$ in an LTL₃ monitor,
 7 can $(\mathcal{E}, \rightsquigarrow)$ reach q_m ? If the SMT instance is satisfiable, then $\lambda(q_m)$ is a
 8 possible verdict. For example, for the monitor in Fig. 5, we consider two paths
 9 $q_0^* q_{\perp}$ and $q_0^* q_{\top}$ (and, hence, two SMT instances). Thus, if both instances turn
 10 out to be unsatisfiable, then the resulting monitor state is q_0 , where $\lambda(q_0) = ?$.

11 We note that LTL₃ monitors may contain non-self-loop cycles. In order
 12 to simplify the SMT instance creation process (for each possible path in the
 13 LTL₃ monitor), we collapse each non-self-loop cycle into one state with a self-
 14 loop labeled by the sequence of events in the cycle using Algorithm 1. As
 15 an example, in Fig. 8, Algorithm 1 first takes an LTL₃ monitor (Fig. 8a) and
 16 adds the necessary self-loops (Fig. 8b). Then it eliminates all non-self-loop
 17 cycles by removing transitions from states with higher identifiers to states with
 18 lower identifiers in cycles (Fig. 8c). The non-deterministic nature of the final
 19 automata ensure that all the transitions and the accepting language of the
 20 automata are preserved.

21 **Lemma 2** Let $\mathcal{M}_\varphi = (\Sigma, Q, q_0, \delta, \lambda)$ be the monitor automaton for LTL formula, φ ,
 22 and $\mathcal{M}'_\varphi = (\Sigma, Q, q_0, \delta', \lambda)$ be the monitor automaton with no non-self loop cycles,
 23 obtained from applying Algorithm 1 on \mathcal{M}_φ . Given a finite trace, $\alpha = a_1 a_2 \cdots a_n$
 24 and a initial state, $q \in Q$, we prove that $\lambda(\delta(q, \alpha)) = \lambda(\delta'(q, \alpha))$.

Algorithm 1: Non-Self Loop Cycle Removal Algorithm

Data: $\mathcal{M}_\varphi = (\Sigma, Q, q_0, \delta, \lambda)$
Result: $\mathcal{M}'_\varphi = (\Sigma, Q, q_0, \delta', \lambda)$

- 1 Let CP be the set of all possible paths containing cycles
- 2 $\delta' \leftarrow \delta$
- 3 **foreach** $q \in Q$ **do**
- 4 **foreach** $q \xrightarrow{s_m} \dots \xrightarrow{s_n} q \in CP$ **do**
- 5 $\delta'(q, s_m \dots s_n) \leftarrow q$
- 6 **foreach** $q_m \xrightarrow{s} q_n \in \{q_i \xrightarrow{s_k} q_j \mid q \xrightarrow{s_m} \dots q_i \xrightarrow{s_k} q_j \dots \xrightarrow{s_n} q \in CP\}$ **do**
- 7 **if** $m > n$ **then**
- 8 $\delta'(q_m, s) \leftarrow \emptyset$
- 9 **return** \mathcal{M}_φ

1 *Proof* We distinguish the following cases:

2

3 **Case 1** (\Rightarrow):

4

5 First we show, $\lambda(\delta(q, \alpha)) \rightarrow \lambda(\delta'(q, \alpha))$, that is, $\forall \alpha, \forall q \in Q . \lambda(\delta(q, \alpha)) \Longrightarrow$
6 $\lambda(\delta'(q, \alpha))$ Let $\alpha = a_1 a_2 \dots a_n$, where $\forall i \in [1, n]. a_i \in \Sigma$. Algorithm 1 removes
7 non-self loop cycles by removing a transition such that the corresponding transi-
8 tion of $\delta(q, a_i)$, $\delta'(q, a_i)$, where $i \in [1, m]$ does not exist. This is such that
9 $\exists k \in [1, i] . q' \xrightarrow{a_{i-k}} \dots q \xrightarrow{a_i} q'$. This transition is same as $\delta'(q', a_{i-k} \dots a_i) = q'$
10 which was one of the added self-loops. The rest of the transitions are maintained
11 such that $\delta(q, a_i) = \delta'(q, a_i)$, where $q \in Q$ and $i \in [1, m]$.

12

13 **Case 2** (\Leftarrow):

14

15 Now, we show, $\lambda(\delta'(q, \alpha)) \rightarrow \lambda(\delta(q, \alpha))$, that is, $\forall \alpha, \forall q \in Q . \lambda(\delta'(q, \alpha)) \Longrightarrow$
16 $\lambda(\delta(q, \alpha))$ Let $\alpha = a_1 a_1 \dots a_n$, where $\forall i \in [1, n]. a_i \in \Sigma$. A self-loop in \mathcal{M}'_φ can be
17 represented by $\exists i \in [1, n], \exists k \in [1, n-i] . \delta'(q, a_i a_{i+1} \dots a_{i+k}) = q$. In another words,
18 there exists a path $q \xrightarrow{a_i} q' \xrightarrow{a_{i+1}} \dots \xrightarrow{a_{i+k}} q$ in \mathcal{M}_φ . The rest of the non-self loop
19 transitions are the same, such that $\delta'(q, a_i) = \delta(q, a_i)$, where $q \in Q$ and $i \in [1, m]$.

20

21 Thus, $\lambda(\delta(q, \alpha)) = \lambda(\delta'(q, \alpha))$ ■

22 **Progression-based approach.** In a synchronous system, verification on a
23 computation can be performed in a state by state approach due to the exist-
24 tence of a total ordering of events [14]. However, in a *partially* synchronous
25 system, no such ordering of events is possible. A distributed computation
26 $(\mathcal{E}, \rightsquigarrow)$ may have different ordering of events dictated by different interleavings
27 of events. Therefore, it is possible to obtain multiple verdicts on the same dis-
28 tributed computation $(\mathcal{E}, \rightsquigarrow)$. In order to explore these verdicts, we propose a
29 monitoring approach based on *formula progression* that, if possible, *partially*
30 evaluates a formula on the current computation, and based on the verdict,

provides a *rewritten* formula that is to be evaluated on the extensions of the computation. As an example, let us consider the formula to be monitored as, $\varphi = \diamond(a \rightarrow \diamond b)$. Now, if in some trace in a computation, the monitor observes a , then for the extensions of computations, it is enough to monitor the rewritten formula, $\varphi' = \diamond b$, as the final verdict is no longer dependent on the occurrence of a . We call this method of rewriting formula **Progression**, which we discuss in length later on. In the next two subsections, we present the SMT entities and constraints with respect to *one* monitor path and a distributed computation.

4.2 SMT Entities

SMT entities represent the sub-formulas of an LTL formula and a distributed computation. After the verdicts from all the sub-formulas are generated, we construct our rewritten formula by attaching the said verdicts to their corresponding parent formulas in the parse tree and then performing an in-order traversal starting from the root of the parse tree. At the end of the traversal, the resulting formula is, in fact, the progression for the next computation. We now introduce the entities that represent a path in an LTL₃ monitor $\mathcal{M}_\varphi = (\Sigma, Q, q_0, \delta, \lambda)$ for LTL formula φ and distributed computation $(\mathcal{E}, \rightsquigarrow)$. It should be noted that the SMT entities in this subsection are used in both the automata-based and the progression-based approaches.

Monitor automaton. Let $q_0 \xrightarrow{s_0} q_1 \xrightarrow{s_1} \dots (q_j \xrightarrow{s_j} q_j)^* \dots \xrightarrow{s_{m-1}} q_m$ be a path of monitor \mathcal{M}_φ , which may or may not include a self-loop. We include a non-negative integer variable k_i for each transition $q_i \xrightarrow{s_i} q_{i+1}$, where $i \in [0, m-1]$ and $s_i \in \Sigma$. This is also true for the self-loop $q_j \xrightarrow{s_j} q_j$, for which we include a non-negative integer k_j .

Distributed computation. In our SMT encoding, the set of events, \mathcal{E} are represented by a bit vector, where each bit corresponds to an individual event in the distributed computation, $(\mathcal{E}, \rightsquigarrow)$. We conduct a *pre-processing* of the distributed computation, during which we create an $\mathcal{E} \times \mathcal{E}$ matrix, **hbSet** to incorporate the additional happen-before relations obtained by the clock-synchronization algorithm. Afterwards, we populate the **hbSet** with 0's and 1's, such that **hbSet** $[i][j] = 1$ if $\mathcal{E}[i] \rightsquigarrow \mathcal{E}[j]$, and **hbSet** $[i][j] = 0$ otherwise. We introduce a function $\mu : \mathcal{E} \times \text{AP} \rightarrow \{\mathbf{true}, \mathbf{false}\}$ in order to establish a relation between each event and the atomic propositions in it. In the event that other variables or constants are used in defining the predicates (e.g. $x_1 + x_2 \geq 2$), μ is constructed accordingly. Finally, we introduce an uninterpreted function $\rho : \mathbb{Z}_{\geq 0} \rightarrow 2^{\mathcal{E}}$ that identifies a sequence of consistent cuts from $\{\}$ to $\{\mathcal{E}\}$ for reaching a verdict, while satisfying a number of given constraints explained in 4.3.

4.3 SMT Constraints

Once we define the necessary SMT entities, we move onto the SMT constraints. We first define the common SMT constraints for consistent cuts that are enforced on both the automata-based and the progression-based approaches. Afterwards we define the SMT constraints that are more dependant on the methodology.

Consistent cut constraints over ρ . In order to ensure that the uninterpreted function ρ identifies a sequence of consistent cuts, we enforce certain consistent cut constraints. The first constraint enforces that each element in the range of ρ is in fact a consistent cut:

$$\forall i \in [0, m]. \forall e, e' \in \mathcal{E}. \left((e' \rightsquigarrow e) \wedge (e \in \rho(i)) \right) \rightarrow \left(e' \in \rho(i) \right)$$

Next, we enforce that the sequence of consistent cuts identified by ρ start from an empty set of events, and each successor cut of the sequence contains one more new event than its predecessor.

$$\forall i \in [0, m]. |\rho(i+1)| = |\rho(i)| + 1$$

Finally, we ensure that each successive consistent cut is immediately reachable in $(\mathcal{E}, \rightsquigarrow)$ by enforcing a subset relation:

$$\forall i \in [0, m]. \rho(i) \subseteq \rho(i+1)$$

Once a sequence of consistent cuts have been generated, we check if the sequence satisfies the specification. This is done using (1) progression-based approach, where the LTL formula is represented by a SMT constrain and (2) LTL₃ automata-based approach, where a path on the automata is represented as an SMT constraint. This is repeated for all sub-formulas of the original LTL formula and all paths in the LTL₃ automata respectively as discussed below.

Constraints for LTL₃ automata over ρ . These constraints are responsible for generating a valid sequence of consistent cuts given a distributed computation $(\mathcal{E}, \rightsquigarrow)$ that runs on monitor path $q_1 \xrightarrow{s_1} q_2 \cdots q_j^* \cdots \xrightarrow{s_{m-1}} q_m$. We begin with interpreting $\rho(k_m)$ by requiring that running $(\mathcal{E}, \rightsquigarrow)$ ends in monitor state q_m . The corresponding SMT constraint is:

$$\mu(\text{front}(\rho(k_m)), s_{m-1})$$

For every monitor state q_i , where $i \in [0, m-1]$, if q_i does not have a self-loop, the corresponding SMT constraint is:

$$\mu(\text{front}(\rho(k_{i+1} - 1)), s_i) \wedge (k_i = k_{i+1} - 1)$$

1 For every monitor state q_j , where $j \in [0, m - 1]$, suppose q_j has a self-loop
 2 (recall that a cycle of r transitions in the monitor automaton is collapsed into
 3 a self-loop labeled by a sequence of r letters). Let us imagine that this self-
 4 loop executed z number of times for some $z \geq 0$. Furthermore, we denote the
 5 sequence of letters in the self-loop as $s_{j_1} s_{j_2} \cdots s_{j_r}$. The corresponding SMT
 6 constraint is:

$$\bigwedge_{i=1}^z \bigwedge_{n=1}^r \mu(\text{front}(\rho(k_j + r(i-1) + n)), s_{j_n})$$

7 Again, since z is a free variable in the above constraint, the solver will identify
 8 some value $z \geq 0$ which is exactly what we need. To ensure that the domain
 9 of ρ starts from the empty consistent cut (i.e., $\rho(0) = \emptyset$), we add:

$$k_0 = 0.$$

10 Finally, let C denote the conjunction of all the above constraints. Recall that
 11 this conjunction is with respect to only one monitor path from q_0 to q_m . Since
 12 there may be multiple paths in the monitor automaton that can reach q_m from
 13 q_0 , we replicate the above constraints for each such path. Suppose there are n
 14 such paths and let C_1, C_2, \dots, C_n be the corresponding SMT constraints for
 15 these n paths. We include the following constraint:

$$C_1 \vee C_2 \vee C_3 \vee \cdots \vee C_n$$

16 This means that if the SMT instance is satisfiable, then computation $(\mathcal{E}, \rightsquigarrow)$
 17 can reach monitor state q_m from q_0 .

18 **Constraints for LTL progression over ρ .** Given a distributed system
 19 $(\mathcal{E}, \rightsquigarrow)$, the aforementioned constraints may generate a valid sequence of con-
 20 sistent cuts that may yield different verdicts based on the ordering of the
 21 concurrent events. Therefore, in order to avoid false positives, all possible out-
 22 comes are explored when evaluating an LTL formula φ on $(\mathcal{E}, \rightsquigarrow)$. We achieve
 23 this by checking for both satisfaction and violation in the sequence of consis-
 24 tent cuts $C_0 C_1 C_2 \cdots C_m$ interpreted by the uninterpreted function $\rho(m)$. Note
 25 that monitoring any LTL formula using our progression rules will result in mon-
 26 itoring sub-formulas with only atomic propositions, globally and eventually
 27 temporal operators:

$$\begin{array}{ll} \varphi = p & \text{front}(\rho_i) \models p, \text{ for } p \in \text{AP} \text{ (satisfaction, i.e., } \top) \\ \varphi = \Box \phi & \exists i \in [0, m]. \text{ front}(\rho_i) \not\models \phi \text{ (violation, i.e., } \perp) \\ \varphi = \Diamond \phi & \exists i \in [0, m]. \text{ front}(\rho_i) \models \phi \text{ (satisfaction, i.e., } \top) \end{array}$$

28 Opposite cases result in a rewritten formula that will progress to the next
 29 segment. In general, the verdict for any LTL formula will be derived using our
 30 progression rules in Section 3.

5 Optimization

5.1 Segmentation of Distributed Computation

RV is known to be an NP-complete problem in the number of processes in a distributed setting [16]. The complexity exhibits even more exponential blowup during verifying formulas with nested temporal operators. In order to cope with this complexity, we divide our computation into smaller *segments*, $(seg_1, \rightsquigarrow) \dots (seg_{l/g}, \rightsquigarrow)$ to create smaller, albeit more SMT problems. Given a distributed computation $(\mathcal{E}, \rightsquigarrow)$ of length l , we divide it into $\frac{l}{g}$ smaller segments length g . The set of events in segment j , where $j \in [1, \frac{l}{g}]$, is the following:

$$seg_j = \left\{ e_{\tau, \sigma, \omega}^n \mid \sigma \in [\max\{0, (j-1) \times g - \epsilon\}, j \times g] \wedge n \in [1, |\mathcal{P}|] \right\}$$

Note that each segment (barring seg_0) has to be constructed starting at ϵ time units before the previous segments ending point. This creates an overlap of ϵ time units between each pair of adjacent segments. Doing so ensures that no pair of possible concurrent become non-concurrent due to the splits caused by segmentation. Therefore, dividing the actual computation into segments does not have any effect on the final verdict of the said computation. We also use parallelization (similar to the one discussed in [8]) to make our algorithm perform faster, while utilizing most of the computation power modern processors are capable of handling.

Lemma 3 A distributed computation, $(\mathcal{E}, \rightsquigarrow)$, of length l satisfies an LTL formula, φ , if and only if the distributed computation, $(\mathcal{E}, \rightsquigarrow)$, is divided into $\frac{l}{g}$ segments of length g satisfies φ using the automata-based approach. That is,

$$[(\mathcal{E}, \rightsquigarrow) \models_3 \varphi] \iff [(seg_1.seg_2 \dots .seg_{\frac{l}{g}}, \rightsquigarrow) \models_3 \varphi]$$

Proof Let us assume $[(\mathcal{E}, \rightsquigarrow) \models_3 \varphi] \neq [(seg_1.seg_2 \dots .seg_{\frac{l}{g}}, \rightsquigarrow) \models_3 \varphi]$, that is, $\{\alpha \models_3 \varphi \mid \alpha \in \text{Tr}(\mathcal{E}, \rightsquigarrow)\} \neq \{\alpha \models_3 \varphi \mid \alpha \in \text{Tr}(seg_1.seg_2 \dots .seg_{\frac{l}{g}}, \rightsquigarrow)\}$ (Recall Section 2.3).

(\Rightarrow) Let C_k be a consistent cut such that C_k is in $\text{Tr}(\mathcal{E}, \rightsquigarrow)$, but not in $\text{Tr}(seg_1.seg_2 \dots .seg_{\frac{l}{g}}, \rightsquigarrow)$ for some $k \in [0, |\mathcal{E}|]$. This implies that the frontier of C_k , $\text{front}(C_k) \not\subseteq seg_1$ and $\text{front}(C_k) \not\subseteq seg_2$ and \dots and $\text{front}(C_k) \not\subseteq seg_{\frac{l}{g}}$. However, this is not possible, as according to the segmentation construction, there must be a seg_j where $1 \leq j \leq \frac{l}{g}$ such that $\text{front}(C_k) \subseteq seg_j$. Therefore, such C_k cannot exist, and $\{\alpha \models_3 \varphi \mid \alpha \in \text{Tr}(\mathcal{E}, \rightsquigarrow)\} \subseteq \{\alpha \models_3 \varphi \mid \alpha \in \text{Tr}(seg_1.seg_2 \dots .seg_{\frac{l}{g}}, \rightsquigarrow)\}$. By extension, $[(\mathcal{E}, \rightsquigarrow) \models_3 \varphi] \Rightarrow [(seg_1.seg_2 \dots .seg_{\frac{l}{g}}, \rightsquigarrow) \models_3 \varphi]$

(\Leftarrow) Let C_k be a consistent cut such that C_k is in $\text{Tr}(seg_1.seg_2 \dots .seg_{\frac{l}{g}}, \rightsquigarrow)$, but not in $\text{Tr}(\mathcal{E}, \rightsquigarrow)$ for some $k \in [0, |\mathcal{E}|]$. This implies, $\text{front}(C_k) \subseteq seg_j$

1 and $\text{front}(C_k) \not\subseteq \mathcal{E}$ for some $j \in [1, \frac{l}{g}]$. However, this is not possible due to
 2 the fact that $\forall j \in [1, \frac{l}{g}] . \text{seg}_j \subseteq \mathcal{E}$. Therefore, such C_k cannot exist, and
 3 $\{\alpha \models_3 \varphi \mid \alpha \in \text{Tr}(\text{seg}_1.\text{seg}_2.\dots.\text{seg}_{\frac{l}{g}}, \rightsquigarrow)\} \subseteq \{\alpha \models_3 \varphi \mid \alpha \in \text{Tr}(\mathcal{E}, \rightsquigarrow)\}$. By exten-
 4 sion, $[(\text{seg}_1.\text{seg}_2.\dots.\text{seg}_{\frac{l}{g}}, \rightsquigarrow) \models_3 \varphi] \Rightarrow [(\mathcal{E}, \rightsquigarrow) \models_3 \varphi]$

5
 6 Therefore, $[(\mathcal{E}, \rightsquigarrow) \models_3 \varphi] \iff [(\text{seg}_1.\text{seg}_2.\dots.\text{seg}_{\frac{l}{g}}, \rightsquigarrow) \models_3 \varphi]$.

7 ■

8 **Lemma 4** A distributed computation $(\mathcal{E}, \rightsquigarrow)$ of length l satisfies an LTL formula
 9 φ if and only if the distributed computation, $(\mathcal{E}, \rightsquigarrow)$, is divided into $\frac{l}{g}$ segments of
 10 length g satisfies φ using the progression-based approach. That is,

$$[(\mathcal{E}, \rightsquigarrow) \models_F \varphi] \iff [(\text{seg}_1.\text{seg}_2.\dots.\text{seg}_{\frac{l}{g}}, \rightsquigarrow) \models_F \varphi]$$

11 *Proof* Using Lemma 1 and Lemma 3, we can trivially prove, $[(\mathcal{E}, \rightsquigarrow) \models_F \varphi] \iff$
 12 $[(\text{seg}_1.\text{seg}_2.\dots.\text{seg}_{\frac{l}{g}}, \rightsquigarrow) \models_F \varphi]$. ■

13 5.2 Parallelized Monitoring

14 Many cloud services use clusters of computers equipped with multiple pro-
 15 cessors and computing cores. This allows them to deal with high data rates
 16 and implement high-performance parallel/distributed applications. Monitor-
 17 ing such applications should also be able to exploit the massive infrastructure.
 18 To this end, we now discuss parallelization of our SMT-based monitoring
 19 technique.

20 Let G be a sequence of g segments $G = \text{seg}_1.\text{seg}_2.\dots.\text{seg}_g$. Our idea is to
 21 create a job queue for each available computing core, and then distribute the
 22 segments evenly across all the queues to be monitored by their respective cores
 23 independently. However, simply distributing all the segments across cores is not
 24 enough for obtaining a correct result. For example, consider formula $\varphi = aUb$
 25 and two segments, seg_1 and seg_2 across two cores, Cr_1 and Cr_2 , respectively. In
 26 order for the monitor running on Cr_2 to give the correct verdict, it must know
 27 the result of the monitor running on Cr_1 . In a scenario, where Cr_1 observes
 28 one or more $\neg a$ in seg_1 , a violation must be reported even if Cr_2 does not
 29 observe b and no $\neg a$. Generally speaking, the temporal order of events makes
 30 independent evaluation of segments impossible for LTL formulas. Of course,
 31 some formulas such as *safety* (e.g., $\Box p$) and *co-safety* (e.g., $\Diamond q$) properties are
 32 exceptions.

33 For our automata-based approach, we address this problem in two steps.
 34 Let $\mathcal{M}_\varphi = (\Sigma, Q, q_0, \delta, \lambda)$ be an LTL₃ monitor. Our first step is to create a
 35 3-dimensional reachability matrix RM by solving the following SMT decision
 36 problem: given a current monitor state $q_j \in Q$ and segment seg_i , can this
 37 segment reach monitor state $q_k \in Q$, for all $i \in [1, g]$, and $j, k \in [0, Q - 1]$. If
 38 the answer to the problem is affirmative, then we mark $RM[i][j][k]$ with **true**,
 39 otherwise with **false**. This is illustrated in Fig. 9 for the monitor shown in

	seg ₁			seg ₂			seg ₃			seg ₄		
	q ₀	q _⊤	q _⊥	q ₀	q _⊤	q _⊥	q ₀	q _⊤	q _⊥	q ₀	q _⊤	q _⊥
q ₀	T	F	F	T	T	F	T	T	T	T	T	T
q _⊤	F	F	F	F	T	F	F	T	F	F	T	F
q _⊥	F	F	F	F	F	T	F	F	T	F	F	T

Fig. 9 Reachability Matrix for $a\mathcal{U}b$

1 Fig. 5, where the grey cells are filled arbitrarily with the answer to the SMT
2 problem. This step can be made embarrassingly parallel, where each element
3 of RM can be computed independently by a different computing core. One can
4 optimize the construction of RM by omitting redundant SMT executions. For
5 example, if $RM[i][j][\top] = \text{true}$, then $RM[i'][\top][\top] = \text{true}$ for all $i' \in [i, Q-1]$.
6 Likewise, if $RM[i][j][\perp] = \text{true}$, then $RM[i'][\perp][\perp] = \text{true}$ for all $i' \in [i, Q-1]$.

7 The second step is to generate a verdict reachability tree from RM . The
8 goal of the tree is to check if a monitor state $q_m \in Q$ can be reached from
9 the initial monitor state q_0 . This is achieved by setting q_0 as the root and
10 generating all possible paths from q_0 using RM . That is, if $RM[i][k][j] = \text{true}$,
11 then we create a tree node with label q_j and add it as a child of the node
12 with the label q_k . Once the tree is generated, if q_m is one of the leaves, only
13 then we can say q_m is reachable from q_0 . In general, all leaves of the tree are
14 possible monitoring verdicts. Note that creation of the tree is achieved using a
15 sequential algorithm. For example, Fig.10 shows the verdict reachability tree
16 generated from the matrix in Fig. 9.

17 For our progression-based approach, we adhere to a similar technique for
18 parallelized monitoring as our automata-based approach. The key difference
19 being, in the progression-based approach subformulas are used, whereas in
20 the automata-based approach different states are used. As an example, the
21 previous formula $\varphi = a\mathcal{U}b$ will be broken into two subformulas $\varphi_1 = \Box a$ and
22 $\varphi_2 = \Diamond b$, before creating the reachability matrix, and then generating the
23 verdict for both these subformulas.

24 **Lemma 5** A distributed computation $(\mathcal{E}, \rightsquigarrow)$ of length l satisfies an LTL formula φ
25 if and only if the parallelized monitoring technique satisfies φ . That is,

$$\top \in [(\mathcal{E}, \rightsquigarrow) \models_3 \varphi] \iff \lambda(q) = \top$$

26 and,

$$\perp \in [(\mathcal{E}, \rightsquigarrow) \models_3 \varphi] \iff \lambda(q) = \perp$$

27 Where $q \in Q$ is some leaf node in the verdict reachability tree generated from RM
28 during the parallelized monitoring process and λ is the labelling function in \mathcal{M}_φ .

29 **Base Case:** Let us first consider the case where there is only one segment. That is,
30 $l = g$.

31
32




Fig. 10 Reachability Tree for aUb

1 (\Rightarrow) If $\top \in [(\mathcal{E}, \rightsquigarrow) \models_3 \varphi]$ (resp., $\perp \in [(\mathcal{E}, \rightsquigarrow) \models_3 \varphi]$), then according to the
 2 construction of the corresponding verdict reachability tree made from the RM ,
 3 the root node q_0 must have a child q_\top (resp., q_\perp), such that, $\lambda(q_\top) = \top$ (resp.,
 4 $\lambda(q_\perp) = \perp$). This child is also a leaf node, as the height of a verdict reachability
 5 tree is 2 when there is only one segment.

6
 7 (\Leftarrow) We can trivially show that if $\lambda(q_\top) = \top$ (resp., $\lambda(q_\perp) = \perp$), that is, if q_\top
 8 (resp., q_\perp) is reachable from q_0 , then $\top \in [(\mathcal{E}, \rightsquigarrow) \models_3 \varphi]$ (resp., $\perp \in [(\mathcal{E}, \rightsquigarrow) \models_3 \varphi]$).
 9

10 **Hypothesis:** Let us assume the proof as been established for $l = g \times k$. Now we
 11 consider $l = q \times (k + 1)$ as the segment length.

12
 13 (\Rightarrow) If $\top \in [(\mathcal{E}, \rightsquigarrow) \models_3 \varphi]$ (resp., $\perp \in [(\mathcal{E}, \rightsquigarrow) \models_3 \varphi]$), then according to our
 14 assumption, there must be at least one node at height $k + 1$ (height of the leaf
 15 nodes where there are k segments), such that $\lambda(q_\top) = \top$ (resp., $\lambda(q_\perp) = \perp$). Now
 16 for $k + 1$ number of segments, according to the construction of the corresponding
 17 verdict reachability tree made from the RM , the node q_\top (resp., q_\perp) can only have
 18 the child q_\top (resp., q_\perp). Therefore, there must be at least one node at height $k + 2$
 19 (height of the leaf nodes when there are $k + 1$ segments), such that $\lambda(q_\top) = \top$
 20 (resp., $\lambda(q_\perp) = \perp$).
 21

22 (\Leftarrow) We can trivially show that if $\lambda(q_\top) = \top$ (resp., $\lambda(q_\perp) = \perp$), that is, if q_\top
 23 (resp., q_\perp) is reachable from q_0 , then $\top \in [(\mathcal{E}, \rightsquigarrow) \models_3 \varphi]$ (resp., $\perp \in [(\mathcal{E}, \rightsquigarrow) \models_3 \varphi]$).
 24

24 6 Case Studies and Evaluation

25 In this section, we emphasize on analyzing our SMT-based solution with-
 26 out digressing into analyzing other dimensions such as instrumentation, data
 27 collection, data transfer, monitoring, etc., as given the distributed setting, run-
 28 time will be the dominant factor over any other kind of overhead. We evaluate

1 our proposed technique using synthetic experiments, Cassandra (a distributed
2 database), and the RACE dataset from NASA [9].

3 **6.1 Implementation and Experimental Setup**

4 Each experiment can be divided into three phases: (1) data generation, (2)
5 data collection and (3) data verification. For data-generation, we develop a
6 synthetic program that randomly generates a distributed computation (i.e.,
7 the behavior of a set of programs in terms of their local computations and
8 inter-process communication). Generating synthetic experimental data offer
9 benefits that enable us to draw comparison between different parameters and
10 their effect on the approach. For example, generating data for different values
11 of ε is beneficial to study its effect on the runtime and the number of false
12 warning verdicts of our approach.

13 When developing the synthetic distributed system as part of our experi-
14 ment, we ensure a partially-synchronous setting by including an HLC imple-
15 mentation. We use a uniform distribution $(0, 2)$ to define the type of event
16 (local computation, send and receive message) and a flip-coin distribution for
17 computing the atomic propositions that are true at each local computation
18 event. Although the events in our synthetic experiments in Section 6.2 are
19 uniformly distributed over the length of the trace, the event distribution as
20 part of the Cassandra experiments in Section 6.3 are affected by the network
21 latency and other external factors. In addition, we assume that there is
22 an external data collection program which keeps track of the data/states of
23 the system under verification. It generates the trace logs which is used by the
24 monitoring program to verify against the given LTL specifications mentioned
25 in Figure 11b.

26 For data verification, we consider the following parameters: (1) number of
27 processes ($|\mathcal{P}|$), (2) computation duration (l secs), (3) segment length (g), (4)
28 event rate (r events/process/sec), (5) maximum clock skew (ϵ), (6) depth of
29 the automaton (d) and number of nested temporal operators ($|\phi|$) for the LTL
30 formula under monitoring. The main *metric* is to measure the runtime of SMT
31 solving for each configuration of the parameters. Note that the time axis is
32 shown in log-scale in all the plots presented in this section. When we analyze
33 the effects of one parameter by holding the value of all the other parameters
34 at a relevant constant value. In all the graphs, we compare the runtime of our
35 automata-based approach against the progression based approach. We use a
36 MacBook Pro with Intel i7-7567U(3.5Ghz) processor, 16GB RAM, 512 SSD
37 and g++ Apple clang version 12.0.5 (clang-1205.0.22.9) interface to the Z3
38 SMT-solver [17] to generate the traces. To evaluate our parallel algorithm, we
39 use a server with 2x Intel Xeon Platinum 8180 (2.5GHz) processor, 768GB
40 RAM, 112 vcores and g++(GCC) 9.3.1 interface to the Z3 SMT-solver [17].
41 Unless specified otherwise, the system under consideration has $|\mathcal{P}| = 2$, $l = 2$
42 sec, $g = 250ms$, $r = 10$ events/process/sec, $\epsilon = 250ms$ and $d = 3$.

6.2 Analysis of Results – Synthetic Experiments

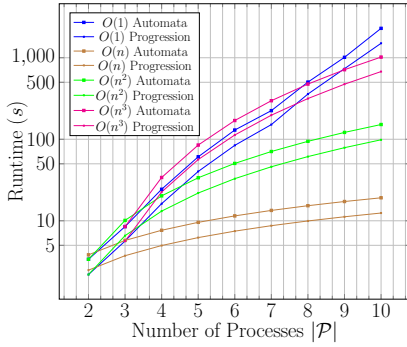
In this set of experiments, we exhaust all the available parameters and note how it affects SMT solving. We test each parameter individually to study its effect on runtime. As our generated synthetic data does not depend on any external factors, we induce a delay to not only limit the number of events happening at every time unit, but also to ensure uniform distribution of events over the execution of each process. We use a uniform distribution $(0, |\Sigma|)$ to assign a value to each local computation event in each process. We only use one CPU core for the following experimental results.

Overall, we notice an improvement of around 35% when the progression based technique is compared to the other automata based approach. This improvement in performance owes to two main reasons: (1) compared to the automata-based approach, the LTL constrains in our progression-based approach is less demanding in terms of computational complexity. Each sub-formula consists of mostly one atomic proposition as opposed to multiple atomic propositions in each path of the automaton, which in turn speeds up the overall verification process, and (2) the total number of SMT-instances needed is fewer due to the less number of sub-formulas compared to automaton paths given the same specification. We now analyze the results in detail.

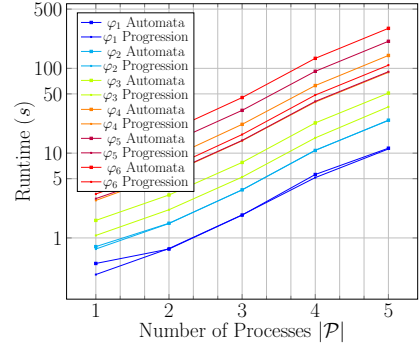
Impact of predicate structure. In this experiment (Figure 11a), we consider different predicate distribution over AP for the formula, φ_1 , i.e., how many processes are involved with a particular predicate. We consider different predicate structures: $O(1)$, $O(n)$, $O(n^2)$ and $O(n^3)$ which signifies the order of the number of SMT-encodings that need to be generated for the given distribution of predicates. As can be seen, the progression based technique outperforms the automata-based technique overall by 35% on average.

Having said that, during our experiments when comparing the runtime of our monitoring approach for increasing number of sub-formulas, we observe a slight decrease in the overall efficiency in runtime when using the progression-based approach compared to the automata-based approach. Since the progression-based approach is based on evaluating each sub-formula, there exists an LTL formula where the number of sub-formulas is more than the number of paths in the corresponding automata, and thus, the the progression-based approach might not be as efficient as the automata-based approach in such a scenario.

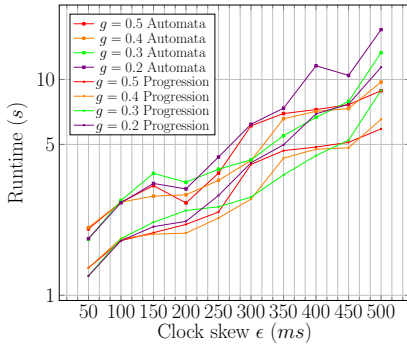
For example, consider a formula, $\varphi = \diamond a \vee \diamond b \vee \diamond c$, where the automata has two states, which makes the number of paths to be 2. However, the progression involves 3 sub-formulas, which makes the progression based approach less efficient than its automata counterpart. We would like to point out that the formula can be rewritten as $\diamond(a \vee b \vee c)$, which makes both the approaches yield similar results. Thus we hypothesize that for all LTL formulas, the progression-based approach will be more (if not equally) efficient to that of the automata-based approach.



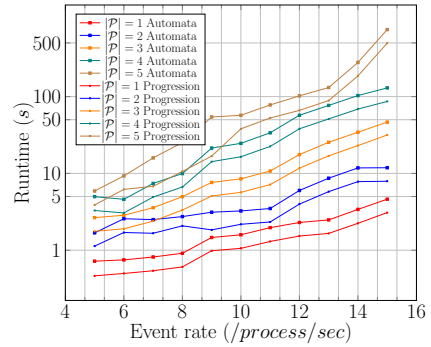
(a) Predicate Structure



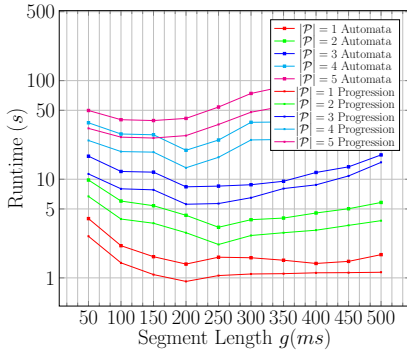
(b) LTL Formula



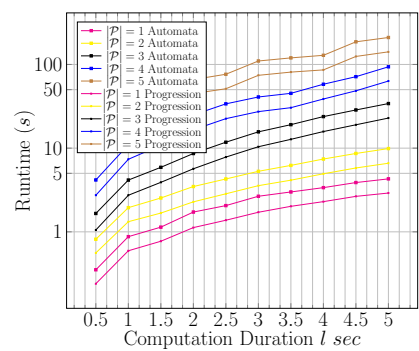
(c) Epsilon



(d) Event Rate



(e) Segment Length



(f) Computation Duration

Fig. 11 Synthetic experiments – impact of different parameters.

- 1 **Impact of LTL formula.** Given an LTL formula, the depth of nested temporal
- 2 operators plays an important role as suggested by Fig. 11b. We experimental
- 3 with the following LTL formula and the progression based technique achieved

1 an average improvement of 32.8% compared to the automata-based one.

$\varphi_1 = \Box p$	$d = 2$	$ \phi = 1$
$\varphi_2 = \Box(q \rightarrow \Box p)$	$d = 3$	$ \phi = 2$
$\varphi_3 = \Box((q \wedge \Diamond r) \rightarrow (\neg p \mathcal{U} r))$	$d = 4$	$ \phi = 3$
$\varphi_4 = \Box((q \wedge \Diamond r) \rightarrow (\neg p \mathcal{U}(r \vee (s \wedge \neg p \wedge \bigcirc(\neg p \mathcal{U} t))))))$	$d = 5$	$ \phi = 8$
$\varphi_5 = \Diamond r \rightarrow (s \wedge \bigcirc(\neg r \mathcal{U} t) \rightarrow \bigcirc(\neg r \mathcal{U}(t \wedge \Diamond p)))$	$d = 6$	$ \phi = 8$
$\varphi_6 = \Box((q \wedge \Diamond r) \rightarrow (s \wedge \bigcirc(\neg r \mathcal{U} t) \rightarrow \bigcirc(\neg r \mathcal{U}(t \wedge \Diamond p)))) \mathcal{U} r$	$d = 7$	$ \phi = 9$

2 **Impact of partial synchrony.** Figure 11c shows an expected result where
 3 increasing clock skew ϵ results in greater runtime as the number of possible
 4 concurrent events across processes increases exponentially. When comparing
 5 with the automata-based approach, the progression-based technique yields us
 6 an improvement of 33.36%.

7 **Impact of event rate.** Figure 11d shows that our approach breaks even with
 8 the computation duration for $|\mathcal{P}| = 3$ for an event rate of $5 \text{ events/process/sec}$.
 9 However, increasing the event rate increases the search space for the SMT
 10 solver. Overall we improve by 34.4% by using the progression-based technique
 11 compared to the automata-based technique.

12 **Impact of segment count.** Increasing the segment length increases the
 13 number of events to be worked with, and therefore, exponentially increasing
 14 the runtime of our approach. In Fig. 11e, we do not see much improvement for
 15 $|\mathcal{P}| = 1, 2$, since the number of events is not large enough to make an impact.
 16 However, we see better performance with low segment length for higher number
 17 of processes. Note, the runtime increases for very small segment length,
 18 since the time taken to generate a higher number of SMT encodings outweigh
 19 the performance gain from smaller segments. Here too, we notice an improve-
 20 ment of 32.6% for the progression-based technique over the automata-based
 21 technique.

22 **Impact of computation duration.** In this experiment (Fig. 11f), we
 23 increase computation duration and measure its effect on runtime. With increas-
 24 ing computation duration, the number of segments needed to verify the longer
 25 computation increases, and thereby resulting in a linear increase of the run-
 26 time. The progression-based approach improves the runtime by 33.1% when
 27 compared to the automata-based approach.

28 **Impact of parallelization.** Distributing the verification among multiple
 29 cores improves the performance of the approach by a considerable amount.
 30 As seen in Figure 12a, increasing the number of cores from 1 to 10 improves
 31 the performance by a huge margin. However, increasing it further shows little
 32 improvement, as the time taken for generating the SMT encodings starts to
 33 dominate the time taken to solve it. An improvement of 33.8% is achieved for
 34 progression-based approach when compared to automata-based approach.




Fig. 12 Impact of parallelization on different data

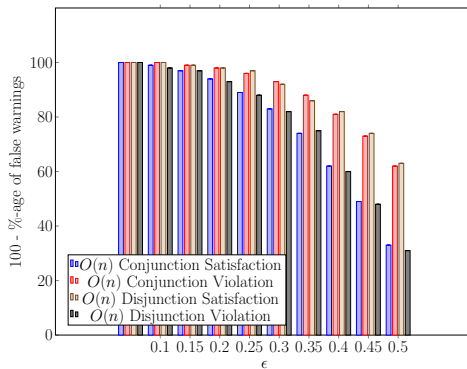


Fig. 13 False Warnings for Synthetic Data

1 **Impact of ϵ on false warnings.** As discussed in Section 2.4, the monitor
 2 does not have access to the global clock, it can report events as concurrent,
 3 when in reality, one happened before the other in the system under observa-
 4 tion. However, during this experiment, we keep track of the global clock values
 5 separately, which gives us full knowledge over the total ordering of all events.
 6 Thus, allowing us to study and report the *real verdicts* alongside the *reported*
 7 *verdicts*. We observe that the monitor sometimes report *false warnings*, that
 8 is, it reports both verdicts (satisfaction and violation), when in reality, only
 9 one has occurred. Note that the monitor never fails to report real verdicts.
 10 However, it may report false warnings alongside real verdicts on some occa-
 11 sions. Although this does not change the correctness of the approach, it may
 12 still include false warnings as part of the set of evaluated results.

13 In Figure 13, we observe that with the increase of the maximum clock skew
 14 ϵ , the number of false warnings increases. The increase in false warnings is
 15 attributed to the fact that as the value of ϵ increases, so does the number of
 16 events considered as concurrent by the monitor.

17 Additionally, we observe that the number of false warning is greatly influ-
 18 enced by the predicate structure of the LTL formula, as evident from Figure 13.
 19 For $O(n)$ conjunctive satisfaction formula monitoring and $O(n)$ disjunctive
 20 violation formula monitoring, false warnings might occur if any one of the n

1 sub-formulas are violated or satisfied, respectively. Therefore, we see a higher
 2 number of false warnings. Similarly, for $O(n)$ disjunctive satisfaction formula
 3 monitoring and $O(n)$ conjunctive violation formula monitoring, false warnings
 4 might occur if all of the n sub-formulas are violated or satisfied, respectively.
 5 Therefore, we see a lower number of false warnings.

6 6.3 Case Study 1: Cassandra

7 Cassandra [18] is a No-SQL distributed database management system. We
 8 simulate a distributed database with two data-centers: one cluster consisting
 9 of 4 nodes, and the other cluster consisting of 3 nodes, with one node from
 10 each cluster serving as the seed node. All data is replicated among every node
 11 in both the clusters. Each node runs on *Red Hat OpenStack Platform* using
 12 4 VCPUs, 4GB RAM, Ubuntu 1804, Cassandra 3.11.6, and Java 1.8.0.252.
 13 We have also simulated a system of multiple processes where each process
 14 is responsible for the basic database operations (read, write and update).
 15 These processes are also capable of inter-process communication that allows
 16 for informing other processes in case of a write of a new entry to the database.

17 To make our simulated database realistic, we tested the latency of our
 18 system to the ones offered by Google Cloud, Microsoft Azure and Amazon
 19 Web Services. The fastest response was clocked at $41ms$ compared to $100ms$
 20 from our system. The reason behind such a high latency when compared to the
 21 industry standard owes to the slow bandwidth and infrastructure differences.
 22 We consider a latency of $100ms$ for all our experiments, and fix maximum
 23 clock skew ϵ at $250ms$.

24 We design the processes such that each process is capable of reading,
 25 writing, or updating the entries in the database. We use a $(0, 2)$ uniform dis-
 26 tribution to select the type of the operation that is to be performed by the
 27 process. Once there is any kind of addition from the write operation, the change
 28 is notified to the other processes using the inter-process communications. We
 29 consider no loss of messages in transmission and all messages are read by the
 30 receiving process immediately once they are received.

31 In a database, consistency level helps maintain the minimum of replications
 32 that needs to be performed on an operation in order to consider the operation
 33 to be successfully executed. According to the recommendations from Cassan-
 34 dra the sum of read and write consistency should be more than the replication
 35 factor so as to remove any chances of read or write anomaly in the database.
 36 We aim to monitor and identify *read/write anomalies* in the database using
 37 runtime monitoring techniques. The corresponding LTL specification becomes:

$$\varphi_{rw} = \bigwedge_{i=0}^n \left(write(i) \rightarrow \diamond read(i) \right)$$

38 where n is the number of read/write operations.

39 One of the challenges for using a distributed database such as Cassandra
 40 is the lack of normalization (of database) capabilities. Therefore, we aim to

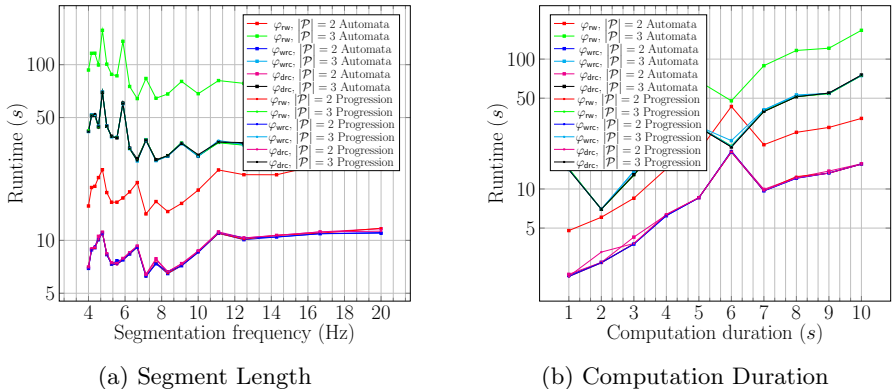


Fig. 14 Cassandra experiments

1 monitor *write reference check* and *delete reference check*. We introduce two
 2 tables:

STUDENT(*id, name*) ENROLLMENT(*id, course*)

3 We enforce the write and delete reference check on the tables above. For a
 4 write in the ENROLLMENT table, it should always be preceded by a write in
 5 the STUDENT table with the same *id*. Similarly, for a delete from the STU-
 6 DENT table, it should always be preceded by a delete from the ENROLLMENT
 7 table with the same *id*. These enforces no insertion and deletion anomaly, and
 8 therefore, leads to the following LTL specification:

$$\varphi_{wrc} = \neg \left(\neg \text{write}(\text{STUDENT.id}) \mathcal{U} \text{write}(\text{ENROLLMENT.id}) \right)$$

9

$$\varphi_{drc} = \neg \left(\neg \text{delete}(\text{ENROLLMENT.id}) \mathcal{U} \text{delete}(\text{STUDENT.id}) \right)$$

10 **Extreme load scenario.** Figure 14b and 14a plot runtime vs computa-
 11 tion duration and runtime vs segmentation frequency respectively, under full
 12 read/write load allowed by our network. When compared with the results
 13 from that of the synthetic experiments, these results are slightly noisier. This
 14 owes to the fact that in the synthetic experiments, the events were evenly
 15 spread over the entire computation duration, whereas here they are not uni-
 16 form. Database operations involving network communications (read, write and
 17 update) takes an average of 100ms, however sending and receiving of messages
 18 are inter-process communication, and takes about 10-15ms, making the over-
 19 all event distribution non-uniform. When comparing with the automata-based
 20 approach, we do not see much improvement when monitoring φ_{wrc} or φ_{drc}
 21 using progression based approach. However, when monitoring φ_{rw} , we observe
 22 an average improvement of 55.53%.

23 **Moderate load scenario.** In Figure 14b, we were able to make even for num-
 24 ber of processes as low as 2. Now, to look for a real-life example with moderate
 25 database operations we consider Google Sheets API, which allows a maximum

of 500 requests per 100 seconds per project and a 100 requests per second per user, i.e., on an average 5 events/sec per project and a user can only generate 1 event/sec. To evaluate how our approach performs in such a scenario, we increase the number of processes and the number of cores available to monitor such a system to study the time taken to verify the trace generated by such a system. We plot our findings in Fig. 12c, and notice that we break even for an event rate of 3 events/sec/user considering the progression-based approach. This is a significant improvement over the automata-based approach, where we could only break even for an event rate of 2 events/sec/user. Our algorithm performs well when the number of processes are 7, 8 and 9 which is much more than what is permitted by Google. This allows for us to be confident that our approach can pave way for implementation in a real-life settings.

6.4 Case Study 2: RACE

Runtime for Airspace Concept Evaluation (RACE) [9] is a framework developed by NASA that is used to build an event based, reactive airspace simulation. We use a dataset developed using this RACE framework². This dataset contains three sets of data collected on three different days. Each set was recorded at around 37° N Latitude and 121° W Longitude. The dataset includes all 8 types of messages being sent by the SBS unit by using a Telnet application to listen to port 30003, but we only use the messages with ID *MSG - 3* which is the Airborne Position Message and includes a flight's latitude, longitude and altitude using which we verify the mutual separation of all pairs of aircraft.

On analyzing the dataset, we observe that the time difference between the time message was generated to the time message was logged is usually less than a second apart, thus we considered an $\epsilon = 1s$ over the time message was generated. Furthermore, calculating the distance between two coordinates is computationally expensive, as we need to factor in parameters such as curvature of earth. In order to speed up distance related calculations, we consider a constant latitude of 111.2km and longitude of 87.62km, at the cost of a negligible error margin. We use these as constants and multiply them by the difference in latitude and longitude, and factor in the altitude to get the distance between two aircrafts. We verify *mutual separation* by considering the minimum separation between every pair of aircrafts to be 500m. From the dataset, we observe that each aircraft generates a message on at least 1 sec intervals. There are 3 separate datasets: sbs-1 consists of 293 aircrafts, 168,283 messages spread over 3 hours and 28 minutes and 58seconds; sbs-2 consists of 110 aircrafts, 64,218 messages spread over 1 hour 1 minute and 46 seconds; sbs-3 consists of 97 aircrafts, 64,162 messages spread over 49 minutes and 42 seconds.

In Fig. 12b, we compare our achieved runtime against the three datasets available from RACE (labelled sbs-1, sbs-2 and sbs-3). We monitor the data in *real time*, with 10s long segments and ϵ of 1s. We test our approach using

²<https://github.com/NASARace/race-data>

1 the parallelization technique introduced in 5.2 by using more number of cores
2 on the processor and utilize all available cores. Our results break even for 4
3 cores. This makes our approach desirable for aircraft monitoring and similar
4 systems such as IoT.

5 **7 Related Work**

6 Both centralized and decentralized monitoring approaches have been exten-
7 sively studied for synchronous and asynchronous systems. In the sequel, we
8 focus on reviewing the existing work on monitoring of distributed system, and
9 subsequently proceed to compare them with ours.

10 *Synchronous Distributed System*

11 Monitoring of LTL formulas in a synchronous distributed system is studied
12 in [2, 19, 20] under the assumption of a global clock shared across all compo-
13 nents. Designing large distributed systems with a common global clock shared
14 across all the processes that are usually in different geographical locations, is
15 challenging and expensive.

16 Other monitoring approaches include [19–22], where the authors employ
17 a decentralized monitor to evaluate the health of a system. In [20], the
18 authors introduce a way of organizing sub-monitors for LTL subformulas in a
19 synchronous distributed system, called choreography. In particular, the mon-
20 itors are organized as a tree across the distributed system, and each child
21 feeds intermediate results to its parent in a manner similar to diffusing using
22 computation. They formalize choreography-based decentralized monitoring by
23 showing how to synthesize a network from an formula, and give a decentralized
24 monitoring algorithm working on top of an LTL network.

25 The authors in [21], study a decentralized monitoring approach for
26 stream-based runtime verification technique with respect to a stream-based
27 specification language LOLA [23]. LOLA allows other data-types and aggre-
28 gated functions like average, median, etc. In [22], the authors deal with an extra
29 challenge where the monitors are only available to read a sub-computation and
30 are vulnerable to crash-faults. The authors show that by replacing few transi-
31 tion in the monitor automata, it is possible to remove non-determinism that
32 were introduced due to monitors only having a partial view of the system. We
33 present a more robust solution in this paper, where $\varepsilon = 0$ would allow for each
34 event occurring in the distributed system to be totally ordered, and therefore
35 replicating the behavior of that of a synchronous distributed system.

36 *Asynchronous Distributed System*

37 Lattice-theoretic centralized and decentralized online predicate detection in
38 asynchronous distributed systems has been extensively studied in [24, 25], and
39 extended to a more generalized temporal operator in [3, 4]. In contrast, our
40 paper is under a more practical assumption, where a clock synchronization

1 algorithm is utilized to limit the time window of asynchrony, as a result, lim-
2 iting the number of possible concurrent events. In [26], predicate detection is
3 shown to be a powerful tool in solving combinatorial optimization problems.
4 In Figure 12c, we show that our approach is effective in solving predicate
5 detection as well. Predicate detection using SMT has been studied in [27, 28],
6 whereas in our work, we use a more expressive temporal logic (LTL) to express
7 the specification of the system. Predicate detection for asynchronous sys-
8 tem is studied in [29], however the assumption to evaluate happened-before
9 relationship between events is too strong.

10 Knowledge-based monitoring of distributed processes is studied in [30],
11 where the authors propose a method for monitoring safety properties in dis-
12 tributed systems using past-time LTL. The main draw of this approach is being
13 able to produce false negative results. In [31], the authors were able to solve the
14 problem of decentralized monitoring of asynchronous distributed system even
15 when the monitors were vulnerable to crash-faults. The methodology include
16 introducing a family of logics called LTL_{2k+4} that refines the 4-valued LTL by
17 incorporating $2k + 4$ truth values, for each $k \geq 0$, where k is the maximum
18 number of crash-faults that the system could face. The truth values of LTL
19 $2k+4$ can be effectively used by each monitor to reach a consistent global set
20 of verdicts for each given formula, provided k is sufficiently large.

21 A synthetic benchmarking tool is proposed in [32] for assessing monitor-
22 ing overhead. While this tool is generally for systems that execute on a single
23 node, the authors show that it is possible to extend its applications to a dis-
24 tributed setting where the system under inspection do not share memory or
25 a global clock, and only communicate through asynchronous message passing.
26 In contrast, we consider a partially synchronous setting in our work ensured
27 by a clock synchronization algorithm. Furthermore, we allow local computa-
28 tion across processes, not just message send and message receive events, which
29 adds to the overall complexity of the system under observation.

30 In our work, we present both the automata-based approach discussed in [8]
31 and a completely new approach of progression, which monitors the same dis-
32 tributed system, but with increased efficiency. We assume that the processes
33 do not share a global clock and there by can be implemented as part of a
34 large geographically separated system. However, the presence of the clock syn-
35 chronization algorithm makes way for a bounded maximum clock skew which
36 reduces the size of the state space considerably when compared to that of an
37 asynchronous system. Through our experiments we study the effects of mul-
38 tiple parameters, including clock-synchronization constant and segmentation,
39 on the overall runtime of the approach. To the best of our knowledge there is
40 no runtime monitoring approach that assumes partial-synchrony among pro-
41 cesses and yet includes system specification in LTL. This is the reason behind
42 us comparing results internally between the automata and progression based
43 approach. Low values of ε correspond to a synchronous system, where as large
44 values of ε can be considered to follow the behavior of an asynchronous dis-
45 tributed system We show an average of 35% speedup using our progression

1 approach when compared with the automata-based approach over a wide range
2 of experiments. Our solution is also SMT based, and therefore, scalable and
3 efficient.

4 **8 Conclusion and Future Work**

5 In this paper, our focus was on distributed runtime monitoring. Both of our
6 proposed techniques take as input an LTL formula and a distributed com-
7 putation, and by assuming a bounded clock skew among all processes, they
8 first chop the computation into multiple segments, and then apply either
9 the automata-based monitoring algorithm, or progression-based monitoring
10 algorithm implemented as an SMT decision problem in order to verify the
11 correctness of the said formula. We conducted rigorous synthetic experiments,
12 as well as case studies on monitoring consistency conditions in Cassandra
13 and a NASA air traffic control dataset. Our experiments demonstrate up to
14 35% improvement in performance in our progression-based algorithm over our
15 automata-based algorithm.

16 For future work, we plan to study the tradeoff between accuracy and scala-
17 bility in our approach. Another important extension of our work is distributed
18 RV for timed temporal logics. Such expressiveness will allow us to monitor dis-
19 tributed applications that are sensitive to explicit timing constraints, such as
20 low-level blockchain and cross-chain functions.

21 **9 Acknowledgment**

22 This work is sponsored by the United States National Science Foundation
23 (NSF) awards FMitF 2102106 and SHF 2118356.

24 **References**

- 25 [1] Corbett, J.C., Dean, J., Epstein, M., Fikes, A., Frost, C., Furman, J.J.,
26 Ghemawat, S., Gubarev, A., Heiser, C., Hochschild, P., Hsieh, W., Kan-
27 thak, S., Kogan, E., Li, H., Lloyd, A., Melnik, S., Mwaura, D., Nagle,
28 D., Quinlan, S., Rao, R., Rolig, L., Saito, Y., Szymaniak, M., Tay-
29 lor, C., Wang, R., Woodford, D.: Spanner: Google’s globally distributed
30 database. *ACM Trans. Comput. Syst.* **31**(3) (2013). <https://doi.org/10.1145/2491245>
31
- 32 [2] Bauer, A., Falcone, Y.: Decentralised LTL monitoring. *Formal Methods*
33 *in System Design* **48**(1-2), 46–93 (2016)
- 34 [3] Ogale, V.A., Garg, V.K.: Detecting temporal logic predicates on dis-
35 tributed computations. In: *Proceedings of the 21st International Sympo-*
36 *sium on Distributed Computing (DISC)*, pp. 420–434 (2007)

- 1 [4] Mostafa, M., Bonakdarpour, B.: Decentralized runtime verification of LTL
2 specifications in distributed systems. In: Proceedings of the 29th IEEE
3 International Parallel and Distributed Processing Symposium (IPDPS),
4 pp. 494–503 (2015)
- 5 [5] Mills, D.: Network time protocol version 4: Protocol and algorithms
6 specification. RFC 5905, RFC Editor (June 2010)
- 7 [6] Bauer, A., Leucker, M., Schallhart, C.: Runtime Verification for LTL
8 and TLTL. *ACM Transactions on Software Engineering and Methodology*
9 (TOSEM) **20**(4), 14–11464 (2011)
- 10 [7] Havelund, K., Rosu, G.: Monitoring Programs Using Rewriting. In:
11 Automated Software Engineering (ASE), pp. 135–143 (2001)
- 12 [8] Ganguly, R., Momtaz, A., Bonakdarpour, B.: Distributed runtime verifi-
13 cation under partial asynchrony. In: Proceedings of the 24nd International
14 Conference on Principles of Distributed Systems (OPODIS), pp. 20–12017
15 (2020)
- 16 [9] Mehlitz, P., Giannakopoulou, D., Shafiei, N.: Analyzing airspace data
17 with race. In: 2019 IEEE/AIAA 38th Digital Avionics Systems Confer-
18 ence (DASC), pp. 1–10 (2019). [https://doi.org/10.1109/DASC43569.](https://doi.org/10.1109/DASC43569.2019.9081664)
19 [2019.9081664](https://doi.org/10.1109/DASC43569.2019.9081664)
- 20 [10] Pnueli, A.: The temporal logic of programs. In: Symposium on Founda-
21 tions of Computer Science (FOCS), pp. 46–57 (1977)
- 22 [11] Manna, Z., Pnueli, A.: *Temporal Verification of Reactive Systems: Safety*.
23 Springer, Berlin, Heidelberg (1995)
- 24 [12] Lamport, L.: Time, clocks, and the ordering of events in a distributed
25 system. *Communications of the ACM* **21**(7), 558–565 (1978)
- 26 [13] Kulkarni, S.S., Demirbas, M., Madappa, D., Avva, B., Leone, M.: Logical
27 physical clocks. In: Proceedings of the 18th International Conference on
28 Principles of Distributed Systems, pp. 17–32 (2014)
- 29 [14] Bauer, A., Falcone, Y.: Decentralised LTL monitoring. *Formal Methods*
30 *in System Design* **48**(1), 46–93 (2016)
- 31 [15] Bauer, A., Leucker, M., Schallhart, C.: Comparing ltl semantics for run-
32 time verification. *Journal of Logic and Computation* **20**(3), 651–674
33 (2010). <https://doi.org/10.1093/logcom/exn075>
- 34 [16] Garg, V.K.: *Elements of Distributed Computing*. John Wiley & Sons,
35 Inc., USA (2002)

- 1 [17] de Moura, L.M., Bjørner, N.: Z3: An efficient SMT solver. In: Tools and
2 Algorithms for the Construction and Analysis of Systems (TACAS), pp.
3 337–340 (2008)
- 4 [18] Lakshman, A., Malik, P.: Cassandra: A decentralized structured storage
5 system. *SIGOPS Oper. Syst. Rev.* **44**(2), 35–40 (2010). [https://doi.org/
6 10.1145/1773912.1773922](https://doi.org/10.1145/1773912.1773922)
- 7 [19] El-Hokayem, A., Falcone, Y.: On the monitoring of decentralized specifica-
8 tions: Semantics, properties, analysis, and simulation. *ACM Transactions
9 on Software Engineering Methodologies* **29**(1), 1–1157 (2020)
- 10 [20] Colombo, C., Falcone, Y.: Organising LTL monitors over distributed sys-
11 tems with a global clock. *Formal Methods in System Design* **49**(1-2),
12 109–158 (2016)
- 13 [21] Danielsson, L.M., Sánchez, C.: Decentralized stream runtime verifica-
14 tion. In: Proceedings of the 19th International Conference on Runtime
15 Verification (RV), pp. 185–201 (2019)
- 16 [22] Kazemloo, S., Bonakdarpour, B.: Crash-resilient decentralized syn-
17 chronous runtime verification. In: Proceedings of the 37th Symposium on
18 Reliable Distributed Systems (SRDS), pp. 207–212 (2018)
- 19 [23] D’Angelo, B., Sankaranarayanan, S., Sanchez, C., Robinson, W.,
20 Finkbeiner, B., Sipma, H.B., Mehrotra, S., Manna, Z.: Lola: runtime
21 monitoring of synchronous systems. In: 12th International Symposium on
22 Temporal Representation and Reasoning (TIME’05), pp. 166–174 (2005).
23 <https://doi.org/10.1109/TIME.2005.26>
- 24 [24] Chauhan, H., Garg, V.K., Natarajan, A., Mittal, N.: A distributed
25 abstraction algorithm for online predicate detection. In: Proceedings of
26 the 32nd IEEE Symposium on Reliable Distributed Systems (SRDS), pp.
27 101–110 (2013)
- 28 [25] Mittal, N., Garg, V.K.: Techniques and applications of computation
29 slicing. *Distributed Computing* **17**(3), 251–277 (2005)
- 30 [26] Garg, V.K.: Predicate detection to solve combinatorial optimization prob-
31 lems. In: Proceedings of the 32nd ACM Symposium on Parallelism in
32 Algorithms and Architectures, pp. 235–245 (2020)
- 33 [27] Valapil, V.T., Yingchareonthawornchai, S., Kulkarni, S.S., Torng, E.,
34 Demirbas, M.: Monitoring partially synchronous distributed systems
35 using SMT solvers. In: Proceedings of the 17th International Conference
36 on Runtime Verification (RV), pp. 277–293 (2017)

- 1 [28] Momtaz, A., Basnet, N., Abbas, H., Bonakdarpour, B.: Predicate moni-
2 toring in distributed cyber-physical systems. In: International Conference
3 on Runtime Verification, pp. 3–22 (2021). Springer
- 4 [29] Stoller, S.D.: Detecting global predicates in distributed systems with
5 clocks. In: Mavronicolas, M., Tsigas, P. (eds.) Distributed Algorithms, pp.
6 185–199. Springer, Berlin, Heidelberg (1997)
- 7 [30] Sen, K., Vardhan, A., Agha, G., G.Rosu: Efficient decentralized moni-
8 toring of safety in distributed systems. In: Proceedings of the 26th
9 International Conference on Software Engineering (ICSE), pp. 418–427
10 (2004)
- 11 [31] Bonakdarpour, B., Fraigniaud, P., Rajsbaum, S., Rosenblueth, D.A.,
12 Travers, C.: Decentralized asynchronous crash-resilient runtime verifica-
13 tion. In: Proceedings of the 27th International Conference on Concurrency
14 Theory (CONCUR), pp. 16–11615 (2016)
- 15 [32] Aceto, L., Attard, D.P., Francalanza, A., Ingólfssdóttir, A.: On bench-
16 marking for concurrent runtime verification. In: Guerra, E., Stoelinga,
17 M. (eds.) Fundamental Approaches to Software Engineering, pp. 3–23.
18 Springer, Cham (2021)