



HAL
open science

Sound Concurrent Traces for Online Monitoring

Chukri Soueidi, Yliès Falcone

► **To cite this version:**

Chukri Soueidi, Yliès Falcone. Sound Concurrent Traces for Online Monitoring. SPIN 2023 - 29th International Symposium on Model Checking Software, Apr 2023, Paris, France. pp.59-80, <10.1007/978-3-031-32157-3_4>. <hal-04381666>

HAL Id: hal-04381666

<https://inria.hal.science/hal-04381666v1>

Submitted on 9 Jan 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire HAL, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons CC BY 4.0 - Attribution - International License

Sound Concurrent Traces for Online Monitoring

Chukri Soueidi[✉] and Yliès Falcone[✉]

Univ. Grenoble Alpes, Inria, CNRS, Grenoble INP, LIG,
38000 Grenoble, France

{chukri.a.soueidi,ylies.falcone}@inria.fr

Abstract. Monitoring concurrent programs typically rely on collecting traces to abstract program executions. However, existing approaches targeting general behavioral properties are either not tailored for online monitoring, are no longer maintained, or implement naive instrumentation that often leads to unsound verdicts. We first define the notion of when a trace is representative of a concurrent execution. We then present a non-blocking vector clock algorithm to collect sound concurrent traces on the fly reflecting the partial order between events. Moreover, concurrent events in the representative trace pose a soundness problem for monitors synthesized from total order formalisms. For this, we extract a causal dependence relation from the monitor to check if the trace has the needed orderings and define the conditions to decide at runtime when a collected trace is monitorable. We implement our contributions in a tool, FACTS, which instruments programs compiling to Java bytecode, constructs sound representative traces, and warns the monitor about non-monitorable traces. We evaluate our work and compare it with existing approaches.

1 Introduction

Guaranteeing the correctness of concurrent programs often relies on dynamic verification which requires observing and abstracting the program behavior. Abstraction is typically provided by *traces* that contain the executed actions which can be analyzed online or post-mortem. Traces serve as a model for property-based *detection* and *prediction* techniques which choose their trace collection approaches differently based on the class of targeted properties. Some approaches target generic classic concurrency errors such as data-races [38,32], deadlocks [14], and atomicity violations [31,74,54]. Other techniques target general behavioral properties; those are typically order violations such as null-pointer dereferences [30], and tpestate violations [39,37,67] and more generally runtime verification [48,27].

When properties require reasoning about concurrency in the program, causality between events must be established during trace collection. Data race detection techniques for instance require it to check for concurrent accesses to shared variables; as well as predictive approaches targeting behavioral properties such as [22,39,67,38] in order to explore other feasible executions. Causality is best

expressed as a partial order over events, which are compatible with various formalisms for the behavior of concurrent programs such as weak memory consistency models [4,7,53], Mazurkiewicz traces [57,34], parallel series [49], Message Sequence Charts graphs [58], and Petri Nets [59].

However, while a concurrent program behaves non-sequentially, trace collection is sequential and the causality between events must be established during or after observation. For many monitoring approaches, the program is a black box and the trace is the sole system information provided. Thus, providing traces with correct and sufficient ordering information is necessary for sound and expressive monitoring. Any *representative* trace for a concurrent execution must only include correct information and must reflect all the ordering information needed to verify a behavioral property. Thus, it is best modeled as a partial order over program actions.

This paper addresses the problem of providing representative traces for existing verification techniques to soundly monitor concurrent programs *online*. We focus on general *behavioral* properties targeting violations that cannot be traced back to classical concurrency errors. These properties are usually expressed using total-order formalisms such as LTL and finite-state machines [51,1]. One example is a precedence property, which specifies that a resource can only be granted (event g) in response to a request (event r), and can be expressed in LTL as $\neg g \mathbf{W} r$. Unfortunately, these properties have received less attention in the literature in the context of concurrent programs.

Existing approaches that monitor such properties are either not tailored for online monitoring, no longer maintained, or implement naive instrumentation. Let us first consider dynamic *predictive* techniques that target behavioral properties, such as [22,39,67,38]. These techniques often rely on vector clock algorithms to timestamp events, construct partial orders and build causal models to explore feasible executions. Although they capture representative traces, they are expensive to use online and often intended for the design phase of the program and not in production environments (see Sec. 8). Moreover, by serializing reads and writes in the program, they assume a sequentially consistent execution model [36]. For simple detection, this forces a specific schedule on the execution, and for prediction, it restricts the space of possible interleavings. Furthermore, such techniques are incomplete and cannot prove the absence of property violations (they never claimed it) which is a reason why online monitoring is inevitable for many systems.

We turn to classical monitoring (*detection*) techniques and we find that many top existing approaches (based on the competitions on RV [28,63,11]) were initially developed to monitor single-threaded programs and later adapted to multithreaded programs. Unfortunately, as shown in [25], these approaches problematically impose a total order on events using incorrect assumptions. They have wrongly taken for granted that the instrumentation always yields representative traces. When events of interest execute in concurrent regions, these tools often produce unsound and inconsistent verdicts. These approaches can leverage representative traces to produce sound verdicts in the presence of concurrent events,

and their expressiveness can be extended to check concurrency-related behavioral properties. For example, one might be interested to verify the correctness of the scheduling algorithms of tasks with data dependencies (e.g., [73]).

We first revisit traces and qualify two properties that determine if they are good representatives of an execution: *soundness* and *faithfulness* (Sec. 4). Soundness holds when the trace does not provide false information about the order of events. Faithfulness holds when the trace contains all the information about the order of events from the execution. **RQ1:** *Can we collect a representative trace of a concurrent execution on the fly with minimal interference on the program?* Vector clock algorithms have been employed and refined for several decades now [18,54,5,65]. However, very few (we know of [64]) are directed towards online monitoring of behavioral properties; where a final trace consists of property-related events only. These algorithms typically require blocking the execution; by synchronizing the instrumentation, program actions, and algorithm’s processing to avoid data races [18]. With the quadratic bound on their runtime complexity and the coarse-grained synchronization they introduce, one might want to run such an algorithm off the critical path of the program. We present a vector clock algorithm (Sec. 5) that does not require blocking the execution and can run either synchronously or asynchronously with the program. Asynchronous trace collection is ideal for scenarios where the monitoring overhead cannot be afforded and a small delay in the verdict can be tolerated. For instance, in real-time systems where the system is expected to produce a result within a defined strict deadline [68]. Our algorithm constructs representative concurrent traces that are sound and very often faithful¹. As far as we know, it is unique in the context of monitoring behavioral properties that can run off the critical path of the execution.

RQ2: *Is the collected concurrent trace good enough to soundly monitor a property?* Monitoring single-threaded programs depend on instrumentation, which is assumed to be correct, to provide all relevant events [12]. For concurrent programs, we notice that monitorability with classical approaches depends also on the ordering information available in the trace (resp. the execution). Firstly, consider for instance the precedence property seen previously. If events r and g happen to execute concurrently in the program, a sound trace would have them as unordered events and a sound monitor should report a violation of the property. However, in practice, when monitoring such property with an automaton, the partial order must be linearized before passing it to the monitor. The linearization will produce an arbitrary order between r and g , for instance, $r.g$ which would make the oblivious monitor miss the violation. In such a case, the trace is not fit for monitoring the property and the monitor should be warned. Secondly, in certain scenarios due to some partial instrumentation or logging failure, some synchronization actions might be missing from the trace, resulting in unfaithful traces where some orderings cannot be established. This also poses soundness problems to the monitor similar to the problems with lossy

¹ The algorithm might miss, in some marginal cases, ordering information resulting in sound but unfaithful traces (see Sec. 5).




Fig. 1: Unsound instrumentation.

traces [40]. To handle the mentioned problems, we extract a *causal dependence relation* from a given property to know which events cannot permute in a trace and check whether a trace contains enough order information (Sec. 6). We then redefine trace *monitorability* for concurrent executions with a necessary condition on the trace to guarantee a sound verdict when monitoring. If the condition is not met, we produce warnings for the monitor.

In this paper we target shared memory concurrency, however, our work can be adapted to message passing. We implement our contributions in a tool, FACTS, which attaches to programs running on the JVM, instruments and collects sound and faithful traces, and checks the monitorability criteria. We evaluate our approach and demonstrate its effectiveness in capturing concurrent traces and provide a cost estimation on the execution time overheads (Sec. 7). We also implement the algorithm from [65], used in [66,39,22], and show how our algorithm has better advantages.

2 Issues with Linear Traces

Monitoring approaches relying on total-order formalisms such as LTL and finite state machines require linear traces to be fed to the monitors as their input consists of words. One might be tempted to immediately capture linear traces from a concurrent execution without reestablishing the causal order. In this section, we discuss issues of collecting linear traces to motivate concurrent traces.

Advice Atomicity Assumption General purpose runtime verification frameworks such as Java-MOP [21], Tracematches [8], and others [11] rely on instrumentation for extracting events for traces. To handle concurrent programs, RV tools provide a feature to *synchronize* the monitor to protect it from concurrent access and data races. As such, threads acquire a lock before notifying the monitor and release it afterward. However, when a program action is executed, its advice will not always execute atomically with it unless both are wrapped in a synchronization block. Consider Fig. 1 and the property r *precedes* g . A context switch happens at t_1 after the $r(x)$ is executed (but before notifying the monitor), allowing the execution of $g(x)$ and its advice before $\text{notify}(r, x)$. It is clear how this would result in unsound monitoring. This observation problem cannot be solved by simply checking for the absence of data races. One needs to guarantee atomicity between all executing program actions and their advice. However, the problem will not appear if r and g are synchronized and executed with their advice within mutually exclusive regions. Nevertheless, we might sometimes want to capture events from concurrent regions.

Forced Atomicity One way to solve the lack of advice atomicity is to force it. Forced atomicity can be achieved by instrumenting synchronization blocks² that wrap the program actions with their advice in mutually exclusive regions. However, forcing atomicity introduces two new problems, one at the program level and the other at the monitor level. First, it forces a total order between concurrent program actions; interfering with the parallelism of the program and changing its behavior. One needs also to minimize the area for which the lock is applied and avoid coarse-grained synchronization. From the monitor side, the verdict will be dependent on the specific scheduling of the execution. Take r precedes g for instance, if these events are concurrent in the program, monitoring will produce a different verdict at each run. Although the verdict is correct, we would like to warn the user. Second, any information about parallel actions in the program is lost and one can no longer determine whether two actions execute concurrently initially in the non-instrumented program. In that case, it becomes impossible to express properties on the concurrent parts of the execution. Furthermore, RV tools mentioned above rely on AspectJ [41] for a high-level specification of instrumentation which is rather unfitting to instrument synchronization blocks as it cannot instrument at the bytecode level.

3 Characterizing a Concurrent Execution

We first choose the smallest observable execution step done by a program, such as a method call or write operation, and refer to this step as an *action*. An *action* is a tuple of attribute-value pairs. Common attributes are *id*, associated *resource* i.e. shared variable, integer *id* to distinguish different actions corresponding to the same instruction³. Actions executing in the same thread follow a *thread order*. The set of all program actions is denoted by A . We distinguish between two types of actions: *regular actions* and *synchronization actions*. Regular actions are relevant to the property we want to monitor. Synchronization actions, denoted by the set SA , are used for synchronization between different threads. They provide *release-acquire* ordering semantics [3] and establish a *synchronization order* on the execution. Essentially, an acquire action a' on a resource by some thread synchronizes with the latest release action a , if it exists, on that same resource by some other thread. Given two threads t and u , we highlight the following synchronization actions that can establish release-acquire:

- $unlock(t, l)/lock(t, l)$: release/acquire of lock l by t ;
- $fork(t, u)/begin(u)$: fork of u by t /first action by u ;
- $end(t)/join(u, t)$: last action t / u blocking until t ends;
- $write(t, x, v)/read(t, x, v)$: value v on shared variable x .
- $notify(t, s)/wait(t, s)$: notify/wait a signal s .

² Such as `synchronized` in Java.

³ We depict actions using the notation $id.label_{resource}^{tid}$ in diagrams.

The *execution order* $\xrightarrow{e} \subseteq (\mathbb{A} \times \mathbb{A})$ is the transitive closure of both *thread* and *synchronization* orders. This order between actions is a partial order⁴ and is general enough to represent various formalisms and models of concurrent systems.

Definition 1 (Concurrent Execution). *A concurrent execution is a partially-ordered set of actions $(\mathbb{A}, \xrightarrow{e})$.*

4 Sound and Faithful Concurrent Traces

Our goal is to capture a representative trace from a concurrent execution. In [69], we introduce two fundamental notions for concurrent traces: *soundness* and *faithfulness*. We first revisit both notions here, then proceed to show how such traces preserve the properties of an execution. We note that the trace actions \mathcal{E} are often called *event* in monitoring and verification approaches.

Definition 2 (Concurrent Trace). *A concurrent trace is a partially ordered set $t = (\mathcal{E}, \xrightarrow{\text{tr}})$ of events such that $\mathcal{E} \subseteq \mathbb{A}$.*

Informally, a trace is a sound trace if it does not provide false information about the execution.

Definition 3 (Trace Soundness). *A concurrent trace t is said to be a sound trace of a concurrent execution e (noted $\text{snd}(e, t)$) iff (i) $\mathcal{E} \subseteq \mathbb{A}$ and (ii) $\xrightarrow{\text{tr}} \subseteq \xrightarrow{e}$.*

To be sound, a trace (i) should not capture an action not found in the execution, and (ii) should not order unrelated actions.

While a trace that does not provide incorrect information about the execution model leads to sound monitoring, a trace can still not provide enough information about the order (for a monitor to determine a verdict). Faithfulness is similar to *completeness*; it is expensive as it requires capturing all relevant causality in the program. Informally, a *faithful* trace contains all information on the order of events that occurred in the execution model.

Definition 4 (Trace Faithfulness). *A concurrent trace t is said to be faithful to an execution e (noted $\text{faith}(e, t)$) iff $\xrightarrow{\text{tr}} \supseteq (\xrightarrow{e} \cap \mathcal{E} \times \mathcal{E})$.*

Example 1 (Soundness and Faithfulness). Consider an execution of a properly synchronized Java program with three threads, a writer, and two readers, where the reader threads performed a concurrent read in between two write operations. We omit the depiction of actions performed by each thread such as locks, etc., for brevity. Four different collected traces from such execution are presented in Figure 2. For a behavioral property such as “no read or write happen concurrently with another read or write”, we are interested in the actions \mathbf{r} and \mathbf{w} . The order relative to these events in the execution is $\rightarrow_{e_0} = \{\langle 0.\mathbf{w}^0, \mathbf{r}^1 \rangle, \langle 0.\mathbf{w}^0, \mathbf{r}^2 \rangle, \langle \mathbf{r}^1, 1.\mathbf{w}^0 \rangle, \langle \mathbf{r}^2, 1.\mathbf{w}^0 \rangle, \langle 0.\mathbf{w}^0, 1.\mathbf{w}^0 \rangle\}$. Fig. 2a presents a linear

⁴ We assume familiarity with partial orders and partially ordered sets.




Fig. 2: Four different collected traces from the execution of 1-Writer 2-Readers.

trace of the execution t_0 as captured by Java-MOP using advice atomicity assumption, see Sec. 2. One can see that the order is $\rightarrow_{t_0} = \rightarrow_{e_0} \cup \{\langle r^1, r^2 \rangle\}$. We notice that we have faithfulness ($\text{faith}(e_0, t_0)$) as $\rightarrow_{e_0} \subset \rightarrow_{t_0}$. However we do not have soundness ($\neg \text{snd}(e_0, t_0)$) as the pair $\langle r^1, r^2 \rangle \notin \rightarrow_{e_0}$, as the reads happened concurrently. Fig. 2b shows a trace that is neither sound nor faithful. Fig. 2c the trace captures thread order only. It is a sound trace, as it only contains $\langle 0.w^0, 1.w^0 \rangle$, and therefore no wrong information. However, it is not faithful, as it is missing order information. Fig. 2d presents a partial trace of the execution t_3 that is both sound and faithful. Ideally, t_3 is the smallest concurrent trace collected to verify behavioral properties on reads and writes.

A property can be used to define the set of correct concurrent executions. For the verification of the temporal behavior of programs, the semantics of matching properties is applied on traces [52]. Effectively, at the semantic level, a property partitions the set of all executions into correct and incorrect ones. Consequently, a property $\varphi(\Sigma)$ defined over $\Sigma \subseteq \mathbb{A}$ is a set of partial orders over Σ .

Definition 5 (Property Satisfaction). *A concurrent execution $(\mathbb{A}, \xrightarrow{e})$ satisfies a property $\varphi(\Sigma)$ (noted $(\mathbb{A}, \xrightarrow{e}) \models \varphi(\Sigma)$) iff $(\xrightarrow{e} \cap \Sigma \times \Sigma) \in \varphi(\Sigma)$.*

To check that a property has been satisfied, we simply “project” the order \xrightarrow{e} on Σ , that is, we restrict our information about the execution to Σ . We then check if the projection belongs to the set of correct concurrent executions ($\varphi(\Sigma)$).

While our goal is verifying properties on the full execution of a program, we generally gather a subset of it as a trace. As such, we are interested in the fact that verifying a property on the trace holds the same as it would on the full execution. By construction (from Def. 3 and 4), we notice that the projections over Σ (Def. 5) of some execution e and a trace t for a property over variables $\Sigma \subseteq \mathcal{E} \subseteq \mathbb{A}$ where we have soundness ($\text{snd}(e, t)$) and faithfulness ($\text{faith}(e, t)$) are the same. We deduce the following theorem.

Theorem 1 (Property Preservation). *Given a concurrent execution $e = (\mathbb{A}, \xrightarrow{e})$, a trace $t = (\mathcal{E}, \xrightarrow{t})$, and a property φ over $\Sigma \subseteq \mathcal{E} \subseteq \mathbb{A}$ we have:*

$$(\text{snd}(e, t) \wedge \text{faith}(e, t)) \implies (e \models \varphi(\Sigma) \text{ iff } t \models \varphi(\Sigma)).$$

We say that t is an appropriate abstraction of e ; e and t can be used interchangeably to verify properties over Σ . Since our notion of a property is simply a set of traces, the presented results apply to any set of monitorable [26,60] properties.

5 Obtaining Sound Concurrent Traces

We here present a vector clock algorithm that constructs sound concurrent traces. The algorithm differs from standard algorithms in that it can run asynchronously, allowing scenarios where a delay can be tolerated. Most, if not all vector clock algorithms in the literature block the execution to process the algorithm at each event. It also does not require advice atomicity. Instead, it just requires instrumenting release and acquire actions with *before* (resp. *after*) directives. When an action a is instrumented with *before* (resp. *after*) direction, the notification action, $\text{notify}(a)$, should execute before a (resp. *after*), and no other action from the same thread should execute in between.

Let us first review notions of vector clocks [55,45]. A vector clock is a function from threads to integers, $V : T \rightarrow \text{Int}$ with three key operations. *Join* is the element-wise maximum of two vector clocks: $V_1 \sqcup V_2 = V_3$ such that $\forall t : V_3(t) = \max(V_1(t), V_2(t))$. *Comparison* is the element-wise comparison of two vector clocks: $V_a \leq V_b$ is defined to mean $\forall t. V_a(t) \leq V_b(t)$. Finally, function *inc* increases a single component of a vector clock: $\text{inc}_t(V) = \lambda u. (\text{if } u = t \text{ then } V(u) + 1 \text{ else } V(u))$.

5.1 The Reordering Algorithm

Algorithm 1 maintains the following data structures. A map $\mathbb{L} : T \rightarrow V$ holding the last timestamp seen by each thread. A map \mathbb{R} holding the last release action per resource, a map \mathbb{W} holding the last write of a value on a shared variable, and a collection \mathbb{T} which represent the concurrent trace and will be populated with timestamped actions. On each received notification action, the algorithm sets its timestamp to the latest timestamp seen by its thread (line 3). Then, synchronization actions in SA (lines 4-5) are send to **ReleaseOrAcquire**, except for *read* or *write*. Actions *unlock*, *fork*, *end* are represented with *release*, and the algorithm puts them in the map entry in \mathbb{R} associated with their resource in R (lines 11-12). Actions *lock*, *begin*, *join* are represented with *acquire*, the algorithm retrieves the last action that released its resource from R ; if found ⁵, their vector clocks are merged (lines 13-15). The *join* is merged with the last action seen by the finished thread. Actions *read* and *write* are handled with **ReadOrWrite**. They are handled with the map \mathbb{W} which is indexed by a shared variable and a value. A write (lines 17-23) is only pushed into \mathbb{W} when it is not conflicting with the entry in \mathbb{W} associated with its variable and value (discussed more in the following paragraph). A read is merged with the latest write (lines 24-26). After that, the timestamp is incremented (line 6), and map \mathbb{L} is updated (line 7). Events will be stored in the concurrent trace, and all synchronization actions will be discarded (lines 8-9).

The *Join*, *Comparison* and *Copy* operations of vector clocks require $\Theta(k)$ time, linear in the number of threads k . *Increment* operations on vector clocks,

⁵ We omit some null checks to simplify the presentation of the algorithm, however, assume that joining $e.VC$ with a null value does not affect it.

Algorithm 1: Vector Clock Algorithm

```

1 procedure ReceiveAction( $e$ )
2    $t = tid(e)$ 
3    $e.VC = copy(\mathbb{L}(t))$ 
4   if  $e$  is a synchronization action then
5     send  $e$  to the appropriate procedure
6    $inc_t(e.VC)$ 
7    $\mathbb{L}(t) = e.VC$ 
8   if  $e$  is not a synchronization action then
9      $\mathbb{T} \leftarrow e$  // Add  $e$  to trace  $\mathbb{T}$ 

10 procedure ReleaseOrAcquire( $e$ )
11   if  $e = release(t, r)$  then
12      $\mathbb{R}(r) = e$  // Update  $\mathbb{R}$ 
13   else if  $e = acquire(t, r)$  then
14      $e' = \mathbb{R}(r)$ 
15      $e.VC = e.VC \sqcup e'.VC$  // Merge vector clocks

16 procedure ReadOrWrite( $e$ )
17   if  $e = write(t, x, v)$  then
18      $e' = \mathbb{W}(x, v)$ 
19     if ( $e' = null$ )  $\vee$  ( $e'.VC \leq e.VC$ ) then
20        $\mathbb{W}(x, v) = e$ 
21     else
22        $\mathbb{W}(x, v) = null$  // Conflicting write
23     end
24   else if  $e = read(t, x, v)$  then
25      $e' = \mathbb{W}(x, v)$ 
26      $e.VC = e.VC \sqcup e'.VC$ 

```

retrieving and inserting elements to the maps require $O(1)$. The reordering algorithm then requires $O(n \times k)$ time for a trace of n actions.

5.2 Correctness Discussion

We now show how the algorithm always produces sound traces. After the algorithm ends, the collected linear trace is timestamped into a concurrent trace represented by \mathbb{T} that contains ordered pairs $\langle a, b \rangle$, such that $a.VC \leq b.VC$. For handling *release* and *acquire* actions (except for reads and writes), these actions already execute in a total order in any execution. The algorithm performs classical operations of timestamp merging for matching actions. However, the correctness of our algorithm is dependent on instrumentation. Since we do not force atomicity between an action and its notification, we must instrument actions that perform vector clock merge operations i.e. *acquire* actions with the *after* directive and *release* actions with the *before* directive. The intuition is that if a *release* action is instrumented with the *after* directive, a context switch between it and its advice can lead to having an unmerged consecutive *acquire*. With

proper instrumentation, whenever we observe an *acquire* action, we are sure that if a matching *release* exists, then it has been already been captured and processed. When the program only uses *lock* and *fork* actions for synchronization, then our algorithm guarantees both soundness and faithfulness of the concurrent trace since it captures all orderings, provided that they are all instrumented and captured.

Handling *reads* and *writes* is more demanding if we want to run the algorithm asynchronously and do not want to force on them a total order in the execution. Reads and writes are instrumented with *after* and *before* directives, respectively, including volatile and atomic variables [47]⁶. When observing a read r on shared variable x with a value v (line 24), we know its matching write was processed. Nevertheless, if prior to r , two (or more) threads performed writes, w and w' , on x with the same value v , then we need to make sure that w and w' are ordered since r could be reading the value written by any one of them; hence its clock cannot be safely merged with any of them. We say the two writes on x , writing the same value v , are *conflicting* when they are not ordered, i.e. $w.VC \not\leq w'.VC$ and $w'.VC \not\leq w.VC$. On a write event, the algorithm always compares it with the previously observed write on the shared variable. If the previous write event had written the same value and the new and previous writes are not ordered (line 19), then we clear the map entry holding the last write (line 22) so that a future read cannot merge with it. To guarantee the correctness of the matching, we push the last write on x into \mathbb{W} if there are no conflicting writes (line 20). Hence, a *read* will either match its write and merge, or not. This might have consequences on the faithfulness of the trace and not its soundness. Faithfulness will only be affected if (1) three or more threads are writing to the same variable the same value and in a concurrent region, and (2) the program relies on those writes to synchronize. We reasonably believe that this is an infrequent case in concurrent programs. However, for such marginal cases, the algorithm can report conflicting writes. Writes can then be instrumented to have advice atomicity forcing a total order on their execution; by adding synchronization blocks that wrap them with their advice in mutually exclusive regions. When this applies, there will be no need for the procedure `ReadOrWrite`, and reads and writes can be treated as *acquire* and *release* where they are handled with procedure `ReleaseOrAcquire`.

6 Criteria For Monitorability

In this section, we discuss the criteria for sound monitoring with automata and concurrent traces. In Sec. 4, we see that a sound and faithful trace represents a program execution and can be used interchangeably when verifying a property. However, concurrent traces (equivalently concurrent executions) might have unordered events. Many monitoring approaches rely on finite state automata as they can be used for most of the specification pattern [24,1]. Such monitors expect a total order of events as their input consists of words. Therefore, concurrent

⁶ Atomic operations such as *compare-and-swap* are handled differently since we check for the result of the operation before emitting an action.

traces must be linearized before monitoring. A *linearization* of a partial order will impose an arbitrary order between these unordered events. However, feeding the monitor with incorrect orderings of events can lead to incorrect verdicts. For the remainder of this section, we set Σ as the set of events over which properties are specified. We distinguish it from \mathcal{E} , which is the set of runtime events that will possibly be projected to events in Σ .

6.1 Monitor Causal Dependence

When observing an automaton, we might find pairs of events whose order is irrelevant to its progress; they can permute without affecting the verdict. The *causal dependence* relation $\mathcal{D} \subseteq \Sigma \times \Sigma$ is a binary relation that is anti-reflexive and symmetric. It contains all pairs of events whose correct order is necessary and their permutation would lead the automaton to a different state. Hence, for all pairs of events that do not belong to \mathcal{D} , their permutation after a linearization (if they occur concurrently) can be safely tolerated. Moreover, $\mathcal{I}_D = (\Sigma \times \Sigma) \setminus \mathcal{D}$ is the independence relation induced by the dependence relation \mathcal{D} . For a deterministic automaton, \mathcal{I}_D is a reflexive and symmetric relation. We extract \mathcal{I}_D by finding pairs of actions that would lead to the same state if permuted. An algorithm that extracts \mathcal{I}_D from a property specified as a DFA is fairly simple (see technical report [70]). It checks all pairs (a, b) in $\Sigma \times \Sigma$, if starting from any state in the automaton $a.b$ would lead to a different state than $b.a$. If so, then the automaton depends on receiving both symbols in order, and the pair is added to the dependency relation. Note that, the *causal dependence* relation resembles trace equivalence from [56] which constrains the allowed linearizations of a partial order. However, here we extract the relation from the monitor itself, and it defines *word* equivalence concerning an automaton, that is, the allowed permutations of letters in a word that would eventually lead to the same verdict.

Example 2 (Monitor Causal Dependence). We demonstrate with two example properties⁷. **P1** states that event s responds to p between q and r . **P2** is a mutual exclusion property that states that no read or write should happen concurrently with another write. To monitor with a finite state machine monitor, the read and writes are instrumented and delimited with events bw and aw represent before and after a write, br and ar for before and after a read. Fig. 3 shows the violation automata of bad prefixes to monitor the properties. For each automaton, we have:

- $\mathcal{I}_{D_1} = \{\{s, q\}\}$
- $\mathcal{I}_{D_2} = \{\{br, ar\}\}$

From \mathcal{I}_{D_1} , we see that the monitor does not depend on the order between s and q . From \mathcal{I}_{D_2} , we see that the monitor depends on the order between writes themselves, reads, and writes and reads.

⁷ In the technical report [70], we extract \mathcal{I}_D from 55 event-based property specifications from [24,1] written as Quantified Regular Expressions (QREs).




Fig. 3: Automata of Bad Prefixes.

6.2 Trace Monitorability of Concurrent Executions

We first define the notion of necessary order for a concurrent trace, which indicates whether the trace has the needed order based on the dependence relation.

Definition 6 (Trace Necessary Order). *We say that a concurrent trace $t = (\mathcal{E}, \xrightarrow{\text{tr}})$ has the necessary orderings w.r.t. a causal dependence relation \mathcal{D} , noted $\text{tno}(t, \mathcal{D})$ when: $\forall e, e' \in \mathcal{E} : \langle e, e' \rangle \in \xrightarrow{\text{tr}} \vee \langle \text{label}(e), \text{label}(e') \rangle \notin \mathcal{D}$*

Example 3 (Trace Necessary Order). Back to property **P2** from Example 2 and the traces depicted in Fig. 2. Trace t_3 is ordered enough for monitoring the property, whereas t_2 , a trace collected considering the thread order only, does not capture order between reads and writes. As such we have:

$$\text{tno}(t_3, \mathcal{D}_2) = \top \qquad \text{tno}(t_2, \mathcal{D}_2) = \perp$$

Let us recall the notion of monitorability from [42,13]. A property φ is *monitorable*, denoted by $\text{Mon}(\varphi)$, if every prefix of every trace has a finite extension that allows its monitor to reach a verdict, be it positive or negative. Monitoring with unsound traces leads to unsound verdicts. We redefine monitorability for concurrent programs by adding necessary conditions on the traces.

Definition 7 (Trace Monitorability of Concurrent Executions). *Given a property φ with its dependency relation \mathcal{D} , and a trace t collected from a concurrent execution e . Property φ is monitorable with t , noted $\mathbf{t}\text{-Mon}(\varphi)$ when $\text{Mon}(\varphi) \wedge \text{snd}(e, t) \wedge \text{tno}(t, \mathcal{D})$.*

First, the property φ should be monitorable in the classical sense, $\text{Mon}(\varphi)$, or else we will not reach a verdict. Second, the trace should be sound, $\text{snd}(e, t)$, or else we will have an unsound verdict. Third, the trace should have all the ordering information needed by the property as per its dependency relation, $\text{tno}(t, \mathcal{D})$, or else a linearization would produce an unsound trace. The above indicates that a concurrent trace does not need to contain all the ordering information between events and that the notion of faithfulness can be relaxed when monitoring. Now, if there is missing information in a sound and faithful concurrent trace, this means that the execution itself does not contain the needed

causality for monitoring the property. As such, the user is warned about the missing order to address the problem, and a tradeoff is presented between concurrency and monitorability. On one hand, they can synchronize the concurrent actions in the program to have them ordered, or they can force linearization of unordered actions via instrumentation, as discussed in Sec. 2. On the other hand, they can leave the actions executing concurrently and afford inconsistent verdicts.

7 Experimentation and Evaluation

We implemented our work in the tool **FACTS** (**F**aithful and **S**ound **C**oncurrent **T**races) for Java programs⁸. **FACTS** consists of three modules; the *instrumentation* module allows the user to specify the instrumentation logic and is built on top of BISM instrumentation framework [71]. The *trace reordering* module implements the vector clock algorithm (Sec. 5) which is responsible for timestamping the events. The *monitorability checker* module, takes as input the causal dependence relation \mathcal{D} . After receiving new events, it checks if the collected concurrent trace is monitorable as discussed in Sec. 6.

Experimental Setup We pick for our evaluation real-world Java applications from Renaissance [61] and DaCapo Benchmarks [16], and synthetic programs from [36]. We monitor properties that can be expressed with total order formalisms to show how concurrent traces can help existing monitoring approaches adapt to concurrent programs. We compare three concurrent trace collection approaches: with **FACTS** in both *asynchronous* and *synchronous* modes, and with Algorithm **A** from [65] (which cannot be run in asynchronous mode)⁹. We instrument the programs to collect the property-related events and the synchronization actions such as thread operations, synchronized blocks, and methods, locks, reads and writes to shared variables, and spawning actors [6] in Akka [2].

Benchmarks We monitor the program **akka-uct** for a response property stating that between the submission and the execution of a task with normal priority (events q and r resp.) if an urgent task is submitted (event p) it should execute (event s) in between q and r . The property is identical to **P1** from Ex. 2. For the independence relation, we have $\mathcal{I}_D = \{\langle s, q \rangle\}$. We extract it beforehand using Algorithm 2 from [70] and pass it to **FACTS** as part of the specification. We monitor program **future-genetic** to check whether dependent tasks execute in parallel. The property is similar to **P2** from Ex. 2 without q_3 and the *read* events. For the Dacapo benchmarks **avroora** and **fop**, we target the classical type-state properties [72]; the `SafeIterator` and `HasNext` properties. For both properties, \mathcal{I}_D is empty and all events are expected to be ordered. We also

⁸ Implementation details can be found in the technical report [70].

⁹ We also collect linear traces as collected by Java-MOP [21] and show the number of ordering corrections made with concurrent traces, details in [70].




Fig. 4: Experimentation Results.

monitor the Bakery locking algorithm [44], **bakery**, and a classic producer-consumer program, **prods-cons**. For both, we monitor a property similar to **P2** from Ex. 2, as such, a similar \mathcal{I}_D .

Results Fig. 4a reports the mean execution time for 20 runs of the benchmarked programs. We first note that FACTS is capable of producing sound and faithful traces from all benchmarks, with no marginal cases (i.e., conflicting writes) reported by our vector clock algorithm. This holds even for the *bakery* benchmark, which relies solely on reads and writes for the synchronization of multiple threads. Secondly, running our algorithm in asynchronous mode, FACTS (async), interferes minimally with the program as it incurs a considerably low overhead in most of the benchmarks, relieving the program from blocking when timestamping events and executing the monitors. Thirdly, for the comparable algorithms, FACTS (sync) performs better than Algorithm A in most of the benchmarks. We fairly believe that our algorithm interferes less with the programs as it imposes finer-grained synchronization than Algorithm A. This is highlighted with the *bakery* algorithm which synchronizes only using shared variables. Algorithm A requires the update of vector clocks associated with a read or write to be atomic through synchronization (as in Sec. 2), while our algorithm does not. This causes the threads to spin more as more contention is added with Algorithm A. Fourthly, for **future-genetic** and **akka-uct**, these programs use parallel tasks and message-passing (resp.) for managing concurrency. We can see how capturing concurrent traces synchronously from them interferes severely with their behavior. Algorithm A and FACTS (sync) timed out with **future-genetic**, while for **akka-uct**, Algorithm A is not intended to

handle message passing. Monitoring programs that utilize concurrency primitives with higher levels of abstraction will require improved adaptation in the future. For the time being, it is preferable to observe and monitor such programs asynchronously.

Table 4b reports on the monitorability of the collected traces. As per our approach, the monitor is warned when $\mathbf{t}\text{-Mon}$ is false, as it may produce unsound verdicts. To test the effectiveness of our approach when used with existing monitoring tools, we introduce two buggy implementations, namely **prods-cons** (**f**) and **bakery** (**f**), where synchronization between threads is faulty, and events can execute without acquiring locks. These programs are intended to have unordered events in the executions (and consequently, the trace), resulting in the trace not having the necessary orderings to monitor the above properties. For each benchmark, we run both a correct implementation where $\mathbf{t}\text{-Mon} = \top$, and a faulty one where $\mathbf{t}\text{-Mon} = \perp$, and monitor them using Java-MOP [21]. We report the results of 100 executions in Table 4c. Our findings reveal that monitoring with $\mathbf{t}\text{-Mon} = \perp$ yields inconsistent results, with some executions producing a *true* verdict (32%) while others yield a *false* verdict (68%). This inconsistency highlights the need for our approach, as FACTS produces warnings in all such executions, ensuring sound verdicts. In contrast, verdicts for correct programs are consistent because the execution of events is linearized, and the traces have the necessary orderings.

8 Related Work

We focus on *property-based* dynamic verification techniques for concurrent programs that rely on traces. More specifically, on techniques developed for monitoring behavioral properties expressed in total order formalisms and refer to [15] for a detailed survey. These techniques typically analyze a trace to either *detect* or *predict* violations.

Detection techniques reason about single runs of a program. We mention runtime monitoring tools, namely Java-MOP [21], Tracematches [8,17], MarQ [62], and LARVA [23] chosen from the RV competitions [28,63,11]. These tools allow different specification formalisms such as finite-state machines, extended regular expressions, context-free grammars, past-time linear temporal logic, and Quantified Event Automata (QEA) [10]. Detection techniques do not establish causal orderings between events and rely on trace collection approaches (discussed in Sec. 2) to order the collected events. We have shown in this paper how this can produce unsound traces leading to unsound and inconsistent monitoring. These tools can benefit from concurrent traces to guarantee the soundness of their verdicts. EnforceMOP [50] for instance, can be used to detect and enforce properties (deadlocks as well). It controls the runtime scheduler and blocks threads that might cause a property violation, sometimes leading to a deadlock. It requires forced atomicity as the scheduler needs to decide at each step if the execution on some thread continues or not. In [19,9], the authors present a monitoring framework for actor-based systems. The tool detectEr monitors Erlang applica-

tions using traces collected using the native logging functionality. Our approach targets generic concurrency primitives and can also be used with actor-based systems.

Predictive techniques reason about all feasible interleavings from a recorded trace of a single execution. Their challenge is to construct sound and maximal causal models [67] that allow exploring flexibly all feasible interleavings. In [65], the authors present an instrumentation algorithm based on vector clocks, used in [66,39,22] for generating the partial order from a running program. The algorithm maintains one vector clock for each thread and two for each shared variable. It executes synchronously with the executing program and is protected using *synchronized* blocks to force an overall sequentially consistent [43] execution. Vector clock algorithms typically require synchronization between the advice, program actions, and algorithm’s processing to avoid data races [18]. Our algorithm can run synchronously or asynchronously with the program depending on the monitoring scenario. As far as we know, it is unique in the context of online monitoring in establishing order off the critical path without the need to block the execution to process. In [39,33] the work targets type-state errors. jPredictor [22] for instance, uses sliced causality [20] to prune the partial order such that only relevant synchronization actions are kept. The tool is demonstrated on atomicity violations and data races; however, we are not aware of an application in the context of generic behavioral properties. In [29], the authors present ExceptioNULL that target null-pointer exceptions. Violations and causality are represented as constraints over actions, and the feasibility of violations is explored via an SMT constraint solver. GPredict [37], for instance, targets generic concurrency properties. It allows the user to express properties with regular expressions and provides explicit concurrency idioms such as atomic and parallel regions. It establishes order by collecting thread-local traces and also producing constraints over actions. In addition to being incomplete due to the possibility of not getting results from the constraint solver, the analysis from these tools might also miss some order relations between events resulting in false positives. Of course, none of the presented predictive techniques are complete, i.e., can produce all possible feasible interleavings that lead to violations, due to the impossibility of constructing a complete causal model of the program. Furthermore, these techniques reason on sequentially consistent execution models [36], restricting the space of possible interleavings of programs. The idea is that if a property is violated in a sequential consistency, then it will surely be violated in a more relaxed execution model. Our work focuses on providing concurrent traces for online detection techniques, and we have yet to explore their applicability in predictive contexts. Unfortunately, many tools from the mentioned approaches [37,22,66,39,22] are not available. Apart from tools dedicated to data races and atomicity violation detection, we found no available tools targeting general behavioral properties to compare with; tools that can establish causal order and instrument various custom regular actions from the program for monitoring.

9 Conclusion And Future Work

Collecting concurrent traces for runtime monitoring opens the path to interesting problems. Firstly, monitoring techniques can be revisited and extended to take into account the partial order. Tools relying on total order in traces can use concurrent traces to check if a trace has the needed causality and, if not, produce warnings. Moreover, specifications (and formalisms) that can match traces obtained from FACTS can be elaborated to extend the expressiveness of monitoring to check concurrency-related behavioral properties. For example, our approach is applicable for OpenMP runtimes; it can be used to verify the correctness of the scheduling algorithms of tasks with data dependencies (e.g., [73]). We can verify that the runtime never schedules two dependent tasks in parallel. Secondly, it is now possible to define and quantify optimizations for capturing sound and faithful traces. How to obtain optimal faithfulness with instrumentation is still an interesting challenge. Thirdly, for scenarios where the execution does not have the causality needed between events, the user can be given facilities to preserve some safety property by enforcing the order. Linearization of concurrent events may be achieved on the fly by JVM hot swapping.

Finally, we will reduce overheads induced by capturing synchronization actions. We aim to include plugins to FACTS to instrument concurrency primitives with higher-level abstraction such as fork-join [46] and software transactional memory [35]. Tailoring instrumentation to these frameworks reduces the number of collected synchronization actions. It requires assumptions about the correctness of these frameworks. These assumptions can be checked using static and dynamic analysis.

References

1. Patterns in property specifications for finite-state verification home page. <https://matthewbdwyer.github.io/psp/patterns.html>, <https://matthewbdwyer.github.io/psp/patterns.html>
2. Akka documentation. <http://akka.io/docs/> (2022), <http://akka.io/docs/>
3. 133, J.S.R.J.: Java Memory Model and Thread Specification Revision, 2004. <http://jcp.org/jsr/detail/133.jsp>
4. Adve, S.V., Gharachorloo, K.: Shared memory consistency models: a tutorial. *Computer* **29**(12), 66–76 (Dec 1996). <https://doi.org/10.1109/2.546611>
5. Agarwal, A., Garg, V.K.: Efficient dependency tracking for relevant events in shared-memory systems. In: Proceedings of the Twenty-Fourth Annual ACM Symposium on Principles of Distributed Computing. p. 19–28. PODC '05, Association for Computing Machinery, New York, NY, USA (2005). <https://doi.org/10.1145/1073814.1073818>, <https://doi.org/10.1145/1073814.1073818>
6. Agha, G.: *Actors: A Model of Concurrent Computation in Distributed Systems*. MIT Press, Cambridge, MA, USA (1986)
7. Ahamad, M., Neiger, G., Burns, J.E., Kohli, P., Hutto, P.W.: Causal memory: definitions, implementation, and programming. *Distributed Computing* **9**(1), 37–49 (Mar 1995). <https://doi.org/10.1007/BF01784241>

8. Allan, C., Avgustinov, P., Christensen, A.S., Hendren, L., Kuzins, S., Lhoták, O., de Moor, O., Sereni, D., Sittampalam, G., Tibble, J.: Adding Trace Matching with Free Variables to AspectJ. In: Proceedings of the 20th Annual ACM SIGPLAN Conference on Object-oriented Programming, Systems, Languages, and Applications. pp. 345–364. OOPSLA '05, ACM (2005). <https://doi.org/10.1145/1094811.1094839>
9. Attard, D.P., Cassar, I., Francalanza, A., Aceto, L., Ingólfssdóttir, A.: A runtime monitoring tool for actor-based systems (2022), <https://www.um.edu.mt/library/oar/handle/123456789/23062>
10. Barringer, H., Falcone, Y., Havelund, K., Reger, G., Rydeheard, D.E.: Quantified Event Automata: Towards Expressive and Efficient Runtime Monitors. In: FM 2012: Formal Methods - 18th International Symposium, Paris, France, August 27-31, 2012. Proceedings. pp. 68–84 (2012). https://doi.org/10.1007/978-3-642-32759-9_9, https://doi.org/10.1007/978-3-642-32759-9_9
11. Bartocci, E., Falcone, Y., Bonakdarpour, B., Colombo, C., Decker, N., Havelund, K., Joshi, Y., Klaedtke, F., Milewicz, R., Reger, G., Rosu, G., Signoles, J., Thoma, D., Zalinescu, E., Zhang, Y.: First international competition on runtime verification: rules, benchmarks, tools, and final results of CRV 2014. International Journal on Software Tools for Technology Transfer (Apr 2017). <https://doi.org/10.1007/s10009-017-0454-5>
12. Bartocci, E., Falcone, Y., Francalanza, A., Reger, G.: Introduction to runtime verification. In: Bartocci, E., Falcone, Y. (eds.) Lectures on Runtime Verification - Introductory and Advanced Topics, Lecture Notes in Computer Science, vol. 10457, pp. 1–33. Springer (2018). https://doi.org/10.1007/978-3-319-75632-5_1
13. Bauer, A., Leucker, M., Schallhart, C.: Runtime verification for ltl and tltl. ACM Trans. Softw. Eng. Methodol. **20**(4), 14:1–14:64 (Sep 2011). <https://doi.org/10.1145/2000799.2000800>
14. Bensalem, S., Havelund, K.: Dynamic deadlock analysis of multi-threaded programs. In: Proceedings of the First Haifa International Conference on Hardware and Software Verification and Testing. p. 208–223. HVC'05, Springer-Verlag, Berlin, Heidelberg (2005). https://doi.org/10.1007/11678779_15
15. Bianchi, F.A., Margara, A., Pezzè, M.: A survey of recent trends in testing concurrent software systems. IEEE Transactions on Software Engineering **44**(8), 747–783 (2018). <https://doi.org/10.1109/TSE.2017.2707089>
16. Blackburn, S.M., Garner, R., Hoffmann, C., Khang, A.M., McKinley, K.S., Bentzur, R., Diwan, A., Feinberg, D., Frampton, D., Guyer, S.Z., Hirzel, M., Hosking, A., Jump, M., Lee, H., Moss, J.E.B., Phansalkar, A., Stefanović, D., VanDrunen, T., von Dincklage, D., Wiedermann, B.: The DaCapo Benchmarks: Java Benchmarking Development and Analysis. In: Proceedings of the 21st Annual ACM SIGPLAN Conference on Object-oriented Programming Systems, Languages, and Applications. pp. 169–190. OOPSLA '06, ACM (2006). <https://doi.org/10.1145/1167473.1167488>
17. Bodden, E., Hendren, L., Lam, P., Lhoták, O., Naeem, N.A.: Collaborative Runtime Verification with Tracematches. Journal of Logic and Computation **20**(3), 707–723 (Jun 2010). <https://doi.org/10.1093/logcom/exn077>
18. Cain, H.W., Lipasti, M.H.: Verifying sequential consistency using vector clocks. In: Proceedings of the Fourteenth Annual ACM Symposium on Parallel Algorithms and Architectures. p. 153–154. SPAA '02, Association for Computing Machinery, New York, NY, USA (2002). <https://doi.org/10.1145/564870.564897>, <https://doi.org/10.1145/564870.564897>

19. Cassar, I., Francalanza, A.: On implementing a monitor-oriented programming framework for actor systems. In: Integrated Formal Methods - 12th International Conference, IFM 2016, Reykjavik, Iceland, June 1-5, 2016, Proceedings. pp. 176–192 (2016). https://doi.org/10.1007/978-3-319-33693-0_12, https://doi.org/10.1007/978-3-319-33693-0_12
20. Chen, F., Roşu, G.: Parametric and sliced causality. In: Proceedings of the 19th International Conference on Computer Aided Verification. p. 240–253. CAV’07, Springer-Verlag, Berlin, Heidelberg (2007)
21. Chen, F., Roşu, G.: Java-MOP: A Monitoring Oriented Programming Environment for Java. In: Tools and Algorithms for the Construction and Analysis of Systems. pp. 546–550. Lecture Notes in Computer Science, Springer (Apr 2005). https://doi.org/10.1007/978-3-540-31980-1_36
22. Chen, F., Serbanuta, T.F., Rosu, G.: Jpredictor: A predictive runtime analysis tool for java. In: Proceedings of the 30th International Conference on Software Engineering. p. 221–230. ICSE ’08, Association for Computing Machinery, New York, NY, USA (2008). <https://doi.org/10.1145/1368088.1368119>, <https://doi.org/10.1145/1368088.1368119>
23. Colombo, C., Pace, G.J., Schneider, G.: LARVA — Safer Monitoring of Real-Time Java Programs (Tool Paper). In: Hung, D.V., Krishnan, P. (eds.) Seventh IEEE International Conference on Software Engineering and Formal Methods, SEFM 2009, Hanoi, Vietnam, 23-27 November 2009. pp. 33–37. IEEE Computer Society (2009). <https://doi.org/10.1109/SEFM.2009.13>, <https://doi.org/10.1109/SEFM.2009.13>
24. Dwyer, M.B., Avrunin, G.S., Corbett, J.C.: Patterns in property specifications for finite-state verification. In: Proceedings of the 21st International Conference on Software Engineering. p. 411–420. ICSE ’99, Association for Computing Machinery, New York, NY, USA (1999). <https://doi.org/10.1145/302405.302672>, <https://doi.org/10.1145/302405.302672>
25. El-Hokayem, A., Falcone, Y.: Can we monitor all multithreaded programs? In: Colombo, C., Leucker, M. (eds.) Runtime Verification - 18th International Conference, RV 2018, Limassol, Cyprus, November 10-13, 2018, Proceedings. Lecture Notes in Computer Science, vol. 11237, pp. 64–89. Springer (2018). https://doi.org/10.1007/978-3-030-03769-7_6, https://doi.org/10.1007/978-3-030-03769-7_6
26. Falcone, Y., Fernandez, J., Mounier, L.: What can you verify and enforce at runtime? *Int. J. Softw. Tools Technol. Transf.* **14**(3), 349–382 (2012). <https://doi.org/10.1007/s10009-011-0196-8>, <https://doi.org/10.1007/s10009-011-0196-8>
27. Falcone, Y., Havelund, K., Reger, G.: A Tutorial on Runtime Verification. In: Engineering Dependable Software Systems, pp. 141–175. IOS Press (2013)
28. Falcone, Y., Nickovic, D., Reger, G., Thoma, D.: Second international competition on runtime verification CRV 2015. In: Bartocci, E., Majumdar, R. (eds.) Runtime Verification - 6th International Conference, RV 2015 Vienna, Austria, September 22-25, 2015. Proceedings. Lecture Notes in Computer Science, vol. 9333, pp. 405–422. Springer (2015). https://doi.org/10.1007/978-3-319-23820-3_27
29. Farzan, A., Madhusudan, P., Razavi, N., Sorrentino, F.: Predicting null-pointer dereferences in concurrent programs. In: Proceedings of the ACM SIGSOFT 20th International Symposium on the Foundations of Software Engineering. FSE ’12, Association for Computing Machinery, New York, NY, USA (2012). <https://doi.org/10.1145/2393596.2393651>, <https://doi.org/10.1145/2393596.2393651>

30. Farzan, A., Parthasarathy, M., Razavi, N., Sorrentino, F.: Predicting null-pointer dereferences in concurrent programs (11 2012). <https://doi.org/10.1145/2393596.2393651>
31. Flanagan, C., Freund, S.N.: Atomizer: A dynamic atomicity checker for multi-threaded programs. *SIGPLAN Not.* **39**(1), 256–267 (jan 2004). <https://doi.org/10.1145/982962.964023>, <https://doi.org/10.1145/982962.964023>
32. Flanagan, C., Freund, S.N.: Fasttrack: Efficient and precise dynamic race detection. In: *Proceedings of the 30th ACM SIGPLAN Conference on Programming Language Design and Implementation*. p. 121–133. PLDI '09, Association for Computing Machinery, New York, NY, USA (2009). <https://doi.org/10.1145/1542476.1542490>, <https://doi.org/10.1145/1542476.1542490>
33. Gao, Q., Zhang, W., Chen, Z., Zheng, M., Qin, F.: 2ndstrike: Toward manifesting hidden concurrency typestate bugs. *SIGPLAN Not.* **46**(3), 239–250 (mar 2011). <https://doi.org/10.1145/1961296.1950394>, <https://doi.org/10.1145/1961296.1950394>
34. Gastin, P., Kuske, D.: Uniform satisfiability problem for local temporal logics over Mazurkiewicz traces. *Inf. Comput.* **208**(7), 797–816 (2010). <https://doi.org/10.1016/j.ic.2009.12.003>
35. Harris, T., Marlow, S., Peyton-Jones, S., Herlihy, M.: Composable memory transactions. In: *Proceedings of the Tenth ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*. p. 48–60. PPOPP '05, Association for Computing Machinery, New York, NY, USA (2005). <https://doi.org/10.1145/1065944.1065952>, <https://doi.org/10.1145/1065944.1065952>
36. Herlihy, M., Shavit, N.: *The Art of Multiprocessor Programming*, Revised Reprint. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1st edn. (2012)
37. Huang, J., Luo, Q., Rosu, G.: Gpredict: Generic predictive concurrency analysis. In: *37th IEEE/ACM International Conference on Software Engineering, ICSE 2015, Volume 1*. pp. 847–857 (2015). <https://doi.org/10.1109/ICSE.2015.96>
38. Huang, J., Meredith, P.O., Rosu, G.: Maximal sound predictive race detection with control flow abstraction. In: *Proceedings of the 35th ACM SIGPLAN Conference on Programming Language Design and Implementation*. p. 337–348. PLDI '14, Association for Computing Machinery, New York, NY, USA (2014). <https://doi.org/10.1145/2594291.2594315>, <https://doi.org/10.1145/2594291.2594315>
39. Joshi, P., Sen, K.: Predictive typestate checking of multithreaded java programs. In: *Proceedings of the 2008 23rd IEEE/ACM International Conference on Automated Software Engineering*. p. 288–296. ASE '08, IEEE Computer Society, USA (2008). <https://doi.org/10.1109/ASE.2008.39>, <https://doi.org/10.1109/ASE.2008.39>
40. Joshi, Y., Tchamgoue, G.M., Fischmeister, S.: Runtime verification of ltl on lossy traces. *Proceedings of the Symposium on Applied Computing* (2017)
41. Kiczales, G., Hilsdale, E., Hugunin, J., Kersten, M., Palm, J., Griswold, W.G.: Getting started with AspectJ. *Commun. ACM* **44**(10), 59–65 (2001)
42. Kupferman, O., Y. Vardi, M.: Model checking of safety properties. *Form. Methods Syst. Des.* **19**(3), 291–314 (Oct 2001). <https://doi.org/10.1023/A:1011254632723>, <https://doi.org/10.1023/A:1011254632723>
43. Lamport: How to make a multiprocessor computer that correctly executes multiprocess programs. *IEEE Transactions on Computers* **C-28**(9), 690–691 (1979). <https://doi.org/10.1109/TC.1979.1675439>
44. Lamport, L.: A new solution of dijkstra's concurrent programming problem. *Commun. ACM* **17**(8), 453–455 (aug 1974). <https://doi.org/10.1145/361082.361093>, <https://doi.org/10.1145/361082.361093>

45. Lamport, L.: Time, Clocks, and the Ordering of Events in a Distributed System. *Commun. ACM* **21**(7), 558–565 (Jul 1978). <https://doi.org/10.1145/359545.359563>, <https://doi.org/10.1145/359545.359563>
46. Lea, D.: A java fork/join framework. In: Proceedings of the ACM 2000 Java Grande Conference, San Francisco, CA, USA, June 3-5, 2000. pp. 36–43 (2000). <https://doi.org/10.1145/337449.337465>, <https://doi.org/10.1145/337449.337465>
47. Lea, D., Bloch, J., Midkiff, S., Holmes, D., Bowbeer, J., Peierls, T.: Jsr 166: Concurrency utilities. <https://www.jcp.org/en/jsr/detail?id=166>
48. Leucker, M., Schallhart, C.: A brief account of runtime verification. *The Journal of Logic and Algebraic Programming* **78**(5), 293–303 (May 2009). <https://doi.org/10.1016/j.jlap.2008.08.004>
49. Lodaya, K., Weil, P.: Rationality in algebras with a series operation. *Inf. Comput.* **171**(2), 269–293 (2001). <https://doi.org/10.1006/inco.2001.3077>
50. Luo, Q., Rosu, G.: Enforcemop: A runtime property enforcement system for multithreaded programs. In: Proceedings of International Symposium in Software Testing and Analysis (ISSTA'13). pp. 156–166. ACM (July 2013). <https://doi.org/http://dl.acm.org/citation.cfm?doid=2483760.2483766>
51. Manna, Z., Pnueli, A.: A hierarchy of temporal properties (invited paper, 1989). In: Proceedings of the Ninth Annual ACM Symposium on Principles of Distributed Computing. p. 377–410. PODC '90, Association for Computing Machinery, New York, NY, USA (1990). <https://doi.org/10.1145/93385.93442>, <https://doi.org/10.1145/93385.93442>
52. Manna, Z., Pnueli, A.: Temporal Verification of Reactive Systems: Safety. Springer-Verlag New York, Inc. (1995)
53. Manson, J., Pugh, W., Adve, S.V.: The Java Memory Model. In: Proceedings of the 32nd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages. pp. 378–391. POPL '05, ACM (2005). <https://doi.org/10.1145/1040305.1040336>
54. Mathur, U., Viswanathan, M.: Atomicity Checking in Linear Time Using Vector Clocks, p. 183–199. Association for Computing Machinery, New York, NY, USA (2020), <https://doi.org/10.1145/3373376.3378475>
55. Mattern, F.: Virtual time and global states of distributed systems. In: PARALLEL AND DISTRIBUTED ALGORITHMS. pp. 215–226. North-Holland (1988)
56. Mazurkiewicz, A.: Trace theory. In: Advances in Petri Nets 1986, Part II on Petri Nets: Applications and Relationships to Other Models of Concurrency. p. 279–324. Springer-Verlag, Berlin, Heidelberg (1987)
57. Mazurkiewicz, A.W.: Trace theory. In: Brauer, W., Reisig, W., Rozenberg, G. (eds.) Petri Nets: Central Models and Their Properties, Advances in Petri Nets 1986, Part II, Proceedings of an Advanced Course, Bad Honnef, Germany, 8-19 September 1986. Lecture Notes in Computer Science, vol. 255, pp. 279–324. Springer (1986). https://doi.org/10.1007/3-540-17906-2_30
58. Meenakshi, B., Ramanujam, R.: Reasoning about layered message passing systems. *Computer Languages, Systems & Structures* **30**(3-4), 171–206 (2004). <https://doi.org/10.1016/j.cl.2004.02.003>
59. Nielsen, M., Plotkin, G.D., Winskel, G.: Petri nets, event structures and domains, part I. *Theor. Comput. Sci.* **13**, 85–108 (1981). [https://doi.org/10.1016/0304-3975\(81\)90112-2](https://doi.org/10.1016/0304-3975(81)90112-2)
60. Peled, D., Havelund, K.: Refining the safety-liveness classification of temporal properties according to monitorability. In: Margaria, T., Graf, S., Larsen, K.G. (eds.) Models, Mindsets, Meta: The What, the How, and the Why Not?

- Essays Dedicated to Bernhard Steffen on the Occasion of His 60th Birthday. *Lecture Notes in Computer Science*, vol. 11200, pp. 218–234. Springer (2018). https://doi.org/10.1007/978-3-030-22348-9_14, https://doi.org/10.1007/978-3-030-22348-9_14
61. Prokopec, A., Rosà, A., Leopoldseder, D., Duboscq, G., Tâma, P., Studener, M., Bulej, L., Zheng, Y., Villazón, A., Simon, D., Würthinger, T., Binder, W.: Renaissance: A modern benchmark suite for parallel applications on the jvm. In: *Proceedings Companion of the 2019 ACM SIGPLAN International Conference on Systems, Programming, Languages, and Applications: Software for Humanity*. p. 11–12. *SPLASH Companion 2019*, Association for Computing Machinery, New York, NY, USA (2019). <https://doi.org/10.1145/3359061.3362778>, <https://doi.org/10.1145/3359061.3362778>
 62. Reger, G., Cruz, H.C., Rydeheard, D.E.: MarQ: Monitoring at Runtime with QEA. In: Baier, C., Tinelli, C. (eds.) *Tools and Algorithms for the Construction and Analysis of Systems - 21st International Conference, TACAS 2015, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2015, London, UK, April 11-18, 2015. Proceedings. Lecture Notes in Computer Science*, vol. 9035, pp. 596–610. Springer (2015). https://doi.org/10.1007/978-3-662-46681-0_55
 63. Reger, G., Hallé, S., Falcone, Y.: Third international competition on runtime verification - CRV 2016. In: Falcone, Y., Sánchez, C. (eds.) *Runtime Verification - 16th International Conference, RV 2016, Madrid, Spain, September 23-30, 2016, Proceedings. Lecture Notes in Computer Science*, vol. 10012, pp. 21–37. Springer (2016). https://doi.org/10.1007/978-3-319-46982-9_3
 64. Roemer, J., Genç, K., Bond, M.D.: Smarttrack: Efficient predictive race detection. In: *Proceedings of the 41st ACM SIGPLAN Conference on Programming Language Design and Implementation*. p. 747–762. *PLDI 2020*, Association for Computing Machinery, New York, NY, USA (2020). <https://doi.org/10.1145/3385412.3385993>, <https://doi.org/10.1145/3385412.3385993>
 65. Rosu, G., Sen, K.: An instrumentation technique for online analysis of multithreaded programs. In: *18th International Parallel and Distributed Processing Symposium, 2004. Proceedings*. pp. 268– (2004). <https://doi.org/10.1109/IPDPS.2004.1303344>
 66. Sen, K., Rosu, G., Agha, G.: Runtime safety analysis of multithreaded programs. *SIGSOFT Softw. Eng. Notes* **28**(5), 337–346 (sep 2003). <https://doi.org/10.1145/949952.940116>, <https://doi.org/10.1145/949952.940116>
 67. Serbanuta, T., Chen, F., Rosu, G.: Maximal causal models for sequentially consistent systems. In: *Runtime Verification, Third International Conference, RV 2012, Istanbul, Turkey, September 25-28, 2012, Revised Selected Papers*. pp. 136–150 (2012). https://doi.org/10.1007/978-3-642-35632-2_16, https://doi.org/10.1007/978-3-642-35632-2_16
 68. Shin, K., Ramanathan, P.: Real-time computing: a new discipline of computer science and engineering. *Proceedings of the IEEE* **82**(1), 6–24 (1994). <https://doi.org/10.1109/5.259423>
 69. Soueidi, C., El-Hokayem, A., Falcone, Y.: Opportunistic monitoring of multithreaded programs (2023), to appear in *Fundamental Approaches to Software Engineering - 26th International Conference, FASE 2023*
 70. Soueidi, C., Falcone, Y.: Sound Concurrent Traces for Online Monitoring Technical Report (1 2023). <https://doi.org/10.6084/m9.figshare.21940205>

71. Soueidi, C., Kassem, A., Falcone, Y.: BISM: bytecode-level instrumentation for software monitoring. In: Deshmukh, J., Nickovic, D. (eds.) Runtime Verification - 20th International Conference, RV 2020, Los Angeles, CA, USA, October 6-9, 2020, Proceedings. Lecture Notes in Computer Science, vol. 12399, pp. 323–335. Springer (2020). https://doi.org/10.1007/978-3-030-60508-7_18, https://doi.org/10.1007/978-3-030-60508-7_18
72. Strom, R.E., Yemini, S.: Typestate: A programming language concept for enhancing software reliability **12**(1), 157–171 (Jan 1986). <https://doi.org/10.1109/TSE.1986.6312929>, <https://doi.org/10.1109/TSE.1986.6312929>
73. Virouleau, P., Broquedis, F., Gautier, T., Rastello, F.: Using data dependencies to improve task-based scheduling strategies on NUMA architectures. In: Euro-Par 2016: Parallel Processing - 22nd International Conference on Parallel and Distributed Computing, 2016, Proceedings. Lecture Notes in Computer Science, vol. 9833, pp. 531–544. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-43659-3_39
74. Wang, L., Stoller, S.: Runtime analysis of atomicity for multithreaded programs. IEEE Transactions on Software Engineering **32**(2), 93–110 (2006). <https://doi.org/10.1109/TSE.2006.1599419>