

Runtime Verification for Decentralized and Distributed Systems *

Adrian Francalanza¹, Jorge A. Pérez², and César Sánchez³

¹ CS@ICT, University of Malta, Malta, adrian.francalanza@um.edu.mt

² University of Groningen and CWI, Amsterdam, The Netherlands,
j.a.perez@rug.nl

³ IMDEA Software Institute, Spain, cesar.sanchez@imdea.org

Abstract. This chapter surveys runtime verification research related to distributed systems. We report solutions that study how to monitor system with some distributed characteristic, solutions that use a distributed platform for performing a monitoring task, and foundational works that present semantics for decomposing monitors or expressing specifications amenable for distributed systems.

We will identify some characteristics that distinguish distributed monitoring from centralized monitoring, and characteristics that allow to classify distributed runtime verification works based on features of the executing platforms, the specification language and the system description. Then, we will use these characteristics to describe and compare the distributed runtime verification solutions proposed in the research literature.

Keywords: monitoring distributed systems, distributed monitoring, decentralized monitoring, monitor decomposition

1 Introduction

This chapter surveys works on runtime verification (RV) related to distributed computing systems. Distributed computing is the area of computer science devoted to the study of *distributed systems*: computational artifacts that run in execution units placed at different locations, and that exchange information using a communication infrastructure, such as a computer network [Coulouris, 2011; Garg, 2002; Attiya and Welch, 2004].

Since distributed systems encompass many different but related classes of systems, the terminology has not been uniformly used. We begin by clarifying what we mean in this chapter by different terms and conventions commonly used in distributed computing, particularly with respect to monitoring.

*Some reviewer suggested to change the title to include *taxonomy* or *classification*, perhaps “Runtime Verification for Decentralized and Distributed Systems: A classification”, but we are not sure about such a change (and the other reviewers did not object).

The computational units that form a distributed system are typically able to execute processes simultaneously, under true concurrency. Each computational unit can run more than one *process*, and independently manage a set of local resources, typically including local memory and a local clock. We call each of these computational units a *location*.

There are two large classes of distributed systems, according to the way in which processes communicate and synchronize: systems that can use shared memory, and systems that can only use some form of message passing as means of communication. It is nowadays widely accepted to refer to the former as *parallel systems* and to the latter as distributed systems, and here we follow this convention. Additionally, some systems assume the existence of a shared clock among the computational units, which is another usual classification criteria. When one assumes the existence of a shared clock, the distributed system is usually called synchronous or *decentralized system*. If the shared clock is not assumed then the system is called asynchronous distributed system or simply a *distributed system*. Sometimes the communication infrastructure within the distributed system is simple, as in the case of buses or broadcast communication, but it is often the case that the network *topology* is relevant for the study of a given class of distributed systems. We follow the convention that, unless specified otherwise, all execution units can talk to all other execution units directly.

In practice, components of distributed systems can fail independently. Locations are typically the units of failure, modeling crashes on the execution platform that cause all processes in the location to stop their execution. Moreover, messages in message passing systems can arrive out-of-order, be duplicated or lost, or experience unbounded delays. The nature of the failures and the high independence of failure between the different components is another factor of complexity when dealing with distributed systems. Unless stated, it is common in distributed systems to assume that the system under study presents no failures. We follow this convention here too.

Due to their concurrent nature and to the other aspects of distribution, it is well-known that distributed systems are notoriously difficult to design and reason about. Throughout the years researchers have proposed many techniques to increase the reliability of distributed algorithms and systems, including dynamic solutions. These efforts include the development of runtime verification techniques for distributed computing, which we report here. We will use *distributed runtime verification* to refer to the broad area of research that studies runtime verification in connection with distributed or decentralized systems. This includes the monitoring of distributed systems as well as the use of distributed systems for monitoring. Due again to these intrinsic difficulties, distributed runtime verification is a very active area of research and new results will be produced in the near future.

Terminology A distributed and decentralized monitoring setting is typically built from subsystems, which we identify with processes for the discussion in this chapter. We use P_1, P_2, \dots to refer to processes. Processes execute indepen-

dently and occasionally synchronize or communicate with each another via the underlying communication platform.

Processes are partitioned across locations, meaning that every process is located at exactly one location for any given instance. We use l, k, \dots to refer to locations. When two processes are co-located at the same location, we say that they are *local* to one another. Otherwise, we say that they are *remote*. Processes may interact and communicate with both local and remote processes. Remote communication is typically assumed to be more expensive than its local counterpart.

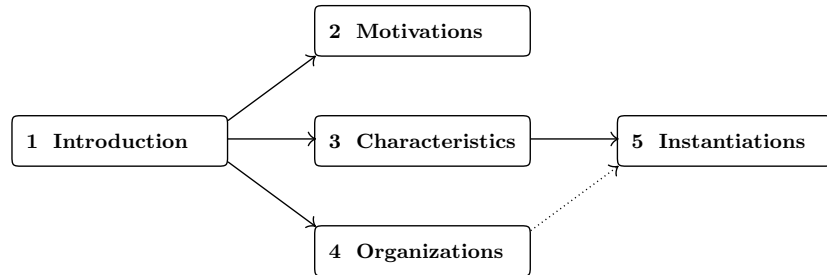
A *local trace* (or simply a *trace*), denoted T_1, T_2, \dots constitutes a log of past behavior used for monitoring purposes. A trace consists of a totally ordered set of trace *events*, each describing discrete computational steps of the monitored system. The ordering of trace events is necessary for the monitoring of temporal properties. A trace can describe events corresponding to a single process or else a group of processes. Although a particular location may host a number of traces (e.g., one per process hosted), we assume that a local trace cannot span across locations.

A monitoring task can be performed *online*, while the system under analysis is running, or *offline* by analyzing the log after the system has finished its execution. Here we consider both kinds of solutions.

In runtime verification, monitors are created from specifications, but we will use monitoring and runtime verification interchangeably in this chapter. In online runtime verification, *monitors*—denoted as M_1, M_2, \dots —are computing entities that check at runtime for the satisfaction or violation of correctness properties of the running system. Different monitors can be created to verify different properties simultaneously, and also in a modular fashion, which generally leads to better separation of concerns. The checking that the monitors perform is carried out by analyzing the traces generated by the executing processes. Similar to processes, monitors are hosted by a single location for any given instance. We allow monitors to analyze multiple traces in order to generate composite traces. We do not put restrictions on whether a monitor is allowed to analyze local or remote traces, but highlight the fact that remote trace analysis may carry additional overhead costs and entail higher security risks. Monitors are allowed to communicate with one another, which gives the flexibility for property checking to be carried out in a decentralized or choreographed manner (see Section 4.2 and Section 4.4).

The rest of the chapter is organized as follows. Section 2 presents a collection of reasons that have been proposed in the literature to motivate the study of distributed runtime verification problems. Section 3 identifies a number of characteristics that are relevant in the study of the solutions proposed; these characteristics serve as a basis to classify and compare the proposed solutions. Section 4 contains a description of the different ways to organize the activities carried out by the monitoring infrastructures. Section 5 describes a collection of solutions proposed in the literature, classified according to the attributes described in the preceding sections. Finally, Section 6 presents current challenges

and concluding remarks. The following diagram illustrates the dependencies between the sections.



Essentially, Section 5 contains the description and comparison of relevant work, using the classification characteristics extracted in Section 3.

2 Motivation and Scenarios

In this section we justify the study of distributed runtime verification. We present different scenarios from the point of view of the real-world problem that motivated the research reported. The list we present here is not intended to be exhaustive but its purpose is to give some practical justifications for the study of distributed runtime verification. Similarly, we do not claim that the papers cited are necessarily the first work to propose the study of a similar class of problems. The works mentioned below are further discussed in Section 5.

Observing Distributed Computations The obvious setting where distributed monitoring arises is when the system under scrutiny is itself distributed. One important problem related to observing distributed computations is that of detecting global predicates, which is recognized as an important problem since the early ages of distributed computing [Cooper and Marzullo, 1991].

It is known that checking general predicates is hard, since one has to store and enumerate all interleavings of the local processes. The so-called *computation slices* can be used for a more efficient detection [see Mittal et al., 2007; Alagar and Venkatesan, 2001; Chauhan et al., 2013]. Slices are abstractions of the distributed computation that guarantee that the predicate is present in a slice if occurred in some state of the original computation. If an algorithm is too general and does not exploit the structure of the predicate under consideration, predicate detection can involve a long runtime and large memory overhead [Chauhan et al., 2013]. Hence, best current solutions for predicate detection consider only fragments of the possible space of global predicates (for example the so-called linear, relational, regular and co-regular, and stable fragments) to gain efficiency. Even though most techniques for predicate detection [Cooper and Marzullo, 1991; Mittal et al., 2007; Alagar and Venkatesan, 2001] send all local events to a central

process for inspection of its interleavings, some modern approaches [Chauhan et al., 2013] consider purely distributed detection. Based on [Chauhan et al., 2013], Mostafa and Bonakdarpour [Mostafa and Bonakdarpour, 2015] adapt the work on predicate detection to detect LTL properties.

Analysis Decomposition Most approaches to runtime verification either consider the system under dynamic evaluation as a black-box, or only inspect the internals of the monitored system for the sake of instrumentation. However, it is common—using design principles like component-based design—that the description of the system is decomposed into different units.

For example, the work [Falcone et al., 2015b] investigates how to use the hierarchical description of the system to generate monitors that are then composed with the original system. This process produces a modified system that shares the original decomposition (and implements its functionality) and also includes the monitors embedded. Within this setting, the authors study how to compile a given design into either a centralized or a decentralized platform by deciding the placement of components using different deployment possibilities. Although the work in [Falcone et al., 2015b] does not specifically target distributed systems, the solution obtained from the compilation of the modified system can lead to a distributed monitoring solution if the target platform is distributed.

A similar approach is exemplified in [Cassar and Francalanza, 2016, 2015], where a framework for monitoring asynchronous component-based systems is presented. Again, the authors do not treat the system under scrutiny as a single monolithic block, but identify its constituent sub-components in the form of independently computing entities, called actors. The resulting monitoring setup generated is also localized to sub-components of the system, mirroring its non-monolithic structure. Even though actor systems are not necessarily executed in distributed fashion, the asynchronous nature of the code generated and its localization lends itself to a straightforward distribution.

Exploiting Parallelism Another justification for studying runtime verification in the context of distributed systems is the exploitation of parallel executing units to perform a monitoring task. For example, Berkovich et al. [Berkovich et al., 2015] propose to use additional hardware (a GPU parallel execution platform) to minimize the impact of online monitors on execution time overhead, reducing the intrusiveness. Moreover, the work [Francalanza and Seychell, 2013, 2015] shows performance gains in terms of lower overheads when monitors are specified as concurrent entities and executed over the prevalent multi-core and multi-processor architectures. This is because a monitor decomposition permits a better exploitation of the underlying processing units, as opposed to monolithic monitor descriptions.

Fault Tolerance Handling failures in distributed systems is challenging because different components can fail independently (e.g., nodes crashing) and the

communication can miss, duplicate or reorder messages or incur in unbounded delays [Francalanza and Hennessy, 2008]. At the same time, failure tolerance can be achieved by replicating components that perform a certain task, including monitors. For example, Fraigniaud et al. [Fraigniaud et al., 2014] study the problem of distributed monitoring with failures, where events can be observed from more than one monitor, but the nodes where the monitors execute can crash. The distributed monitoring algorithm then tries to reach a verdict among the surviving monitors.

The work in [Basin et al., 2013] targets the incomplete knowledge caused by network failures and message corruptions and attempts to handle the resulting disagreements. A subsequent work investigates how to handle network failures, and proposes algorithms that can reach verdicts when some information is missing and messages are reordered [Basin et al., 2015]. Since message losses are also considered, this approach can also model node crashes, which are simulated by all messages from the crashed node being lost.

Efficiency In many distributed systems scenarios, a simple monitoring solution can be obtained by implementing a central monitor that all other entities communicate with. However, distribution itself can be exploited to coordinate the monitoring task more efficiently. Many works attempt to provide more efficient solutions by exploiting the locality in the observations to also perform partially the monitoring task. For example, the works in [Falcone et al., 2015b; Cassar and Francalanza, 2015, 2016], already mentioned, exploit the hierarchical structure of the system to generate local monitors. On the other hand, the work in [Cassar et al., 2015] and in [Francalanza and Seychell, 2013, 2015] exploits the structure and semantics of the correctness property from which the monitors are synthesized to generate monitor organizations that use the underlying hardware efficiently. Concretely, the generated monitors minimize idle computing units and improve memory management via redundant monitor deallocations and monitor network reorganizations. These works pursue a more efficient monitoring where less communication and execution overhead is needed.

The pursuit of lowering overheads has also led Colombo et al. [Colombo et al., 2012] to consider distribution as a means of offloading part of the monitoring computation to the computing resources of another machine. They provide handles that allow the specifier to dictate whether a property is to be runtime-checked locally, as inlined code within the monitored system, or remotely via an independent monitoring unit located on a separate machine. In separate work [Colombo et al., 2015], the authors investigate various instrumentation techniques in Enterprise-Service Bus (ESB) distributed architectures, so as to determine which of them lead to lower monitoring overheads.

As observed in [Bauer and Falcone, 2012] and in [Francalanza et al., 2013], when atomic observations of the monitored system occur locally, one can organize the monitors hierarchically according to the structure of the original specification. This can lead to substantial savings in communication overheads because a verdict of a subformula can often be reached further down hierarchically. From

the practical point of view, the authors of [Bauer and Falcone, 2012] claim that many cyber-physical systems, like distributed systems found in the automotive and avionics industries, fulfill the requirement that both observations and their placement to local nodes are known at deployment time.

In the context of multithreaded programs with shared memory, the work of Luo and Roşu [Luo and Roşu, 2013] proposes to decompose a given property into local decentralized monitors for each of the threads, which again helps to reduce monitoring overheads.

Monitoring Expressivity Some approaches borrow directly monitoring languages from non-distributed computing, and study how to exploit or adapt the methods for distributed systems. Other approaches present new formalisms or extend existing ones with specific capabilities for distributed systems. For example, the work [Sen et al., 2004, 2006] proposes a method to check for violations of safety properties in distributed systems, using a variation of LTL that is suitable to describe (past time) properties of distributed systems. This extension essentially allows to express the knowledge of particular agents. The work in [Francalanza et al., 2013] proposes and formalizes a migrating monitor setup so as to better handle the open-ended and dynamic nature of distributed systems. This helps monitoring to adapt to locations that are learnt dynamically and to varying correctness specifications over the course of long-running distributed computations.

The efficiency of migrating monitors is investigated in [Bauer and Falcone, 2012] for fixed-location setups. The subsequent work [Colombo and Falcone, 2016] extends these results and compares them to choreographic solutions (see Section 4.4).

Testing and Enforcement Testing multithreaded programs is in general a challenging task because often concurrency errors arise only under specific interleavings and execution conditions, which are hard to cause and reproduce due to the non-determinism introduced by the scheduler. The work in [Luo and Roşu, 2013], already mentioned, presents an *enforcement* mechanism that exploits user-specified properties to generate local monitors that block individual threads whose execution may lead to a violation of the specified property. This approach intends (1) to force schedulers to explore properties during the testing of multithreaded programs; and (2) to automatically guarantee the enforcement of properties in a multi-threaded program, which is otherwise typically implemented using ad-hoc manual synchronization. The monitoring generation described in [Luo and Roşu, 2013] includes the decomposition of the property into local decentralized monitors for each of the threads.

3 Characteristics of Distributed Runtime Verification

In this section we capture some challenges that distributed systems impose on monitoring and the main difficulties that must be tackled by solutions to dis-

tributed runtime verification. We begin in Section 3.1 by describing some key characteristics of distributed system monitoring, particularly following a historical perspective. Overall, we consider 14 characteristics, denoted (C1)–(C14). Some of them (in particular (C1)–(C5)) are common to most distributed monitoring solutions, but are not typically a concern for non-distributed systems. Other criteria are not oblivious to all distributed monitoring cases, but identify aspects that will allow us to extract some classification dimensions, according to the approach taken by each solution. Most of these characteristics are also either unique to distributed systems or more challenging and important in distributed systems than in non-distributed systems. The classification aspects are listed later in Section 3.2.

3.1 Common Characteristics

Already in the late 1980s, Joyce et al. [Joyce et al., 1987] identified five issues in monitoring distributed systems, in an early attempt to characterize the key constraints that distinguish monitoring in sequential settings from monitoring in distributed systems:

- (C1) The fact that distributed systems have *many foci of control*;
- (C2) The presence of *communication delays* among nodes, which makes it difficult to determine a system’s state at any given time;
- (C3) The inherent *non-determinism* in distributed and asynchronous systems;
- (C4) The fact that monitoring a distributed system *alters its behavior*;
- (C5) The *complexity of the interactions* between the system and the system developer.

Aspect (C1) captures the idea that a distributed system is composed of processes running independently in distributed execution units. Issue (C2) refers to one of the aspects of message passing systems. We will later refer to this aspect that allows to distinguish between systems that are not synchronized (see **Global Clock** below) and where messages can be unboundedly delayed or be lost (see **Failures** below). Not all current research in distributed monitoring assumes that messages can suffer independent delays. Issue (C3) refers to the non-deterministic and asynchronous nature common to many distributed systems. Issue (C4) refers to the intrusiveness of monitoring in the system under analysis, which is not a unique characteristic of monitoring distributed systems. We consider here intrusiveness as a key characteristic (see **Intrusiveness** below). Finally, issue (C5) refers to the additional complexity (when compared with non-distributed systems) for the engineer exercising the monitoring infrastructure, in terms of deploying the monitors and collecting and analyzing the reported data. We do not develop (C5) further in this chapter as we focus runtime verification, and not on software engineering aspects.

Another work that explores monitoring distributed systems and identifies common and classifying criteria is [Francalanza et al., 2011, 2013], where the following characteristics are extracted:

- (C6) Difficulties in keeping a *global state*;
- (C7) Confidentiality of the information collected and communicated;
- (C8) Trace analysis *locality*;
- (C9) *Dynamic* aspects of specifications;
- (C10) Locations constitute *units of failure*.

Maintaining a global state in a distributed system under observation is impractical for several reasons, captured by aspect (C6). One reason is that sometimes it is even theoretically impossible to build and maintain a global view, due to the lack of global clocks, asynchrony, message loss and reordering, etc. Even when it is theoretically possible, it is common that the volume of event messages that are required to build such a global view would substantially increase the monitoring overhead, making it impractical. Most works recognize that although such a central solution would greatly simplify monitoring, it is either too complex or too intrusive. This difficulty will be captured as **Global Clock** and **Failures** below.

Aspect (C7) is related to security. Every time a trace of events is communicated across locations, the confidentiality of the information contained may be compromised. Solutions that encode and decode this information can further increase the monitoring overhead. However, we will not discuss this security aspects in this chapter.

Aspect (C8) refers to where the monitors are placed and where the events from the observed system are collected. Ideally, local monitors should analyze events locally and then communicate analysis summaries across locations. On the other hand, placement sometimes involves additional restrictions. For example, certain locations may not allow monitoring to be carried out locally due to resource constraints. Placement is often at odds with locality, which sometimes involves dynamic aspects. There are cases when it is difficult to anticipate the location where certain computations will be executed because this location depends on some runtime information that is hard to infer statically. Aspect (C8) is related to the distribution of the monitoring process, and in particular refers to the preference of decentralizing it (see **Centralization** below).

Aspect (C9) considers that in long-running applications without a central authority, correctness specifications may not be all available prior to deployment. Some specifications are added at runtime, while the system is already executing, which disables the static placement of monitors. Dynamic aspects of monitoring are considered in (C8) and (C9), caused by either unpredictable aspects at deployment time, or constraints in the execution platform which restrict installing monitors dynamically. Finally, aspect (C10) considers again the issue of failures (see **Failures** below).

In a recent short paper, [Bonakdarpour et al., 2016b], the authors discuss the following four issues as distinctive, characteristic challenges of distributed runtime verification:⁴

⁴The distributed RV considered in [Bonakdarpour et al., 2016b] is a general monitoring solution that runs on an infrastructure that is unreliable and unable to solve consensus.

- (C11) *Modeling* a distributed RV system;
- (C12) Defining and evaluating distributed correctness *specifications*;
- (C13) Using different *verdicts* on the state of the monitored system;
- (C14) Giving semantics to the different verdicts.

Issue (C11) concerns both the actual implementation of a distributed systems’s description (including whether it is used in the monitoring process, see issue **Exploiting System Description** below), as well as efforts devoted to describing the monitoring solutions (see Section 4). Aspect (C12) is related to the formalism used to describe monitors (see **Distributed Specifications**). Finally, the last two issues (C13) and (C14) are more specific to the solution provided in [Bonakdarpour et al., 2016b]. The first issue (C13), is already witnessed in monitoring non-distributed systems using LTL₃ [Bauer et al., 2011], where the semantics of LTL in finite traces is expressed using a 3 valued domain (the third value captures the possibility of expressing an *unknown* verdict, which may become later true or false when new observations are made). Issue (C13) refers to the use of multi-valued domains as verdicts emitted the local monitors in the distributed systems. Issue (C14) refers to how these multiple verdicts can be combined during the creation of a final verdict.

3.2 Distinguishing Characteristics

We now list six dimensions that will allow us to distinguish the different lines of research and classify the solutions proposed.

Exploiting System Description Most work in RV focuses on building monitors that can analyze any system, that is, the system is consider as a black-box that emits the necessary signals to the monitor. On the other hand, some other approaches exploit the system’s description to generate specialized monitors. Examples of system’s descriptions proposed include models of the system, abstractions or even full descriptions as programs. In this case, the monitors generated are only guaranteed to be correct for the specific system analyzed, and in case a different system is finally deployed with the monitor, the verdicts of this monitor may not be correct. On the other hand, solutions that consider the system as a black-box generate monitors that are correct for every system at the price of potentially less efficiency. For example, algorithms that generate monitors as finite state systems from LTL specifications work for all systems as sources of traces. If the monitor can rule out certain paths using concrete facts of the system under observation, obtained by static analysis for example, then the monitor can be specialized into a smaller finite state machine.

In some cases only certain aspects of the system description are used to build the solution, like the number of distributed nodes, the location of the individual predicates emitted by the running system, or the topology of the network.

Centralization Even if the system under observation is intrinsically distributed, the monitoring task can be performed in a central location that collects information from the remote units. However, solutions with a *central monitor* have many drawbacks from the points of view of overhead, efficiency, tolerance to faults and security. For these reasons, many solutions attempt to divide the monitors into local monitors and perform part of the monitoring activities locally, in a distributed fashion.

Global Clock There are two large classes of distributed runtime verification techniques depending on whether it is assumed that all nodes have access or not to a global clock (or a perfectly synchronized local clock). In case a global clock is assumed, the system under analysis is equivalent to a synchronous system (following distributed computing terminology). In this case, we call the problem *decentralized monitoring*. Similarly, when monitors do not have access to a global clock we refer to the problem as *distributed monitoring*. Another characteristic feature of monitoring distributed systems is *asynchrony*, both between the monitors and the distributed system under scrutiny, and among the distributed monitors themselves.

Monitoring a distributed system often amounts to monitoring a message passing system. We reserve the term non-distributed systems for those systems that have a global clock and direct access shared memory between all computational units. For example, parallel systems (as defined above) are non-distributed systems with several concurrent execution units.

Distributed Specifications One key classification criteria is whether the specification language from which monitors are generated has specific features for distributed systems, that is whether the formalism allows to refer to characteristics of the distributed platform. Some approaches borrow directly a language originally proposed for non-distributed systems, like LTL, and attack the problem of monitoring distributed systems against specifications written in this language. Other approaches start by introducing a modified specification language with some distributed feature, and then develop specific monitoring algorithms for this language.

Failures In practice, both non-distributed and distributed systems are subject to failures. However, failures in distributed systems can be more subtle than in non-distributed systems due to the physical independence of the executing units. Even though most monitoring solutions assume that no component can fail, some approaches consider the possibility of some part of the distributed system failing. In particular, some of the failing aspects considered are network delays in the transmission of the messages, message loss or duplication, message corruption and node crashes. Even though Byzantine failures have been thoroughly studied in distributed systems, this aspect has received little attention in the area of monitoring distributed systems.

Intrusiveness As already identified in early surveying efforts (see (C4) above), the monitoring process typically modifies the behavior of the monitored system. Naturally, most works focus on the effectiveness of the monitoring solution proposed, that is, on proving that the monitoring process actually detects the intended property. Some research also considers the efficiency of the combined solution (in terms of running time, number of messages, etc) and in some few cases how the monitoring process affects the running system (that is, how intrusive monitoring is). Moreover, some works are intrusive *on purpose*, trying to reduce the intrinsic non-determinism of the running system with the goal of avoiding failures (like in enforcement) or provoking failures (for testing purposes).

4 Monitor Organizations

In this section we explain and compare the various ways in which monitoring distributed system activities can be organized. The various monitoring organizations can be explained in terms of the different configurations used to compose these components together as a monitoring infrastructure contributing towards a common goal.

The analysis of correctness properties concerning different processes, possibly spanning across different locations, often requires the aggregation of traces into composite traces. We will generally assume that the composition of two remote traces does not necessarily yield a total ordering among the events of the resulting composite trace, but instead gives a partial ordering. Monitors can communicate with each other to coordinate the monitoring task.

4.1 Traditional Monitoring

A traditional monitoring setup, depicted in Figure 1, typically consists of a group of processes (P_1 , P_2 and P_3 in the figure) that reside at one location (l). These processes generate a single local trace (T_1) that is analyzed by a single monitor (M_1), also located at the same location. Even if these processes execute concurrently and are subject to a different interleaving every time the system is executed, the monitoring setup will always report a trace with a total ordering of events reflecting the executed interleaving.

4.2 Decentralized Monitoring

As depicted in Figure 2, a decentralized monitoring setup resembles traditional monitoring in that all process executions and trace events are governed by a single global clock. Moreover, processes and monitors can communicate using synchronous channels, and computations are totally ordered. Consequently, traces can also be totally ordered, either explicitly as one data structure or locally via using time-stamps.

In contrast to traditional monitoring which is typically performed by a single monolithic monitor, monitoring in a decentralized and distributed setup is

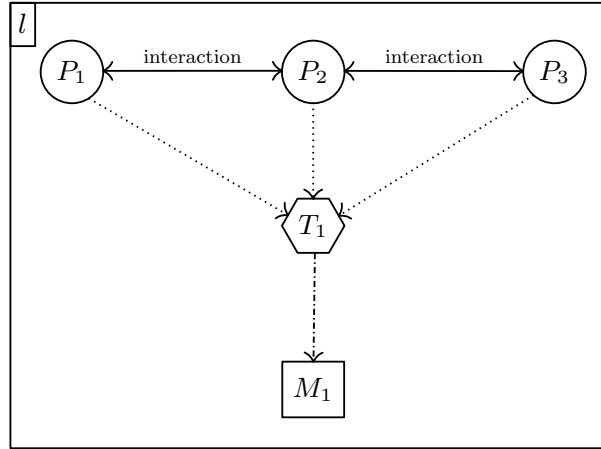


Fig. 1. A traditional monitoring setup where processes generate a single trace observed by a single monitor.




Fig. 2. In a decentralized monitoring setup the synchronized trace can be processed by several independent monitors that share a global clock.

decomposed into different sub-components (M_1, M_2, M_3) reflecting the fact that a global correctness property may be decomposed into smaller properties. For instance, in cases where trace events may be attributed to different system units (e.g. classes or objects), each monitor may selectively analyze events pertaining to a particular unit entity (e.g. all the method call invocations on a particular object) and then communicate aggregate monitoring information to other monitors in order to verify a global property. It is common that sub-monitors reflect some decomposition of the specification, but sometimes sub-monitors are obtained directly by the placement of parts of the specification into locations without much decomposition.




Fig. 3. In an orchestrated monitoring setup, traces are independently produced locally at the location of processes, but can be processed by remote monitors.

There are also cases in which the correctness properties are inherently separate and concern only a subset of processes as in the case of parametric monitoring where the property can be evaluated independently for different parameter instances (see, e.g., [Chen and Roşu, 2009]). In this case, monitoring may be decentralized in a natural manner without the need for the individual monitors to communicate among themselves.

4.3 Orchestrated Monitoring

Orchestrated and choreographed monitoring approaches are used in settings where more than one process is dispersed across more than one location. The set of processes generate more than one trace that can only be partially ordered due to the absence of a global clock.

In an orchestration all monitoring is ultimately performed *centrally* by a single monitor, accessing the respective trace events from different locations. The approach is depicted in Figure 3, which shows two sub-systems located at l and k , each producing local traces of events (T_1, T_2 and T_3 respectively), subsequently analyzed by monitors M_1, M_2 and M_3 from a remote location g . Each of these monitors analyze an independent correctness property.

On the one hand, the centralization of the analysis simplifies the logic of the monitor, which is conducive to a decrease in errors in the monitor code itself. However, these benefits come at a cost in distributed settings such as the one depicted in Figure 3. First, the approach leads to a substantial increase in the volume of trace information which has to be transmitted remotely for monitoring. The considerable increase in communication overhead across locations affects monitoring scalability when the number of processes and locations increases. The approach is also susceptible to data exposure when the trace events transmitted across locations contain private information. Adding additional security layers via mechanisms such as encryption further increases the monitoring overhead. Finally, the architecture poses a security risk by exposing

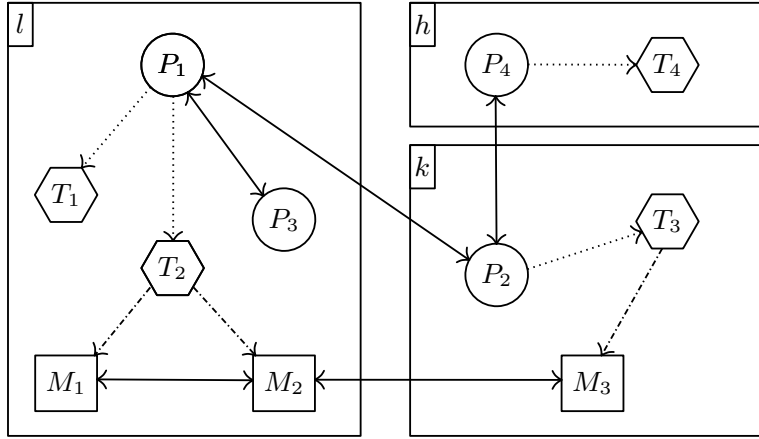


Fig. 4. In an choreographed monitoring setup, traces are independently produced locally at the location of processes and can be processed by different monitors independently.

the monitor as a central point of attack from which sensitive information can be tapped. Nevertheless, restricted forms of the orchestrated monitoring approach can be suitable when dealing with public information that is communicated over a relatively safe medium.

4.4 Choreographed Monitoring

A choreographed monitoring approach also targets system settings consisting of multiple processes dispersed across more than one location. In contrast to orchestrated monitoring, choreography-based approaches push the runtime verification activities locally to the location where the traces are generated, as shown in Figure 4. The diagram depicts four processes, located at three locations l , k , and h , each generating local traces, with monitors M_1 , M_2 placed at l , and M_3 placed at k . The monitor decomposition is not only due to the independence of the correctness properties being checked. In fact, monitors M_2 and M_3 could be verifying the same global property and eventually interact with each other in order to synchronize their monitoring effort.

The appeal of localizing monitoring is the potential minimization of data exposure and communication overhead. By verifying locally, we avoid having to transmit trace information to a remote monitor. Moreover, localized monitors typically require less communication than remote monitoring using a central monitor. However, choreography is more complex to instrument, since correctness properties need to be decomposed into coordinated local monitors. Furthermore, choreographed monitoring is also more intrusive, by burdening the monitored subsystems with additional local computation, and is thus applicable only when the hosting locations allow local instrumentation of monitoring code.

5 Instantiations

In this section we describe and compare research solutions proposed in the literature, using the characteristics captured in Section 3 and the organizations described in Section 4. To ease the description we group the papers as follows:

- Decentralized Monitoring
- Distributed Monitoring
- Fault Tolerance
- Monitor Decomposition
- Predicate Detection and Distributed Monitoring
- Intrusiveness
- Behavioral Type Systems for Distributed Monitoring

Other features cross-cut papers across different characteristics and are mentioned in each particular case. Tables 1 and 2 summarize the papers according to the main characteristics considered. In the characteristics shown in the tables, **Global Clock**, **Failures**, **System Description**, **Distributed Specs** and **Intrusiveness** are directly characteristics captured in Section 3. **Asynchronous Msgs** refers to whether the underlying platform is a message passing system. **Asynchronous Msgs**, **Decentralized Monitoring** and **Distributed Monitoring** are characteristics considered within **Global Clock** in Section 3.

The entries **LTL** and **Predicate Detection** are included because these specification languages have been thoroughly considered in many works. Finally, **Types** refer to a line of research based on process algebras and session types.

Decentralized Monitoring The works [Bauer and Falcone, 2012, 2016] study the problem of decentralized monitoring. The starting point is a specification expressed in LTL without any specific extension for distributed systems, except for the static mapping of atomic predicates to individual processes. Note that individual state predicates of the specification may be split into more than one process. The solution synthesizes a monitor for each process, under the assumption that components communicate synchronously with a global clock. Each component has a local monitor attached, and emits events synchronously after every global clock tick. By design, the solution to a verdict is taken as combination of the execution of the local monitors, lacking a central decision-making point.

The main advantage of a decentralized solution over a non-distributed one is that not all events must be sent to the location of the central monitor. The challenge is that local monitoring must be performed with only partial observations of the global trace. The algorithm progresses by rewriting the specification at each node, with the partial information available. When local monitors are unable to evaluate a specification given their local view of the computation, they communicate their residual formulas to the other monitors. An alternative approach would use a central monitor that receives information about the local states of all other locations. One of the main practical concerns is how the decentralized approach compares with this alternative central approach. The empirical

Publications	Global Clock	Asynchronous Msgs	Decentralized Monitoring	Distributed Monitoring	Failures	System Description	Distributed Specs.	LTL	Predicate Detection	Intrusiveness	Types
[Bauer and Falcone, 2012]	✓	.	✓	.	.	.	✓
[Bauer and Falcone, 2016]	✓	.	✓	.	.	.	✓
[Colombo and Falcone, 2014]	✓	.	✓	.	.	.	✓
[Colombo and Falcone, 2016]	✓	✓	✓	.	.	.	✓
[Bartocci, 2013]	✓	✓	✓	.	.	.	✓
[Sen et al., 2004]	.	✓	.	✓	.	.	✓
[Fractalanza et al., 2011]	.	✓	.	✓	.	✓	✓
[Fractalanza et al., 2013]	.	✓	.	✓	.	✓	✓
[Basin et al., 2013]	✓	.	.	.	✓
[Basin et al., 2015]	✓	✓	.	✓	✓
[Fraigniaud et al., 2014]	.	✓	.	.	✓	.	✓
[Bonakdarpour et al., 2016a]	.	✓	.	.	✓	.	✓
[Falcone et al., 2015b]	✓
[Bonakdarpour et al., 2010a]	.	✓	.	.	.	✓
[Bonakdarpour et al., 2010b]	.	✓	.	.	.	✓
[Berkovich et al., 2015]	✓	✓	.	✓	.	.
[Fractalanza and Seychell, 2013]	.	✓	.	✓
[Fractalanza and Seychell, 2015]	.	✓	.	✓
[Attard and Fractalanza, 2016]	.	✓	.	✓
[Chase and Garg, 1998]	.	✓	✓	.	.	.
[Cooper and Marzullo, 1991]	.	✓	✓	.	.	.
[Garg and Waldecker, 1994]	.	✓	✓	.	.	.
[Garg and Mittal, 2001]	.	✓	✓	.	.	.
[Mittal and Garg, 2005]	.	✓	✓	.	.	.
[Mittal et al., 2007]	.	✓	✓	.	.	.
[Sen and Garg, 2007]	.	✓	✓	.	.	.
[Ogale and Garg, 2007]	.	✓	✓	✓	.	.	.
[Chauhan et al., 2013]	.	✓	.	✓	.	.	.	✓	✓	.	.
[Mostafa and Bonakdarpour, 2015]	.	✓	.	✓	.	.	✓	✓	.	.	.
[Sen and Garg, 2003a]	.	✓	✓	✓	.	.	.
[Luo and Roşu, 2013]	✓	.	✓	.	.	✓	.	.	✓	.	.
[Cassar and Fractalanza, 2014]	.	✓	.	✓	✓	.	.
[Zhang et al., 2016]	.	✓	.	✓	✓	.	.
[Colombo et al., 2012]	.	✓	✓	.	.
[Colombo et al., 2011]	.	✓	.	✓	✓	.	.
[Cassar and Fractalanza, 2016]	.	✓	.	✓	✓	.	.
[Bocchi et al., 2013]	.	✓	.	✓	.	.	✓	.	.	.	✓
[Bocchi et al., 2017]	.	✓	.	✓	.	.	✓	.	.	.	✓
[Hu et al., 2013]	.	✓	.	✓	.	.	✓	.	.	.	✓

Table 1. State-of-the-art on distributed monitoring. Each paper is classified according to the characteristics considered (part 1) .

Publications	Global Clock	Asynchronous Msgs	Decentralized Monitoring	Distributed Monitoring	Failures	System Description	Distributed Specs.	LTL	Predicate Detection	Intrusiveness	Types
[Demangeon et al., 2015]	.	✓	.	✓	.	.	✓	.	.	.	✓
[Neykova et al., 2013]	.	✓	.	✓	.	.	✓	.	.	.	✓
[Neykova et al., 2014]	.	✓	.	✓	.	.	✓	.	.	.	✓
[Jia et al., 2016]	.	✓	.	✓	.	.	✓	.	.	✓	✓
[Di Giusto and Pérez, 2015]	.	✓	.	✓	✓	.	✓	.	.	✓	✓
[Di Giusto and Pérez, 2016]	.	✓	.	✓	✓	.	✓	.	.	✓	✓
[Castellani et al., 2014]	.	✓	.	✓	✓	.	✓	.	.	✓	✓
[Castellani et al., 2016]	.	✓	.	✓	✓	.	✓	.	.	✓	✓
[Mezzina and Pérez, 2016]	.	✓	.	✓	.	.	✓	.	.	✓	✓
[Mezzina and Pérez, 2017]	.	✓	.	✓	.	.	✓	.	.	✓	✓

Table 2. State-of-the-art on distributed monitoring. Each paper is classified according to the characteristics considered (part 2).

evaluation reported in [Bauer and Falcone, 2012, 2016] suggests that the overhead introduced is lower in the distributed solution. There is also an economic advantage in the decentralized solution, because in a distributed solution there is no need to add a central processor. Practical applications of this approach involve monitoring the behavior of embedded systems that are distributed by nature, like cars and airplanes where the different distributed components are known upfront. These include typical field-busses like EtherCAT, ProfiBus and ProfiNet (also known as “Industrial Ethernet” [Felser, 2005]). In these systems, processes communicate over a synchronous bus, so the global clock assumption is justified.

The works [Colombo and Falcone, 2014, 2016] start from similar assumptions and goals: there is a global clock and one local monitor per executing component. The number of executing components is also known upfront. The work in [Colombo and Falcone, 2016] removes the assumption of instantaneous communication from [Colombo and Falcone, 2014] and enables a solution with reliable messages with any delay. Still, a global clock is assumed because the specification logic is LTL and individual predicates sensed are totally ordered. The solution proposed is a choreographed decentralized monitoring algorithm, where each local monitor senses a collection of local predicates. The local monitors use the rewriting approach (also known as formula progression) by which the state of the monitor is the LTL formula that results by expanding the LTL formula to the residual formula in the next state, simplified with the acquired

knowledge. A key element in the solution is that a network of monitors is statically built by assigning each subformula of the original formula to a node in the distributed system. The hierarchical description inherent by the sub-formula relation in turn dictates the communication pattern between the local monitors. Consider a formula ψ and let φ be a sub-formula of ψ . The monitor M_1 for φ informs the monitor M_2 for ψ about the verdict of φ which, in turn, is used by M_2 to compute the verdict of ψ . If a synchronous clock is assumed, the root formula verdict is guaranteed to be reached within at most k steps of delay, where k is the height of the original formula.

Bartocci [Bartocci, 2013] extends the work [Bauer and Falcone, 2012] to real-time embedded systems by considering the maximum duration of the computation and communication. The main result is the ability to calculate a sampling ratio above which the decentralized monitoring process is guaranteed to generate the correct outcome.

Distributed Monitoring The work in [Sen et al., 2004] proposes a method to detecting violations of safety properties in an asynchronous distributed system, where no global clock is assumed. The method proposed generates, given a specification, local monitors for all distributed nodes. These local monitors communicate only by piggybacking additional information in the messages sent by existing processes in the system, so the shape of the history of messages exchanged is not modified by the actions taken by the monitors. The logic used in [Sen et al., 2004] extends past time LTL with features for distributed systems:

$$\begin{aligned} \varphi_i &:= true \mid P(\psi_i, \dots \psi_i) \mid \neg\varphi_i \mid \varphi_i \wedge \varphi_i \mid \exists\varphi_i \mid \diamond\varphi_i \mid \ominus\varphi_i \mid \varphi_i \mathcal{S} \varphi_i \mid @_j\varphi_j \\ f_i &:= c \mid v_i \mid @_j f_j \end{aligned}$$

where formula $@_j\varphi_j$ refers to the most recent value of formula φ_j according to process j .

The algorithm uses vector clocks [Lamport, 1978; Mattern, 1989; Fidge, 1989, 1991] to transmit the most recent value of sub-formulas needed to compute the outcome of their containing formulas. Then, at deployment time, the monitor specification is decomposed into local monitors that collect information locally and compute the current value of formulas, based on this local information and on the information received in messages about the causal past of remote processes. This approach allows to generate monitors without inspecting the internal behavior of each process.

Francalanza et al. [Francalanza et al., 2011, 2013] present a formal model for distributed monitoring. System computations are described as π -calculus processes [Milner et al., 1992] hosted at different locations and interacting with one another via message passing. When systems compute, they generate residual trace events that are only locally ordered (with respect to the other events generated at the same location) but globally unordered (with respect to events generated at other locations), thereby modeling the absence of global clocks. Distributed monitors, also residing at different locations, are then tasked with analyzing local traces and interacting with one another in order to perform a global

analysis of system computation. The model is equipped with a bisimulation-based equivalence relation that is used to reason about different distributed monitoring strategies such as those discussed in Section 4. The model is also used to define and evaluate a new migrating monitor strategy that better handles the dynamic nature of open distributed systems.

Fault tolerance Not many works attack the problem of monitoring distributed systems considering that components can fail. Notable exceptions are [Basin et al., 2013, 2015] and [Fraigniaud et al., 2014; Bonakdarpour et al., 2016a].

In [Basin et al., 2013] the authors present a policy language, a variant of FOLTL with three-valued semantics, and an algorithm that allows to reason about incomplete knowledge and handle disagreements. The main practical motivation is to handle errors in the observed trace, for example due to corruption or loss of part of logs files in complex IT systems, crashes in running systems, or network failures. Another motivation is to reconcile different views or verdicts obtained from monitors that observe different parts of the logs. The key idea is to equip the execution with features to enable monitors to distinguish between an event not being observed and the event not existing at all. The authors claim that any approach that solves this problem must satisfy that, once a definite verdict is given, providing more compatible information cannot retract the definite verdict. They manage to provide a complete algorithm for a fragment of the policy language. A similar work on compliance checking is [Garg et al., 2011]. Even though these works do not explicitly handle distributed systems, they handle runtime verification under incomplete information and incorrect information, which can be used to encode problems for distributed systems.

Influenced by [Basin et al., 2013], the work in [Basin et al., 2015] considers the problem of monitoring distributed systems in the presence of network failures. The authors also consider the case that the monitor itself is distributed for the purposes of efficiency, performing the monitoring computation closer to the observation point and reducing the communication overhead. The paper [Basin et al., 2015] deals with MTL, a logic that allows to express real-time properties. The algorithm is designed based on the *timed asynchronous* model for distributed systems [Cristian and Fetzer, 1999], which assumes the availability of highly synchronized local clocks but permits crash failures in the processes and in the network. Another assumption is that components are known at deployment time.

In [Basin et al., 2015] processes time-stamp their observation before communicating them to the local monitors. The time-stamp allows components to compute precise delays between events, and to totally order the events. It is interesting to point out that even without failures, reliable asynchronous networks allow messages to arrive in different orders. Forcing messages to arrive in order requires buffering messages to ensure proper delivery order, which in turn prevents the early detection of some violations that would be possible with out-of-order delivery.

The algorithm in [Basin et al., 2015] uses a richer value to encode the absence of knowledge when evaluating part of the specification. When the missing

information is finally received, the monitor can precisely resolve the uncertainty. Sometimes, a monitor can reach a precise verdict only with the partial information received in a timely manner. Consequently, the algorithm can monitor MTL properties tolerating the out-of-order arrival of partial observations.

Concerning organization, the monitors in [Basin et al., 2015] are distributed in a directed acyclic graph (DAG) where each monitor handles a subformula of the given formula, and children nodes handle subformulas of the formula handled by their parent node. The root of the DAG handles the original formula. During execution, messages are sent from children to parent monitors to inform about the verdicts reached in the subformula handled at the given point in time. When a process performs an atomic observation it also equips the time-stamp with an additional sequence number, which is locally unique. This sequence number allows monitors to infer the existence or absence of unknown intermediate samples between two observations. Intermediate nodes can also send heart-beat messages, which serve the purpose of informing about the absence of verdicts and the health of the intermediate node. Heart-beats also allow to infer the existence or absence of intermediate meaningful observations or verdicts, and in turn compute timeouts.

The problem of distributed monitoring for asynchronous distributed systems with node crashes is considered in [Fraigniaud et al., 2014; Bonakdarpour et al., 2016a]. Monitors can either work correctly or fail, but after a fail, monitors do not perform any action for the remainder of the execution. The solution is based on the asynchronous “wait-free” communicating infrastructure. It is well known from the research area of distributed algorithms that the wait-free model of computation [Attiya and Welch, 2004] can simulate many other models of crash-fail asynchronous distributed systems. The main result in [Bonakdarpour et al., 2016a; Fraigniaud et al., 2014] is an algorithm and a lower-bound on the number of different verdicts that monitors need to communicate with each other to correctly detect the violation of an LTL property. The lower bound on the number of verdicts reveals that monitors need to communicate complex information in order to compute a global outcome. The final verdict reached by the cooperating monitors, in turn, will be that of LTL_3 . The following three options are possible: (1) the property is satisfied in all continuations; (2) the property is violated in all continuations; (3) the outcome is unknown. These papers do not assume that the observations of the distributed monitors are disjoint. Even though monitors may only be observing part of the global input alphabet, several monitors may overlap in their partial observation.

Monitor Decomposition The work [Falcone et al., 2015b] targets the problem of monitoring component based-systems, that is, systems that are described by the composition of components. More precisely, in [Falcone et al., 2015b] systems are described using the BIP paradigm. Even though this paper does not attack explicitly the problem of monitoring a distributed system, it is nowadays well understood that component-based descriptions can be compiled into distributed implementations (see [Bonakdarpour et al., 2010a,b]) and not only into non-

distributed platforms. Consequently, the monitors generated at the component level following [Falcone et al., 2015b] are attached to the system generating a modified BIP description that can subsequently be compiled into a distributed system.

Monitor decentralization can also be inferred from the specification formula from which a monitor is synthesized. This line of research is explored extensively in [Francalanza and Seychell, 2013, 2015; Attard and Francalanza, 2016] for both safety and co-safety properties of logics involving conjunctions, disjunctions and recursion. Conjunctions and disjunctions are synthesized into concurrent monitors that analyze sub-parts of the system, whereas recursion leads to the dynamic generation of concurrent monitors, generated lazily only when required to minimize monitoring overheads. In every case, the concurrent monitors generated lead to self-contained localized monitoring that can be readily distributed. The automated synthesis function is proved correct in each of these cases (see [Francalanza et al., 2017] for the correctness proof of [Attard and Francalanza, 2016]). The work in [Cassar et al., 2015] considers a refined implementation where the concurrent sub-monitors cooperate among themselves and reorganize their interconnection so as to optimize the resources used for monitoring, thus reducing monitoring overheads.

Predicate Detection and Distributed Monitoring Predicate detection (see [Chase and Garg, 1998]) consists on checking whether a certain predicate occurred during the distributed execution, or more formally, whether the predicate holds in some *consistent cut* of the execution. In this context, predicates are state formulas (and consequently safety properties) even though some work has extended predicate detection to richer temporal formulas.

All algorithms for predicate detection assume that the collection of executing processes is known a-priori, that processes do not fail and that all messages eventually arrive. Predicate detection can be performed offline, when all events are available before the detection algorithm starts running, or online, when one event at a time is processed. There are three main techniques for predicate detection. The first technique uses the global snapshots proposed by Chandy and Lamport in [Chandy and Lamport, 1985], which can only detect *stable* predicates, which are predicates that remain true after becoming true (like termination, but unlike mutual exclusion). The second technique consists in an explicit construction of the lattice of global states proposed in [Cooper and Marzullo, 1991]. This technique can detect unstable predicates but it is exponential in the number of local states and processes. Finally, the third technique exploits the specific structure of the predicate to provide efficient solutions. Examples include conjunctions of local predicates [Garg and Waldecker, 1994] and relational predicates of the form $\sum_i x_i < C$, where x_i are local variables.

Even if one had access to all the local histories of the execution of all processes, detecting a predicate is hard because—for general predicates—one needs to enumerate and search all possible interleavings of the local executions. In [Chase and Garg, 1998], Chase and Garg show that detection of 2-CNF predi-

cates is an NP-hard problem, even when assuming a central monitor. A solution to this explosion problem is a technique called *slicing* [Garg and Mittal, 2001]. Slices are abstractions of the computation that guarantee that the predicate is detected in a slice if and only if the predicate holds in some consistent cut of the original computation. Computing a slice for a general predicate is still an NP-hard problem [Mittal and Garg, 2005], but when efficient slices exist, these are much smaller than actual explicit histories. Consequently, a line of work has focused on identifying classes of predicates for which efficient slicing procedures exist. These slices are based on fragments of the logic used to express the global state predicates. These fragments include regular, co-regular, linear, relational and stable predicates [Mittal and Garg, 2005; Mittal et al., 2007; Sen and Garg, 2007; Ogale and Garg, 2007]. Some of these solutions construct the slices offline, assuming that the whole histories are available to the slicing algorithms, while others work online, building the slice incrementally. Similarly, most of the solutions are still centralized [Cooper and Marzullo, 1991; Mittal et al., 2007; Sen and Garg, 2007; Mittal et al., 2007] in the sense that all histories are sent to a central monitor that computes the slice and detects the predicate.

The first distributed solution to slice-based predicate detection is [Chauhan et al., 2013]. The solution is online and distributed, in the sense that the slicing is computed by the distributed monitors. The guarantee is that if the predicate exists in a consistent cut of the computation, then it is detected by some monitor. The algorithm exploits both the structure of the property (in [Chauhan et al., 2013] regular properties) and epistemic information about what the knowledge that the different monitors acquire.

Also, even though most approaches are restricted to state predicates (or more precisely, fragments of the propositional logic for state predicates), some approaches tract richer temporal properties. For example, the works [Sen and Garg, 2003b; Ogale and Garg, 2007] present methods for sliced based predicate detection for a fragment of temporal logic that includes invariants (AG) and possible reachable (EF) operators, which extends the applicability beyond safety properties into a subclass of CTL formulas called Regular CTL [Sen and Garg, 2003a]. The restrictive use of negation in [Sen and Garg, 2003b] is relaxed in [Ogale and Garg, 2007]. Even though the work in [Sen and Garg, 2003b; Ogale and Garg, 2007] is applicable to a richer fragment of temporal logic, these algorithms work with a central monitor.

More recently, Mostafa and Bonakdarpour [Mostafa and Bonakdarpour, 2015] provide a solution for monitorable LTL_3 temporal properties, but in this case extending [Chauhan et al., 2013] so the solution obtained is distributed. This solution inserts additional messages in the network and is not restricted to only piggybacking information in existing messages.

Intrusiveness It is often desirable that the monitoring process perturbs the execution of the system under analysis in the least possible manner. Typically, either the system is instrumented by embedding monitors in the code itself, or monitors and processes share resources because they execute in the same plat-

form. These changes affect the behavior of the system, sometimes in a significant manner.

In [Berkovich et al., 2015] the authors propose to use additional hardware, and in particular a GPU parallel execution platform, to minimize the impact of online monitoring. The authors show how to generate parallel monitors from temporal logic specifications and evaluate empirically that the obtained parallel monitors together with the additional GPU hardware alleviate the effect of monitoring on the execution of the original system.

Other times, it is desirable that the monitoring process perturbs the execution of the system. One example is runtime enforcement, where the objective of the “monitoring” is to guarantee that the system stays within a safe region of states. Consequently, the enforcement system uses the information provided by the monitor to prevent an error before it occurs (see the chapter in this monograph about runtime enforcement). Another example is testing of multithreaded programs, which is in general a very hard task, due to the non-deterministic nature of the execution of concurrent programs, and the difficulty to reproduce erroneous behaviors. In this context it is desirable to guide the system towards executions that are more likely to produce an error. The work in [Luo and Roşu, 2013] consists of an *enforcement* mechanism that uses user-specified monitors to generate local monitors. Such local monitors block individual threads that violate the specified properties. This enforcement pursues two objectives: (1) to guarantee the enforcement of properties in a multi-threaded program in a systematic way, which is typically implemented using ad-hoc synchronization manually; and (2) to force schedules that test properties during the testing of multithreaded programs. The monitor generation described in [Luo and Roşu, 2013] includes the decomposition of the property into local decentralized monitors for each of the threads.

The body of work [Roşu and Havelund, 2005; Cassar and Francalanza, 2014; Zhang et al., 2016] explores the idea of decoupling the execution of monitors from the systems under scrutiny. This approach uses a mixture of synchronous and asynchronous monitoring, in order to obtain a feasible instrumentation setup that distribute monitors and systems at different locations, such as in the case of [Colombo et al., 2012] and other orchestrated monitoring setups. Asynchronous monitoring, used in various monitoring tools such as [Colombo et al., 2011; Francalanza and Seychell, 2015; Zhang et al., 2016; Attard and Francalanza, 2016], minimizes monitor intrusiveness because it requires less instrumentation effort. Moreover, in [Cassar and Francalanza, 2014; Zhang et al., 2016] the authors show that this method of instrumentation can substantially reduce monitoring overhead. By using hybrid solutions, they also show how one need not compromise on the timeliness of detections.

In [Cassar and Francalanza, 2016, 2015], Cassar and Francalanza extend the concept of non-intrusiveness to runtime adaptation via hybrid asynchronous monitoring. The goal is to design monitors that intervene with the execution of the system under scrutiny, and apply these interventions (i.e. system adaptations) with minimal overheads. In particular, the work [Cassar and Francalanza,

2016] implements a framework where the monitors for system components can act at varying degrees of synchrony with respect to the observed components. Some parts of the system can be executed in a decoupled fashion with their monitors when no adaptations on that sub-system are required. Later, these sub-systems can be incrementally synchronized with the respective monitor when an adaptation is about to be applied. The entire framework is implemented atop a completely asynchronous actor computational model, which eases the distribution over remote locations.

Behavioral Type Systems for Distributed Monitoring Many large-scale systems consist of heterogeneous, distributed software artifacts (processes) that interact following some precise protocols. In these communication-centric settings, processes communicate asynchronously, without a global clock, and are prone to local failures. These characteristics make distributed monitoring a suitable approach to enforce system correctness by complementing the static verification techniques that are typically applied individually to each process. As we detail next, monitoring for communication-centric systems is an instance of the choreographed monitoring organization described in Section 4.

A productive research strand to the analysis of communication-centric software systems uses *process calculi* (such as the π -calculus) as minimal specification languages. These formal calculi provide an unambiguous setting in which the communication correctness of these systems can be compositionally established. In particular, coupling process calculi with so-called *behavioral type systems* allows to (statically) enforce safety and liveness properties associated to protocol conformance. Rather than classifying data values, behavioral types define abstractions of the protocols that a communication entity (say, a socket or a channel) should respect throughout its execution (see [Hüttel et al., 2016] for a survey).

Several works have explored the interplay of behavioral types and mechanisms for distributed monitoring. In particular, monitoring frameworks based on *session types*, a particular class of behavioral types, have been put forward. Session types organize a series of communication actions corresponding to the same reciprocal protocol into a structure called *session* [Honda et al., 1998]. While typed process frameworks for binary session types can analyze two-party protocols, more general type theories for multiparty session types cover the case of protocols with three or more participants [Honda et al., 2016]. Both binary and multiparty session types start to make their way into mainstream programming languages and frameworks [Ancona et al., 2016]. In the multiparty case, a global type entirely describes the intended communication scenario. By projecting this global type onto each protocol participant, one may obtain its corresponding local type, which abstracts a participant’s contribution to the protocol. This collection of local types thus offers a key reference for obtaining correct implementations for all participants.

Communication-centric systems often comprise components made available as grey- and black-boxes, with limited communication interfaces. As such, static

verification techniques are unsuitable for their validation. Motivated by this observation, several works develop abstract frameworks based on process calculi in which monitors are terms of the specification language. The formal semantics of these calculi uses these monitor terms to enable process behavior according to the intended protocol. Rather than a logical specification (say, an LTL formula), each monitor uses a behavioral type (e.g., a local protocol) to guide a participant’s behavior. These works define a special case of choreographed monitoring: the coupling of processes and monitors at the same level of abstraction makes the notion of local trace implicit. Monitors do not communicate to each other, nor perform autonomous actions. The global type through its projections is used to synthesize a monitor for each participant. This way, even untyped processes can be used to implement a protocol participant as long as they offer the right communication actions at the right time, in accordance with the governing local protocols.

Based on this general setup, Bocchi et al. [Bocchi et al., 2013, 2017] develop a monitored π -calculus with dynamic usage of multiparty session types, offering local and global safety assurance of distributed components. In their model, a network is a collection of processes (one per participant) that communicate via asynchronous message passing. Each participant is equipped with a trusted monitor that guards the run-time behavior of both the principal and its environment—this is realized by the evaluation of incoming and outgoing messages. Monitors regulate the creation of sessions and movement of messages within sessions. This dynamic checking can be switched off when processes have been statically verified. A series of queues shared between principals is assumed to support message passing, together with a global transport that abstracts distributed communication.

Building upon [Bocchi et al., 2013, 2017], the works [Hu et al., 2013; Demangeon et al., 2015] propose a dynamic verification framework for multiparty session types that admit interruptions. This a practical framework, which relies on the Scribble protocol language (an implementation of multiparty session types [Yoshida et al., 2013]) to specify global protocols, and on a Python API for conversation programming. In this framework, the monitor that tracks the progress of each participant within a session is represented using a finite state machine (FSM), generated from the local type. By independently monitoring each session endpoint at runtime, this framework ensures global communication safety even in the presence of asynchronous interruptions.

Other works on a practical strand are [Neykova et al., 2013, 2014]. In [Neykova et al., 2013], the authors propose a toolchain for designing deadlock-free multiparty global protocols. Using automatically generated monitors for each session endpoint, this toolchain can detect illegal communication actions and mistaken message types that go against protocol conformance. The work in [Neykova et al., 2014] extends preceding works with timed information: Scribble specifications are extended with clocks, resets, and clock predicates that constrain the occurrence of protocol interactions.

Recent work introduces a framework for monitoring interacting processes that follow binary session protocols, building upon a logically motivated theory of session types [Jia et al., 2016]. As in several of the works mentioned above, in this framework monitors are placed next to communication endpoints. A distinguishing aspect is blame assignment: in case processes deviate from the prescribed session protocols, monitors may halt the execution, raise an alarm, and assign blame. The authors prove that their dynamic monitoring is not intrusive in the sense that it does not change the behavior of well-typed processes. Also, they show that in case of alarm one of an indicated set of possible culprits must have been compromised.

Finally, we mention some works in which the concept of monitor as a process term, in the sense just described, has been exploited. Even though the main purpose of these works is not run-time verification, they can be seen as applications of choreographed monitoring. The works in [Di Giusto and Pérez, 2015; Di Giusto and Pérez, 2016] use this kind of monitors to support the run-time adaptation of session-typed processes in both binary and multiparty settings. There is exactly one monitor per session. By combining monitor information and event-based constructs, one may specify the reaction to unanticipated circumstances (for example, local failures) by means of adaptation steps. An associated type system ensures communication safety and consistency properties: while safety guarantees absence of run-time communication errors, consistency ensures that adaptation steps do not disrupt already established session protocols. In a similar line, the monitors defined in [Castellani et al., 2014, 2016] play a dual role: they enforce run-time adaptation policies, and ensure secure information flow in multiparty exchanges. Recent work by Mezzina and Pérez uses monitors as the memories required to support models of concurrency in which actions are reversible and causally consistent [Mezzina and Pérez, 2016, 2017].

6 Challenges and Conclusion

6.1 Challenges

We list here some challenges for future research in distributed runtime verification.

Fault tolerance One of the key characteristics of distributed systems is that, in practice, different parts of the system can fail independently. However, most approaches consider that the system does not fail. Some future problems include the following.

The theoretical approach in [Fraigniaud et al., 2014; Bonakdarpour et al., 2016a] (discussed in Section 5) has two major obstacles to become practical:

- First, after the distributed verdicts are emitted, there is a phase in which a global function is applied to the collection of verdicts emitted. This function must be implemented somehow by a central computational infrastructure

which must receive all verdicts and produce an outcome. However, a general implementation of this function requires a non-failing central monitor. But the existence of such a central unit would greatly simplify the initial monitoring problem, and in fact, the basic starting point of [Bonakdarpour et al., 2016a; Fraigniaud et al., 2014] is to design distributed fault tolerant solutions.

- Second, the work [Bonakdarpour et al., 2016a] only presents an algorithm for the processing of a one letter observation, under the assumption that the processes are perfectly synchronized at the beginning of such an observation. To process a subsequent observation, the monitors that survive the first round must somehow re-synchronize, but again, a synchronization procedure would provide a much simpler solution to the monitoring problem at hand. In summary, A general fault-tolerant solution for sequences of observations is still an open problem.

Also, there are very few results in runtime verification that can handle network failures (most notably, the work in [Basin et al., 2015]). It would be very interesting to extend these approaches to other logics and distributed system assumptions.

Global atomic observations Specification formalisms for non-distributed systems assume that atomic predicates are testable, which is not a restriction. In distributed systems, in general, predicates are global in the sense that they can involve different parts of the system. Then, not all global predicates are Boolean combinations of local predicates. For example, one restriction in [Bauer and Falcone, 2012] is that the individual global observations are Boolean combinations of local observations performed in each of the processes, whose observations do not overlap. More formally, each process j can emit a collection of local propositions AP_j (such that $AP_j \cap AP_i = \emptyset$ whenever $i \neq j$). The alphabet of atomic observations is then $\Sigma_j = 2^{AP_j}$. Note how the global alphabet $\Sigma = 2^{\cup_i AP_i}$ is strictly larger, in general, than $\cup_i \Sigma_i$ because it can contain relational symbols like $p_i \vee p_j$ where p_i and p_j are local observations at different processes. We use *relational observations* to refer to atomic propositions whose truth value depends on the observations made at more than one process. For example, consider the numeric variables x_i and y_j where the sub-index indicates the process at which the variable is observed (P_i and P_j resp). The atomic predicate $x_i < y_j$ cannot be evaluated at P_i or P_j alone, and it cannot be decomposed into a Boolean combination of local predicates either.

As discussed earlier, even though research in predicate detection has considered classes of predicates richer than individual observations (regular, linear, etc) and has characterized that detecting a predicate in 2-CNF is already NP-hard, it would be interesting to extend other techniques for decentralized and distributed monitoring beyond combinations of local predicates.

Monitor Orchestrations In [Colombo and Falcone, 2016], a choreographed decentralized monitoring solution is obtained from a network of local monitors,

which is statically computed by mapping every subformula to a distributed system node. There are many possible ways to create such a network, even if one restricts the map (as in [Colombo and Falcone, 2016]) to one of the nodes with the highest number of propositions locally involved in the subformula, because there can be more than one such node. Even though all choices could lead to a correct monitoring solution, for a given trace of execution, the choice of network has an impact in the communication overhead. For every input trace, one could calculate a-posteriori the best network in the sense of the network that would have produced the lowest overhead. However, even for the fixed parameter assumed in [Colombo and Falcone, 2016] (e.g., static number of locations, fixed specifications, no dynamic remote spawning of new computation, the assumption of a global clock) it is not clear how to pre-compute an optimal network, even how to approximate it. Nevertheless, there are alternatives worth investigating. One plausible solution is to exercise the system in a test-bed to obtain input traces and compute the optimal network for the observed set of traces, with the assumption that the traces after deployment will involve similar communication flows. However, this kind of approach is not considered in [Colombo and Falcone, 2016].

Adequate solutions to this problem are probably even harder to come up with when proper distributed system constraints are considered, such as computation asynchrony, distributed clocks, and the possibility of partial failure. In practical settings, cases may even arise whereby one has to content with conflicting criteria. For instance, certain locations may not allow monitor processing and analysis to be carried out locally, forcing events to be communicated remotely to the analyzing monitor. This, in turn, may conflict with confidentiality and security concerns.

Monitorability and Correctness In general, the use of runtime analysis impinges on the extent to which a correctness property can be verified.

This aspect is often referred to as *monitorability*. One of the first works that introduces monitorability is [D’Angelo et al., 2005] for stream runtime verification on finite traces as a class of specification for which efficient monitors can be generated. Later, Pnueli and Zaks [Pnueli and Zaks, 2006] formalized monitorability for LTL as the possibility of a finite trace to be extended to a finite witness of a specification satisfaction or violation. A similar notion was presented in [Bauer et al., 2011] and proved equivalent in [Falcone et al., 2012]. This notion was generalized to ω -regular languages in [Bauer, 2010], and later extended and corrected in [Diekert et al., 2015]. The tight complexity of this notion of monitorability was finally captured in [Baader and Lippmann, 2014]. An alternative definition of monitorability is given in [Francalanza et al., 2017] where the fragment of formulas of a given branching time logic that can be monitored at runtime is captured.

Decentralized and distributed monitoring introduces further restrictions and raises additional issues that may affect the monitorability of certain correctness

properties. It is worth investigating these limits and establishing firm results that can guide the engineering of distributed monitoring.

Concurrent and distributed systems are notoriously hard to get right and these complications extend also to distributed runtime verification: errors arise only for particular sequences of events that are hard to simulate using pre-deployment techniques such as testing, and are also hard to trace and reproduce for analysis once they occur. It is thus imperative to continue to extend existing work on developing methods for ascertaining the correctness of the decentralized and distributed monitoring setups constructed along the lines of [Falcone et al., 2015a; Bocchi et al., 2017; Francalanza, 2016].

6.2 Conclusion

In this chapter we have surveyed the literature on runtime verification for distributed systems. After showing some practical motivations that have justified the study of monitoring techniques for distributed and decentralized systems, we identified a series of features that characterize and that allow to classify the different problems and approaches. These criteria include whether the solution involves or exploits the description of the system under analysis, whether there is a single central monitor or the monitoring task is distributed, whether there is an assumption on a global clock, and whether the system tolerates failures or perturbs the execution. Finally, we showed a comprehensive list of results proposed in the literature and listed some challenges for future work.

Acknowledgments We are grateful to the anonymous reviewers for their useful remarks and suggestions, which led to significant improvements.

[Financial acknowledgements to be inserted here in the final version]

Bibliography

- S. Alagar and S. Venkatesan. Techniques to tackle state explosion in global predicate detection. *IEEE Transactions on Software Engineering (TSE)*, 27(8):704–714, 2001.
- Davide Ancona, Viviana Bono, Mario Bravetti, Joana Campos, Giuseppe Castagna, Pierre-Malo Deniérou, Simon J. Gay, Nils Gesbert, Elena Giachino, Raymond Hu, Einar Broch Johnsen, Francisco Martins, Viviana Mascardi, Fabrizio Montesi, Rumyana Neykova, Nicholas Ng, Luca Padovani, Vasco T. Vasconcelos, and Nobuko Yoshida. Behavioral types in programming languages. *Foundations and Trends in Programming Languages*, 3(2-3):95–230, 2016.
- Duncan Paul Attard and Adrian Francalanza. A Monitoring Tool for a Branching-Time Logic. In *Proc. of RV'16*, volume 10012 of *LNCS*, pages 473–481. Springer, 2016.
- Hagit Attiya and Jennifer L. Welch. *Distributed computing: fundamentals, simulations and advanced topics*. Wiley, 2004.
- Franz Baader and Marcel Lippmann. Runtime verification using the temporal description logic ALC-LTL revisited. *Journal of Applied Logic*, 12(4):584–613, 2014.
- Ezio Bartocci. Sampling-based decentralized monitoring for networked embedded systems. In Luca Bortolussi, Manuela L. Bujorianu, and Giordano Pola, editors, *Proc. of the 3rd Int'l Workshop on Hybrid Autonomous Systems (HAS'13)*, volume 124 of *EPTCS*, pages 85–99, 2013.
- David Basin, Felix Klaedtke, Srdjan Marinovic, and Eugen Zălinescu. Monitoring compliance policies over incomplete and disagreeing logs. In *Proc. of RV'13*, volume 8174 of *LNCS*, pages 151–167. Springer, 2013. doi: 10.1007/978-3-642-35632-2_17.
- David A. Basin, Felix Klaedtke, and Eugen Zălinescu. Failure-aware runtime verification of distributed systems. In *Proc. of FSTTCS'15*, pages 590–603, 2015.
- Andreas Bauer, Martin Leucker, and Christian Schallhart. Runtime verification for LTL and TLTL. *ACM Trans. Softw. Eng. Methodol.*, 20:14:1–14:64, 2011.
- Andreas Klaus Bauer. Monitorability of ω -regular languages. arXiv:1006.3638v1, 2010.
- Andreas Klaus Bauer and Yliès Falcone. Decentralised LTL monitoring. In *Proc. of the 18th Int'l Symp. on Formal Methods (FM'12)*, volume 7436 of *LNCS*, pages 85–100. Springer, 2012.
- Andreas Klaus Bauer and Yliès Falcone. Decentralised ltl monitoring. *Formal Methods in System Design*, 48(1–2):49–93, 2016.
- Shay Berkovich, Borzoo Bonakdarpour, and Sebastian Fischmeister. Runtime verification with minimal intrusion through parallelism. *Formal Methods in System Design*, 46(3):317–348, 2015. doi: 10.1007/s10703-015-0226-3.

- Laura Bocchi, Tzu-Chun Chen, Romain Demangeon, Kohei Honda, and Nobuko Yoshida. Monitoring networks through multiparty session types. In *Proc. of FMOODS/FORTE'13*, pages 50–65, 2013.
- Laura Bocchi, Tzu-Chun Chen, Romain Demangeon, Kohei Honda, and Nobuko Yoshida. Monitoring networks through multiparty session types. *Theor. Comput. Sci.*, 669:33–58, 2017.
- Borzoo Bonakdarpour, Marius Bozga, Mohamad Jaber, Jean Quilbeuf, and Joseph Sifakis. Automated conflict-free distributed implementation of component-based models. In *Proc. of the IEEE 5th Int'l Symp. on Industrial Embedded Systems (SIES'10)*, pages 108–117. IEEE, 2010a.
- Borzoo Bonakdarpour, Marius Bozga, Mohamad Jaber, Jean Quilbeuf, and Joseph Sifakis. From high-level component-based models to distributed implementations. In *Proc. of EMSOFT'10*, pages 209–218. ACM, 2010b.
- Borzoo Bonakdarpour, Pierre Fraigniaud, Sergio Rajsbaum, David Rosenblueth, and Corentin Travers. Decentralised asynchronous crash-resilient runtime verification. In *Proc. of the 27th Int'l Conf. on Concurrency Theory (CONCUR'16)*, volume 59 of *LIPICs*, pages 16:1–16:15. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2016a.
- Borzoo Bonakdarpour, Pierre Fraigniaud, Sergio Rajsbaum, and Corentin Travers. Challenges in fault-tolerant distributed runtime verification. In Tiziana Margaria and Bernhard Steffen, editors, *Proc. of ISoLA 2016 Part II*, volume 9953 of *Lecture Notes in Computer Science*, pages 363–370, 2016b.
- Marco Carbone, editor. *Proc. of the Third Workshop on Behavioural Types, (BEAT'14)*, volume 162 of *EPTCS*, 2014.
- Ian Cassar and Adrian Francalanza. On Synchronous and Asynchronous Monitor Instrumentation for Actor Systems. In *Proc. of FOCLASA'14*, volume 175, pages 54–68, 2014.
- Ian Cassar and Adrian Francalanza. Runtime Adaptation for Actor Systems. In *RV*, volume 9333 of *LNCS*, pages 38–54. Springer, 2015.
- Ian Cassar and Adrian Francalanza. On implementing a monitor-oriented programming framework for actor systems. In *Proc. of iFM'16*, 2016.
- Ian Cassar, Adrian Francalanza, and Simon Said. Improving runtime overheads for detector. In Bara Buhnova, Lucia Happe, and Jan Kofron, editors, *Proc. of the 12th Int'l Workshop on Formal Engineering approaches to Software Components and Architectures (FESCA'15)*, volume 178 of *EPTCS*, pages 1–8, 2015.
- Ilaria Castellani, Mariangiola Dezani-Ciancaglini, and Jorge A. Pérez. Self-adaptation and secure information flow in multiparty structured communications: A unified perspective. In Carbone [2014], pages 9–18.
- Ilaria Castellani, Mariangiola Dezani-Ciancaglini, and Jorge A. Pérez. Self-adaptation and secure information flow in multiparty communications. *Formal Asp. Comput.*, 28(4):669–696, 2016.
- K. Mani Chandy and Leslie Lamport. Distributed snapshots: Determining global states of distributed systems. *ACM Transactions on Computer Systems*, 3(1): 63–75, 1985.

- Craig M. Chase and Vijay K. Garg. Detection of global predicates: Techniques and their limitations. *Distributed Computing*, 11(4):191–201, 1998.
- Himanshu Chauhan, Vijay K. Garg, Aravind Natarajan, and Neeraj Mittal. A distributed abstraction algorithm for online predicate detection. In *IEEE 32nd Symp. on Reliable Distributed Systems, (SRDS'13)*, pages 101–110. IEEE Computer Society, 2013.
- Feng Chen and Grigore Roşu. Parametric trace slicing and monitoring. In *Proc. of the Int'l Conf. on Tools and Algorithms for the Construction and Analysis of Systems (TACAS'09)*, volume 5505 of *LNCS*. Springer, 2009.
- Christian Colombo and Yliès Falcone. Organising LTL monitors over distributed systems with a global clock. In *Proc. of the 5th Int'l Conference on Runtime Verification (RV'14)*, pages 140–155, 2014.
- Christian Colombo and Yliès Falcone. Organising LTL monitors over distributed systems with a global clock. *Formal Methods in System Design*, 49(1–2):109–158, 2016.
- Christian Colombo, Adrian Francalanza, and Rudolph Gatt. Elarva: A monitoring tool for Erlang. In *Proc. of RV'11*, volume 7186 of *LNCS*, pages 370–374. Springer, 2011.
- Christian Colombo, Adrian Francalanza, Ruth Mizzi, and Gordon J. Pace. polylarva: Runtime verification with configurable resource-aware monitoring boundaries. In *Proc. of SEFM'12*, volume 7504 of *LNCS*, pages 218–232. Springer, 2012.
- Christian Colombo, Gabriel Dimech, and Adrian Francalanza. Investigating instrumentation techniques for ESB runtime verification. In *Proc. of SEFM'15*, volume 9276 of *LNCS*, pages 99–107. Springer, 2015.
- Robert Cooper and Keith Marzullo. Consistent detection of global predicates. In *Proc. of the ACM/ONR Workshop on Parallel and Distributed Debugging*, pages 163–173, 1991.
- George Coulouris. *Distributed Systems: Concepts and Design*. Addison-Wesley, 2011.
- Flaviu Cristian and Christof Fetzer. The timed asynchronous distributed system model. *IEEE Transactions on Parallel and Distributed Systems*, 10(6):642–657, 1999.
- Ben D'Angelo, Sriram Sankaranarayanan, César Sánchez, Will Robinson, Bernd Finkbeiner, Henny B. Sipma, Sandeep Mehrotra, and Zohar Manna. LOLA: Runtime monitoring of synchronous systems. In *Proc. of the 12th Int'l Symp. of Temporal Representation and Reasoning (TIME'05)*, pages 166–174. IEEE CS Press, 2005.
- Romain Demangeon, Kohei Honda, Raymond Hu, Rumyana Neykova, and Nobuko Yoshida. Practical interruptible conversations: distributed dynamic verification with multiparty session types and python. *Formal Methods in System Design*, 46(3):197–225, 2015.
- Cinzia Di Giusto and Jorge A. Pérez. An event-based approach to runtime adaptation in communication-centric systems. In Thomas T. Hildebrandt, António Ravara, Jan Martijn van der Werf, and Matthias Weidlich, editors, *Proc. of the 11th Int'l Workshop on Web Services, Formal Methods, and Behavioral Types*

- (*WS-FM'14*), volume 9421 of *Lecture Notes in Computer Science*, pages 67–85. Springer, 2015.
- Cinzia Di Giusto and Jorge A. Pérez. Event-based run-time adaptation in communication-centric systems. *Formal Aspects of Computing*, pages 1–36, 2016. doi: 10.1007/s00165-016-0377-z.
- Volker Diekert, Anca Muscholl, and Igor Walukiewicz. A note on monitors and Bèchi automata. In *Proc. of In'11 Colloquium on Theoretical Aspects of Computing (ICTAC'15)*, volume 9399 of *LNCS*, pages 39–57. Springer, 2015.
- Yliès Falcone, Jean-Claude Fernandez, and Laurent Mounier. What can you verify and enforce at runtime? *STTT*, 14(3):349–382, 2012.
- Yliès Falcone, Mohamad Jaber, Thanh-Hung Nguyen, Marius Bozga, and Saddek Bensalem. Runtime verification of component-based systems in the bip framework with formally-proved sound and complete instrumentation. *Software & Systems Modeling*, 14(1):173–199, 2015a. doi: 10.1007/s10270-013-0323-y.
- Yliès Falcone, Mohamad Jaber, Thanh-Hung Nguyen, Marius Bozga, and Saddek Bensalem. Runtime verification of component-based systems in the BIP framework with formally-proved sound and complete instrumentation. *Software and System Modeling*, 14(1):173–199, 2015b.
- Max Felser. Real-time ethernetindustry prospective. *Proceedings of the IEEE*, 93(6):1118–1129, 2005.
- Colin Fidge. Timestamps in message-passing systems that preserve the partial ordering. In *Proc. of the 11th Australian Computer Science Conf.*, pages 55–66, 1989.
- Colin Fidge. Logical time in distributed computer systems. *Computer*, 1991.
- Pierre Fraigniaud, Sergio Rajsbaum, and Corentin Travers. On the number of opinions needed for fault-tolerant run-time monitoring in distributed systems. In *RV*, 2014. doi: 10.1007/978-3-319-11164-3_9.
- Adrian Francalanza. A Theory of Monitors. In *FoSSaCS*, volume 9634 of *LNCS*, pages 145–161. Springer, 2016.
- Adrian Francalanza and Matthew Hennessy. A Theory of System Behaviour in the presence of Node and Link failure. *Information and Computation*, 206(6): 711 – 759, 2008.
- Adrian Francalanza and Aldrin Seychell. Synthesising correct concurrent run-time monitors (Extended Abstract). In *Proc. of RV'13*, volume 8174 of *LNCS*, pages 112–129. Springer, 2013.
- Adrian Francalanza and Aldrin Seychell. Synthesising Correct concurrent Runtime Monitors. *FMSD*, 46(3):226–261, 2015. doi: 10.1007/s10703-014-0217-9.
- Adrian Francalanza, Andrew Gauci, and Gordon J. Pace. Distributed system contract monitoring. In Ernesto Pimentel and Valentin Valero, editors, *Proceedings Fifth Workshop on Formal Languages and Analysis of Contract-Oriented Software, FLACOS 2011, Málaga, Spain, 22nd and 23rd September 2011.*, volume 68 of *EPTCS*, pages 23–37, 2011.
- Adrian Francalanza, Andrew Gauci, and Gordon J. Pace. Distributed system contract monitoring. *Journal of Logic and Algebraic Programming*, 82(5–7): 186–215, 2013.

- Adrian Francalanza, Luca Aceto, and Anna Ingólfssdóttir. Monitorability for the Hennessy–Milner logic with recursion. *FMSD*, pages 1–30, 2017.
- Deepak Garg, Limin Jia, and Anupam Datta. Policy auditing over incomplete logs: Theory, implementation and applications. In *Proc. of CCS’11*, pages 151–162, 2011.
- Vijay K. Garg. *Elements of Distributed Computing*. Wiley-IEEE Press, 2002.
- Vijay K. Garg and Neeraj Mittal. On slicing a distributed computation. In *Proc. of the 21st IEEE Int’l Conf. on Distributed Computing Systems (ICDCS’01)*, pages 322–329. IEEE CS Press, 2001.
- Vijay K. Garg and Brian Waldecker. Detection of weak unstable predicates in distributed programs. *IEEE Transactions on Parallel and Distributed Systems*, 5(3):299–307, 1994.
- Kohei Honda, Vasco Thudichum Vasconcelos, and Makoto Kubo. Language primitives and type discipline for structured communication-based programming. In Chris Hankin, editor, *Proc. of the 7th European Symp. Programming Languages and Systems (ESOP’98)*, volume 1381 of *Lecture Notes in Computer Science*, pages 122–138. Springer, 1998.
- Kohei Honda, Nobuko Yoshida, and Marco Carbone. Multiparty asynchronous session types. *J. ACM*, 63(1):9:1–9:67, 2016.
- Raymond Hu, Rumyana Neykova, Nobuko Yoshida, Romain Demangeon, and Kohei Honda. Practical interruptible conversations - distributed dynamic verification with session types and python. In Legay and Bensalem [2013], pages 130–148.
- Hans Hüttel, Ivan Lanese, Vasco T. Vasconcelos, Luís Caires, Marco Carbone, Pierre-Malo Deniérou, Dimitris Mostrous, Luca Padovani, António Ravara, Emilio Tuosto, Hugo Torres Vieira, and Gianluigi Zavattaro. Foundations of session types and behavioural contracts. *ACM Comput. Surv.*, 49(1):3:1–3:36, 2016.
- Limin Jia, Hannah Gommerstadt, and Frank Pfenning. Monitors and blame assignment for higher-order session types. In Rastislav Bodík and Rupak Majumdar, editors, *Proceedings of the 43rd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2016, St. Petersburg, FL, USA, January 20 - 22, 2016*, pages 582–594. ACM, 2016.
- Jeffrey Joyce, Greg Lomow, Konrad Slind, and Brian W. Unger. Monitoring distributed systems. *ACM Trans. Comput. Syst.*, 5(2):121–150, 1987.
- Leslie Lamport. Time, clocks and the ordering of events in distributed systems. *Communications of the ACM*, 1978.
- Axel Legay and Saddek Bensalem, editors. *Proc. of the 4th Int’l Conf. on Runtime Verification (RV’13)*, volume 8174 of *LNCS*, 2013. Springer.
- Qingzhou Luo and Grigore Roşu. EnforceMOP: A Runtime Property Enforcement System for Multithreaded Programs. In *ISSTA*, New York, NY, USA, 2013. ACM.
- Friedemann Mattern. Virtual time and global states of distributed systems. In *Proc. of the Workshop on Parallel and Distributed Algorithms*, pages 215–226. Elsevier, 1989.

- Claudio Antares Mezzina and Jorge A. Pérez. Reversible sessions using monitors. In Dominic A. Orchard and Nobuko Yoshida, editors, *Proc. of the Ninth Workshop on Programming Language Approaches to Concurrency- and Communication-centric Software, (PLACES'16)*, volume 211 of *EPTCS*, pages 56–64, 2016.
- Claudio Antares Mezzina and Jorge A. Pérez. Reversibility in session-based concurrency: A fresh look. *Journal of Logical and Algebraic Methods in Programming*, 90:2–30, 2017.
- Robin Milner, Joachim Parrow, and David Walker. A calculus of mobile processes, I. *Inf. Comput.*, 100(1):1–40, 1992.
- Neeraj Mittal and Vijay K. Garg. Techniques and applications of computation slicing. *Distributed Computing*, 17(3):251–277, 2005.
- Neeraj Mittal, Alper Sen, and Vijay K. Garg. Solving computation slicing using predicate detection. *IEEE Transactions on Parallel and Distributed Systems (TPDS)*, 18(12):1700–1713, 2007.
- Menna Mostafa and Borzoo Bonakdarpour. Decentralized runtime verification of LTL specifications in distributed systems. In *Proc. of the 2015 IEEE Int'l Parallel and Distributed Processing Symp. (IPDPS'15)*, pages 494–503. IEEE Computer Society, 2015.
- Rumyana Neykova, Nobuko Yoshida, and Raymond Hu. SPY: local verification of global protocols. In Legay and Bensalem [2013], pages 358–363.
- Rumyana Neykova, Laura Bocchi, and Nobuko Yoshida. Timed runtime monitoring for multiparty conversations. In Carbone [2014], pages 19–26.
- Vinit A. Ogale and Vijay K. Garg. Detecting temporal logic predicates on distributed computations. In *Proc. of the Symposium on Distributed Computing (DISC'07)*, volume 4731 of *LNCS*, pages 420–434. Springer, 2007.
- Amir Pnueli and Aleksandr Zaks. PSL model checking and run-time verification via testers. In *Proc. of the 14th Int'l Symp. on Formal Methods (FM'06)*, volume 4085 of *LNCS*, pages 573–586. Springer, 2006.
- Grigore Roşu and Klaus Havelund. Rewriting-based techniques for runtime verification. *Automated Software Engineering*, 12(2):151–197, April 2005.
- Alper Sen and Vijay K. Garg. Detecting temporal logic predicates in distributed programs using computation slicing. In *Proc. of the 7th Int'l Conf. on Principles of Distributed Systems (OPODIS'03)*, volume 3144 of *LNCS*, pages 171–183. Springer, 2003a.
- Alper Sen and Vijay K. Garg. Partial order trace analyzer (POTA) for distributed programs. *ENTCS-Proc. Workshop on Runtime Verification (RV'03)*, 89(2), 2003b.
- Alper Sen and Vijay K. Garg. Formal verification of simulation traces using computation slicing. *IEEE Transactions on Computers*, 2007.
- Koushik Sen, Abhay Vardhan, Gul Agha, and Grigore Rosu. Efficient decentralized monitoring of safety in distributed systems. In *Proc. of ICSE'04*. IEEE CS Press, 2004.
- Koushik Sen, Abhay Vardhan, Gul Agha, and Grigore Rosu. Decentralized runtime analysis of multithreaded applications. In *IPPS*, 2006.

- Nobuko Yoshida, Raymond Hu, Rumyana Neykova, and Nicholas Ng. The scribble protocol language. In Martín Abadi and Alberto Lluch-Lafuente, editors, *Trustworthy Global Computing - 8th International Symposium, TGC 2013, Buenos Aires, Argentina, August 30-31, 2013, Revised Selected Papers*, volume 8358 of *Lecture Notes in Computer Science*, pages 22–41. Springer, 2013. ISBN 978-3-319-05118-5. doi: 10.1007/978-3-319-05119-2_3. URL https://doi.org/10.1007/978-3-319-05119-2_3.
- Teng Zhang, Peter Gebhard, and Oleg Sokolsky. SMEDL: Combining synchronous and asynchronous monitoring. In *Proc. of the 16th Int'l Conf on Runtime Verification (RV'16)*, volume 10012 of *LNCS*, pages 482–490. Springer, 2016.