

Mind the Gap! Exploiting Probabilistic Trace Expressions for Decentralized Runtime Verification with Gaps

Davide Ancona, Angelo Ferrando, and Viviana Mascardi

University of Genova, Genova 16146, Italy
name.surname@dibris.unige.it

Abstract. Multiagent Systems (MASs) are distributed systems composed by autonomous, reactive, proactive, heterogeneous communicating entities. In order to dynamically verify the behavior of such complex systems, a decentralized solution able to scale with the number of agents is necessary. When, for physical, infrastructural, or legal reasons, the monitor is not able to observe all the events emitted by the MAS, gaps are generated. In this paper we present a runtime verification decentralized approach to handle observation gaps in a MAS.

Keywords: Decentralized Runtime Verification · Probabilistic Trace Expressions · Observation Gaps · Multiagent Systems · Agent Interaction Protocols

1 Introduction and Motivations

Distributed Runtime Verification (DRV) is a relatively new research sub-field aimed at designing fault-tolerant distributed algorithms that monitor other distributed algorithms, with the end goal of developing lightweight software systems that are more efficient than traditional verification techniques [17,18]. The literature on DRV is almost limited [26,27,28,31,24,15] and becomes even more limited when we consider DRV of a special kind of systems: multiagent systems (MASs [36]). In the MAS area, in fact, we are only aware of our own previous works [25,4,7].

Another sub-field which is raising more and more attention in the RV area concerns partial observability of the monitored events which can cause gaps in the event traces [16,33,29,12]. Also in this case, when we consider MASs as the target system of the verification activity, we find very few works, all related with norm monitoring [22,23].

This paper addresses the two issues above, decentralized runtime verification of partially observable systems, in a MAS context. The findings presented in this work can be generalized and applied to other kinds of systems, but – for presentation purposes – we concentrate our investigation on MASs.

The main source of inspiration for our work is the paper by Stoller et al. [33], where the authors introduce *runtime verification with state estimation*. With

respect to a more standard RV approach, they are interested in checking system executions (traces) containing *gaps*. A gap represents the absence of information in the trace of observed events and corresponds to an execution point where the monitor knows that the system emitted some event, but does not know which one. In offline RV gaps in the trace logs are due to the process of sampling observed events in order to reduce the monitoring overhead. Gaps can also be met in online RV, where the system behavior is analyzed while the system is running and problems with the infrastructure, privacy and legal issues that prevent the monitor to observe some kind of events, faults in the monitor observation capabilities, may generate gaps. Although the problems raised by online and offline RV with gaps share many similarities, the online setting is much more challenging. Each time a gap is perceived, the monitor must make guesses on the possible actual events that the gap represents and save all the states generated by these guesses. A possibly huge logical tree-like structure with states as nodes, and moves from states to states as edges, represents the open possibilities¹. In offline RV, this logical tree-like structure can be explored following a depth-first search, requiring a limited amount of memory. If the RV takes place online, its exploration must follow a breadth-first strategy, with much more space needed to save the states, as the final trace of events is unknown and the levels of the structure are generated and explored at the same time. In order to cope with the state space explosion due to guesses in the online RV scenario, we propose to *decentralize* the monitoring activity.

RV decentralization is a very natural choice when the system under monitoring is a MAS, which is *distributed by definition*, and may improve *efficiency*, as the verification process can be spread on different machines improving performance; *scalability*, as under some conditions depending on the protocol [7,25] it is possible to associate one monitor with each agent in the MAS, keeping under control the RV complexity even when the number of agents grows; *feasibility*, as for physical/logical/legal reasons one single monitor might not be able to observe all the events generated by the MAS.

The feature that is usually subject to verification (both static and dynamic) in a MAS is its *communicative behavior* [13,14,21,34,37,20,32,35,10]. With respect to [33], in this work we do not aim at verifying temporal properties. Rather, we want to check the conformance of the MAS actual communicative behavior to an Agent Interaction Protocol (AIP) that models the allowed interactions among agents, under the hypothesis that some interactions could not be observed. The research question we address is thus *how to evaluate the probability that a MAS satisfies an AIP, in the presence of gaps*.

In [11] we introduced Probabilistic Trace Expressions (PTEs) and the theory behind them. In this work we take a more pragmatcal perspective and we show how to use PTEs for decentralized RV of AIPs within MASs with gaps.

¹ In the remainder we will use the term “branch” to denote paths in this logical structure, and we will sometime use “states” meaning “the final states of all the possible branches”, when this does not generate confusion.

2 Background

2.1 Probabilistic Trace Expressions

Trace expressions [5,1,19,3,2,8,9,6] are based on the notions of *event* and *event type*. We denote by \mathcal{E} the fixed universe of events subject to monitoring. An event trace over \mathcal{E} is a possibly infinite sequence of events in \mathcal{E} , and a trace expression over \mathcal{E} denotes a set of event traces over \mathcal{E} . Trace expressions are built on top of event types (chosen from a set \mathcal{ET}), each specifying a subset of events in \mathcal{E} . A trace expression $\tau \in \mathcal{T}$ represents a set of possibly infinite event traces, and is defined on top of the following operators:

- ϵ (empty trace), denoting the singleton set $\{\epsilon\}$ containing the empty event trace ϵ .
- $\vartheta:\tau$ (*prefix*), denoting the set of all traces whose first event e matches the event type ϑ , and the remaining part is a trace of τ .
- $\tau_1\cdot\tau_2$ (*concatenation*), denoting the set of all traces obtained by concatenating the traces of τ_1 with those of τ_2 .
- $\tau_1\wedge\tau_2$ (*intersection*), denoting the intersection of the traces of τ_1 and τ_2 .
- $\tau_1\vee\tau_2$ (*union*), denoting the union of the traces of τ_1 and τ_2 .
- $\tau_1|\tau_2$ (*shuffle*), denoting the set obtained by shuffling the traces of τ_1 with the traces of τ_2 .

Trace expressions support recursion through cyclic terms expressed by finite sets of recursive syntactic equations, as supported by modern Prolog systems.

A probabilistic trace expression is a trace expression where event types have a probability associated with them [11]. PTEs are suitable to manage guesses in the presence of observation gaps; in order for this management to work, we assume that *each gap represents one single unobserved event*.

As an example, the probabilistic trace expression

$$\tau = e_1[0.2]:\tau_1\vee e_2[0.8]:(\tau_2|\tau_3)$$

represents the protocol where we can accept the event e_1 with probability 0.2, or, the event e_2 with probability 0.8. If we consume the event e_1 , we go to the new state τ_1 , while, if we consume e_2 , we go to a state where we can have all possible interleaving of τ_2 and τ_3 . If there is a gap in the monitoring activity and the monitor is not able to observe which event took place, it can nevertheless make its guesses which involve e_1 and e_2 , associate a probability with each of them, and keep both possibilities.

Like a “normal” trace expression, a probabilistic trace expression τ can be seen as the current state of a protocol that started in some initial state τ_{init} and reached τ after n events $O_1\dots O_n$ took place. These events moved τ_{init} to τ through intermediate states $\tau_{q1}, \tau_{q2}, \dots, \tau_{qn} = \tau$. If we denote with $\tau \xrightarrow{O} \tau'$ the transition from state τ to state τ' due to the event O taking place and being observed, we may write

$$\tau_{init} \xrightarrow{O_1} \tau_{q1} \xrightarrow{O_2} \tau_{q2} \xrightarrow{O_3} \tau_{q3} \dots \xrightarrow{O_n} \tau_{qn}, \text{ where } \tau_{qn} = \tau.$$

In order to properly manage probabilities, it is convenient to associate with τ – in an explicit and easily computable way – the probability of the protocol to have reached τ starting from τ_{init} and having observed $O_1 \dots O_n$.

We define a “probabilistic trace expression state” the triple consisting of a trace expression τ , a sequence of events $O_1 \dots O_n$ observed before reaching τ , and the probability π_τ that the protocol reached τ . We represent the state with the notation $\langle \tau, \pi_{tr}, O_1 \dots O_n \rangle$.

2.2 Decentralized MAS Monitoring with DecAMon

In [25] we presented the DecAMon algorithm to decentralize agent interaction protocols modeled using trace expressions. There, we defined the notion of “monitoring safe” partition. A partition can be used to drive the distribution process. To decentralize the monitoring activity, we project the global AIP onto each subset of agents belonging to the partition, where by “projection” we mean that we maintain only the interactions involving agents in the chosen subset. In general, not all the partitions can be used for the RV decentralization. A partition that can be used to decentralize the RV of a protocol is called “monitoring safe” and the algorithm presented in [25] generates all the monitoring safe partitions for a given AIP.

Since under the conditions considered in this paper we may observe gaps, we could not have only one single state representing the current situation of the protocol, like it happens in our previous works; instead, we have to maintain all the states that may be possibly reached “via the gaps”. As already anticipated, each state can be represented as a tuple $\langle \tau, \pi, evs \rangle$, where τ is the PTE representing the current state of the protocol and π is the joint probability that the sequence of events evs is compliant with τ [11].

Let us name M_0 the set of possible initial states of the monitor (as there may be more than one). The number 0 stands for the *0th* iteration, since at the beginning we have not consumed any event yet. We can first run DecAMon on the global AIP to find a good set of monitoring safe partitions and, after that, we can use one of them to project the τ s in M_0 onto the subsets of the agents. Once we have obtained the distributed versions of the initial τ s via projection, we can generate one monitor for each partition, and decentralize the RV.

The combination of decentralization and lack of information calls for a synchronized management of gaps. Since each monitor has a different state representing its current protocol evolution, when there is an observation gap, each monitor can have different opinions about which are the correct events that might suitably “fill the gap”. The local perspectives can be compared and used by the monitors to cut wrong guesses, and hence wrong states, on the basis of distributed knowledge. Despite the overhead due to synchronization, this approach may dramatically improve performance, as discussed in the next sections.

3 Handling Gaps in Decentralized RV

Gaps represent lack of information, thus a point (or points) in the event trace where the monitor does not know what event had been actually generated by the system under monitoring. In the remainder we will write that “gaps can be observed”, in the sense that a monitor can realize that something went wrong and that an event was generated by the system, and not correctly observed. We also assume that, in a decentralized setting, when one monitor “observes a gap”, all the monitors “observe a gap” as well. From a technical viewpoint, this could be obtained by forcing one monitor to inform the others when it observes a gap. This would require some shared clock among the monitors as, in order for our algorithm to work, the gap must take place at the same time for all the monitors hence raising clock synchronization issues. Given that these issues are well known and well studied in distributed systems [30], we leave them out of our investigation.

When a centralized monitor observes a gap, since it is the only monitor checking the event trace w.r.t. the AIP specification, it can make guesses on what the gap is and reason on its own guesses, eventually tagging some of them as wrong due to successive observations. When there are many monitors, each one monitoring a subset of the agents, and hence a sub-protocol of the global AIP, each monitor can still suppose what the observed gap is, but the reasoning on its suppositions must be shared with the others. This sharing phase among the monitors is crucial, because it allows them to cut wrong branches on the basis of what other monitors suppose, or what they are fully sure of.

Let us consider two monitors m_1 and m_2 that observe a gap. Given that the protocols driving the two monitors are different, although being derived via projection from the same global protocol, m_1 might suppose that the events admissible for filling the observed gap are e_1 and e_2 , while m_2 could instead suppose that admissible events are e_2 and e_3 . Both m_1 and m_2 must keep track of these possibilities in their local knowledge bases, and – so far – they do not need to share they guesses.

Let us now suppose that in the current state of m_1 , in the branch where e_1 was supposed to have taken place, the only successive possible event is e_4 , while in the branch for e_2 the only possible event is e_5 . If, after the gap, m_1 observes e_5 , it can cut the branch where the gap was associated with e_1 , because e_5 would not be allowed after e_1 . The gap before e_5 , that could be filled in principle by e_1 and e_2 , becomes bound - “without any doubt”² - to e_2 . After having found

² Modulo the assumption that observed events are compliant with the foreseen protocol. Gaps may inevitably generate false negatives. In this case, m_1 assumes that the gap was e_2 because this would be consistent with the successive observation of e_5 and with the protocol to be respected. If the gap were any other event, a protocol violation would have taken place and m_1 should have raised a protocol monitoring exception. Depending on the protocol, the violation could be recognized later on, or never. Suppose for example an infinite protocol where only a s are allowed. A gap will be necessarily filled with a even if the actual event was b , and if the successive observed events are all a s, the violation will never be discovered.

the right value for the gap and cut one branch, m_1 informs m_2 allowing it to cut the branch where the value for that gap was guessed to be e_3 . In this way, both m_1 and m_2 can continue the verification process supposing that the unobserved event represented by the gap was e_2 , with some given probability due to the probability associated with e_2 in the PTE modelling the protocol.

Before presenting the decentralized monitoring algorithm, we make some considerations on the kind of gaps a monitor can observe. So far, we considered generic events. This is correct and consistent with the general approach presented in [11], but in a MAS scenario where PTEs model agent interaction protocols we can be more specific. In this scenario, in fact, the universe of events is *Msgs*, namely the universe of the possible messages among agents. Such special events can be represented as $a_1 \xrightarrow{c} a_2$, meaning that agent a_1 sends a message to a_2 with content c . Since messages are composed by (at least) three mandatory components, sender, receiver and content, there can be many partially instantiated gaps such as:

- $gap(a_1 \xrightarrow{\quad} a_2)$, where the content of the message is unknown;
- $gap(- \xrightarrow{m} a_2)$, where the sender is unknown;
- $gap(a_1 \xrightarrow{m} -)$, where the receiver is unknown.

Although, for sake of clarity, in the sequel we consider gaps where neither the sender, nor the receiver, nor the content are known (total absence of information), all the combinations of “information holes” are possible, and partially instantiated gaps may be exploited to reduce branches due to guesses. The algorithm presented in the next section can be easily adjusted to take partially instantiated gaps into account.

3.1 Synchronizing Decentralized Gaps Management

We present the algorithm used by the decentralized monitors to synchronize the gaps management, in order to cut useless branches and check the compliance of interactions with the protocol. When an event is generated by the system, two different situations can take place.

Case 1: The event is not a gap

If the event is not a gap, each monitor that observed it can use the event for updating its local state(s). If some branches have been removed as in the previous example involving m_1 and m_2 , the monitor has to inform the other monitors of the associations between gaps and events that are not admissible any longer. This phase can be reiterated until all the monitors have cut all the possible wrong branches, and have nothing more to say. After this synchronization stage, the monitoring process continues in the normal way.

Case 2: The event is a gap

To keep the presentation simple, we assume that gaps are observed by all the monitors at the same time. Each monitor guesses the events admissible to fill the gap, according to its local states. If the gap is partially instantiated (some of its components were correctly observed, like the sender, or the content, or both),

the monitor can use this information to reduce the set of possible candidate events.

The two cases can be seen as a *reduce* and *extend* stages, respectively. When the monitor observes a fully instantiated event it can invalidate zero, one or more branches. If the invalid branches contain gaps, the monitor can also invalidate the associations between these gaps and the guessed events, and can allow the other monitors to invalidate these associations as well via communication. On the other hand, observation of gaps generates as many branches as the events that, according to the AIP, could fill the gap. We can formalize this intuition in the following way.

Given M_0 as the set of global states $\{\langle \tau_1, \pi_1, [] \rangle, \dots, \langle \tau_n, \pi_n, [] \rangle\}$.

1. Distribute M_0 with respect to a given partition $P = \{\{ags_1\}, \dots, \{ags_{np}\}\}$, projecting the states onto subsets of the agents involved (the function Π projects an AIP τ onto a set of agents ags removing all the events whose sender and receiver do not belong to ags), obtaining

$$M_{0,\{ags_1\}} = \{\langle \Pi(\tau_1, \{ags_1\}), \pi_1, [] \rangle, \dots, \langle \Pi(\tau_n, \{ags_1\}), \pi_n, [] \rangle\}$$

...

$$M_{0,\{ags_{np}\}} = \{\langle \Pi(\tau_1, \{ags_{np}\}), \pi_1, [] \rangle, \dots, \langle \Pi(\tau_n, \{ags_{np}\}), \pi_n, [] \rangle\}$$

2. Each monitor observes only the event messages involving the agents belonging to its set ags_i :
 - (a) if the event message is a gap, the monitor guesses what it could be and generates as many states as the possible events (*extend*);
 - (b) if the event message is ground, the monitor can cut branches, and in this case it communicates with other monitors the gap values that are no longer admissible (*reduce*).
3. If, after observation of an event or because of information received from other monitors, the set of possible current states for a monitor m becomes empty, m stops the monitoring process, informs all the other monitors, and they also stop monitoring. The absence of possible current states for a monitor is due to a protocol violation that took place, preventing at least one monitor to move a further step. So, the system checked does not satisfy the agent interaction protocol and the associated probability is 0.
4. Else,
 - (a) if there are no events left to analyze, the monitoring process ends and the resulting probability is evaluated (see after how);
 - (b) else, repeat from step 2.

To be more clear, in step 2, given the current event message, each monitor queries its current state following the PTE operational semantics presented in [11] in order to check if the event message is admissible or not. In the updating phase, the monitors inform the others trying to cut not admissible branches.

If the monitoring process ends without violations detected and there are no more events left to analyze, each monitor stops with at least one admissible branch. Each monitor states its own evaluation of the probability that the system's behavior satisfies the agent interaction protocol. This probability can be computed summing up all the joint probabilities contained in all the final states, corresponding to the last nodes of the admissible branches. This leads to having one estimated value for each monitor: we can adopt different strategies to summarize the final, and global, one. One way could be to take the smallest value among all those estimated by all the monitors, meaning that we want to be cautious and we consider the lowest probability of acceptance; otherwise we could take the biggest value, meaning that we want to be optimistic since we trust the probabilities used in our specification. In other scenarios we could take the means of the values computed by the monitors, or a weighted means where weights model each monitor's trustability, or other domain-dependent strategies.

4 Example

We present a simple example helping us to show how the *extend* and *reduce* steps work. We consider a scenario involving a MAS involving four agents: $\{alice, bob, charlie, dave\}$. The set of events of our interest is the set of messages that these agents can use to communicate with each other.

Given the PTE

$$\begin{aligned}\tau &= \tau_1 \vee \tau_2 \\ \tau_1 &= alice \xrightarrow{msg_1} bob[0.7]:(bob \xrightarrow{msg_2} charlie[0.6]:\tau_1 | bob \xrightarrow{msg_3} dave[0.4]:\epsilon) \\ \tau_2 &= alice \xrightarrow{msg_4} dave[0.3]:(charlie \xrightarrow{msg_5} dave[0.3]:\epsilon | bob \xrightarrow{msg_3} dave[0.7]:\tau_2)\end{aligned}$$

We decentralize τ on each single agent, obtaining³:

$$\begin{aligned}M_{0,\{alice\}} &= \{\langle \Pi(\tau, \{alice\}), 1, [] \rangle\} = \{\langle \tau_{alice}, 1, [] \rangle\} \\ M_{0,\{bob\}} &= \{\langle \Pi(\tau, \{bob\}), 1, [] \rangle\} = \{\langle \tau_{bob}, 1, [] \rangle\} \\ M_{0,\{charlie\}} &= \{\langle \Pi(\tau, \{charlie\}), 1, [] \rangle\} = \{\langle \tau_{charlie}, 1, [] \rangle\} \\ M_{0,\{dave\}} &= \{\langle \Pi(\tau, \{dave\}), 1, [] \rangle\} = \{\langle \tau_{dave}, 1, [] \rangle\}\end{aligned}$$

where

$$\begin{aligned}\tau_{alice} &= \tau_{1_{alice}} \vee \tau_{2_{alice}} \\ \tau_{1_{alice}} &= alice \xrightarrow{msg_1} bob[0.7]:\tau_{1_{alice}} \\ \tau_{2_{alice}} &= alice \xrightarrow{msg_4} dave[0.3]:\tau_{2_{alice}} \\ \tau_{bob} &= \tau_{1_{bob}} \vee \tau_{2_{bob}} \\ \tau_{1_{bob}} &= alice \xrightarrow{msg_1} bob[0.7]:(bob \xrightarrow{msg_2} charlie[0.6]:\tau_1 | bob \xrightarrow{msg_3} dave[0.4]:\epsilon)\end{aligned}$$

³ The initial probability of each state is 1, since we do not want to influence the probability evaluation process (multiplication of probabilities).

$$\begin{aligned}
\tau_{2_{bob}} &= bob \xrightarrow{msg3} dave[0.7]:\tau_{2_{bob}} \\
\tau_{charlie} &= \tau_{1_{charlie}} \vee \tau_{2_{charlie}} \\
\tau_{1_{charlie}} &= bob \xrightarrow{msg2} charlie[0.6]:\tau_{1_{charlie}} \\
\tau_{2_{charlie}} &= charlie \xrightarrow{msg5} dave[0.3]:\tau_{2_{charlie}} \\
\tau_{dave} &= \tau_{1_{dave}} \vee \tau_{2_{dave}} \\
\tau_{1_{dave}} &= bob \xrightarrow{msg3} dave[0.4]:\epsilon \\
\tau_{2_{dave}} &= alice \xrightarrow{msg4} dave[0.3]:(charlie \xrightarrow{msg5} dave[0.3]:\epsilon | bob \xrightarrow{msg3} dave[0.7]:\tau_{2_{dave}})
\end{aligned}$$

Let us suppose that the monitors observe a *gap* now. Each monitor moves to a new set of states corresponding to the possible values for the *gap*.

$$\begin{aligned}
&M_{0,\{alice\}} \xrightarrow{gap} \{ \\
&\langle \tau_{1_{alice}}, 0.7, [gap(alice \xrightarrow{msg1} bob)] \rangle, \\
&\langle \tau_{2_{alice}}, 0.3, [gap(alice \xrightarrow{msg4} dave)] \rangle, \\
&\langle \tau_{alice}, 1, [gap(none)] \rangle \\
&\} = M_{1,\{alice\}} \\
&M_{0,\{bob\}} \xrightarrow{gap} \{ \\
&\langle (bob \xrightarrow{msg2} charlie[0.6]:\tau_1 | bob \xrightarrow{msg3} dave[0.4]:\epsilon), 0.7, [gap(alice \xrightarrow{msg1} bob)] \rangle, \\
&\langle \tau_{2_{bob}}, 0.7, [gap(bob \xrightarrow{msg3} dave)] \rangle, \\
&\langle \tau_{bob}, 1, [gap(none)] \rangle \\
&\} = M_{1,\{bob\}} \\
&M_{0,\{charlie\}} \xrightarrow{gap} \{ \\
&\langle \tau_{1_{charlie}}, 0.6, [gap(bob \xrightarrow{msg2} charlie)] \rangle, \\
&\langle \tau_{2_{charlie}}, 0.3, [gap(charlie \xrightarrow{msg5} dave)] \rangle, \\
&\langle \tau_{charlie}, 1, [gap(none)] \rangle \\
&\} = M_{1,\{charlie\}}, \\
&M_{0,\{dave\}} \xrightarrow{gap} \{ \\
&\langle \epsilon, 0.4, [gap(bob \xrightarrow{msg3} dave)] \rangle, \\
&\langle (charlie \xrightarrow{msg5} dave[0.3]:\epsilon | bob \xrightarrow{msg3} dave[0.7]:\tau_{2_{dave}}), 0.3, [gap(alice \xrightarrow{msg4} dave)] \rangle, \\
&\langle \tau_{dave}, 1, [gap(none)] \rangle \\
&\} = M_{1,\{dave\}}
\end{aligned}$$

the related gap-events associations, and may need to inform the others of some association which is no longer possible. In the previous example, one single iteration was enough.

As we already anticipated, the proposed approach may lead to false negatives, due to an optimistic approach of the monitors that stubbornly assume that observed events are compliant with the protocol, if there is just one possibility left to make such an assumption. Also in this example, the monitors gave the correctness of the ground event msg_2 (the second event observed) for granted. But let us suppose that the actual event masked by the *gap* was not msg_1 , but msg_4 , and that the successive message msg_2 was sent from *bob* to *charlie* by mistake and did not comply with the protocol. In this scenario, since the monitors do not know for sure what the first *gap* was, it is reasonable to consider msg_2 a valid message and hence cut the branch where the gap has been supposed to be msg_4 . This is a problem intrinsically related to the state estimation approach, since until it is acceptable to observe an event in a state, the monitors keep track of the related branch. Only when a monitor, observing an event, loses all its branches it can conclude that a protocol violation took place because some wrong assumption on gaps – confirmed by successive observations – had been made in the past. This delay in the error detection, which could also be infinite, can be reduced introducing a threshold on the probability that a branch must have to be considered valid. In this way, if after observing an event the probability associated with a branch becomes lower than a chosen threshold, the monitor can cut that branch and make error detection possibly quicker.

5 Experimental results

In our experiments we have considered the four following features:

1. the number of agents involved in the MAS we want to verify at runtime;
2. the number of *shuffled sub-PTEs* due to shuffle operators $|$ in the AIP: we name shuffled sub-PTE each portion of the PTE composed via a $|$, so for example $\tau_3 = alice \xrightarrow{msg_1} bob[0.7]:\epsilon \mid bob \xrightarrow{msg_3} dave[0.4]:\epsilon$ consists of 2 *shuffled sub-PTEs*; we point out that when decentralizing the monitoring, we can associate one different monitor with each shuffled sub-PTE, as shuffled sub-PTE are independent one from the other and can be monitored in a fully decentralized way;
3. the number of operators for each shuffled sub-PTE in the AIP;
4. the number of gaps contained in the analyzed traces.

In Table 1, we report the results of our experiments. For each row, we keep the number of shuffled sub-PTE, agents and operators fixed, while we change the length of the traces and the percentage of gaps inside each trace. For each row we executed many different runs and we have measured the total time required for recognizing the set of 300 randomly generated traces. We changed the number of gaps contained inside the traces and we tested both the centralized [11] and

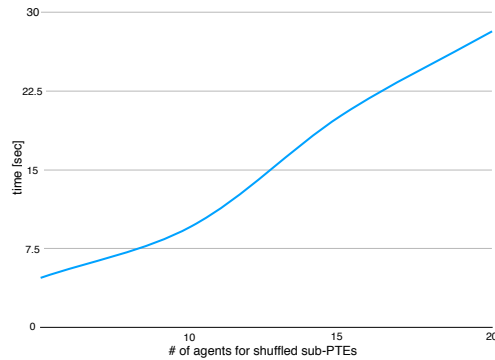
Table 1. Average time of the centralized and decentralized algorithms; “sh. PTE” stands for “shuffled sub-PTE”.

# sh. PTEs	# agents for sh. PTE	# operations for sh. PTE	Centralized [sec]	Decentralized [sec]
10	10	20	6.64	1.26
10	10	15	8.26	1.04
10	5	20	9.85	1.49
10	5	15	9.92	1.28
10	15	15	14.86	1.23
10	5	10	18.35	1.08
10	15	10	20.25	1.61
10	10	10	29.59	1.98
15	5	15	93.34	2.73
15	15	10	116.61	3.56
10	15	20	126.31	25.32
15	10	10	283.70	4.14
15	5	10	349.30	2.23
20	10	10	355.90	3.99
15	5	20	363.67	5.83
20	5	15	558.59	9.28
20	5	20	801.37	7.82
15	20	10	952.43	12.36
20	5	10	1223.85	10.64
20	15	10	1340.29	9.57
20	20	10	1727.26	2.89

the decentralized algorithms. In the following, we reported the graphics obtained from such executions.

Concerning the figures, *the traces used in our experiments contain only gaps* (namely, we run experiments in the worst possible scenario), so the algorithm makes only expansions and never reductions. We chose traces with only gaps to stress the algorithms as much as possible. In real scenarios gaps should be the exceptions, and perfectly observable events the norm.

In Figures 1 and 2, both the centralized and the decentralized algorithms seem to show linear complexity with respect the number of the agents involved, even if the decentralized algorithm has better performances.

**Fig. 1.** Centralized algorithm: changing number of agents.

In Figures 3 and 4, we can observe that the complexity of the centralized algorithm seems to grow in a quadratic way, while the decentralized one seems




Fig. 2. Decentralized algorithm: changing number of agents.

to grows linearly. This can be explained by the decentralization of the monitoring of shuffled sub-PTEs, as if we add one operator to each shuffled sub-PTE, the monitor in charge for that shuffled sub-PTE will need to manage one more operator only, whereas the centralized monitor will cope with as many new operators as the shuffled sub-PTEs in the trace expression. We point out that we use “seems to” to reflect that the complexities emerging from the figures have not been computed on the basis of the algorithm, but have been estimated on the basis of the experiments, and the behaviour in situations involving a limited number of agents, operators, shuffled sub-PTEs, might not be the actual asymptotic behaviour of the algorithm.




Fig. 3. Centralized algorithm: changing number of operators.

In Figures 5 and 6, we can appreciate the real advantages of decentralization, as – from the figures – it seems that we have an exponential complexity for the




Fig. 4. Decentralized algorithm: changing number of operators.

centralized algorithm and a pseudo-quadratic complexity for the decentralized one. We emphasise that in the decentralized case (Figure 6) we were able to run experiments with 40 shuffled sub-PTEs, while in the centralized case we had to stop with half shuffled sub-PTEs, and with an execution time hundred times higher. The number of shuffled sub-PTEs is indeed the feature which most impacts the algorithms performance, and this is not a surprise; intuitively, when we add a new shuffled sub-PTE we have to interleave it with all the already existent shuffled sub-PTEs. In the centralized case, this brings to a state explosion, while in the decentralized one, since we can decentralize the monitoring of each shuffled sub-PTEs, we simply have to add a new monitor. In this way, we can avoid the state explosion, even if the presence of a new monitor increases the exchange of messages among the monitors needed to synchronize information about gaps.

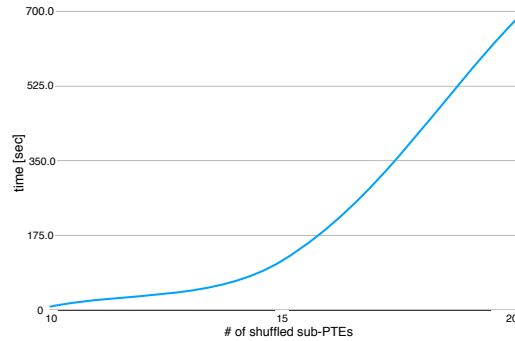


Fig. 5. Centralized algorithm: changing number of shuffled sub-PTEs.




Fig. 6. Decentralized algorithm: changing number of shuffled sub-PTEs.

6 Conclusions and Future Work

In this paper we presented a distributed approach to runtime verification where we may lack some pieces of information about observed events. With respect to standard runtime verification, the state estimation approach allows us to be more reliable, especially in scenarios where partial or total absence of information is frequent.

For the sake of clarity, we considered only totally uninstantiated gaps. This choice has been made to make the development of monitors easier. Naturally, the presence of part of information about the event could be used by the monitors in order to cut useless branches. We will extend our implementation to cope with partially instantiated gaps.

Another future work will be to consider a threshold in order to cut branches that are unreasonable to maintain, as the probability to be correct is too low. Fixed a threshold, a monitor will be able to remove all the branches with a joint probability associated with them lower than the chosen threshold. This will bring the advantage of anticipating the error detection and to prune useless branches related to unreasonable possibilities.

References

1. Ancona, D., Barbieri, M., Mascardi, V.: Constrained global types for dynamic checking of protocol conformance in multi-agent systems. In: SAC. pp. 1377–1379. ACM (2013)
2. Ancona, D., Bono, V., Bravetti, M., Campos, J., Castagna, G., Deniélou, P., Gay, S.J., Gesbert, N., Giachino, E., Hu, R., Johnsen, E.B., et al., F.M.: Behavioral types in programming languages. *Foundations and Trends in Programming Languages* **3**(2-3), 95–230 (2016)

3. Ancona, D., Briola, D., Ferrando, A., Mascardi, V.: Global protocols as first class entities for self-adaptive agents. In: AAMAS. pp. 1019–1029. ACM (2015)
4. Ancona, D., Briola, D., Ferrando, A., Mascardi, V.: MAS-DRiVe: a practical approach to decentralized runtime verification of agent interaction protocols. In: Santoro, C., Messina, F., Benedetti, M.D. (eds.) From Objects to Agents, 17th Workshop, WOA 2016. Proceedings. CEUR Workshop Proceedings, vol. 1664, pp. 35–43. CEUR-WS.org (2016), <http://ceur-ws.org/Vol-1664>
5. Ancona, D., Drossopoulou, S., Mascardi, V.: Automatic generation of self-monitoring mass from multiparty global session types in jason. In: Declarative Agent Languages and Technologies X - 10th International Workshop, DALT 2012, Valencia, Spain, June 4, 2012, Revised Selected Papers. pp. 76–95 (2012)
6. Ancona, D., Ferrando, A., Franceschini, L., Mascardi, V.: Parametric trace expressions for runtime verification of Java-like programs. In: FTfJP@ECOOP. pp. 10:1–10:6. ACM (2017)
7. Ancona, D., Ferrando, A., Franceschini, L., Mascardi, V.: Coping with bad agent interaction protocols when monitoring partially observable multiagent systems. In: PAAMS. Lecture Notes in Computer Science, vol. 10978, pp. 59–71. Springer (2018)
8. Ancona, D., Ferrando, A., Mascardi, V.: Comparing trace expressions and linear temporal logic for runtime verification. In: Theory and Practice of Formal Methods. LNCS, vol. 9660, pp. 47–64 (2016)
9. Ancona, D., Ferrando, A., Mascardi, V.: Parametric runtime verification of multi-agent systems. In: AAMAS. pp. 1457–1459. ACM (2017)
10. Ancona, D., Ferrando, A., Mascardi, V.: Improving flexibility and dependability of remote patient monitoring with agent-oriented approaches (2018), IJAOSE. To appear.
11. Ancona, D., Ferrando, A., Mascardi, V.: Probabilistic trace expressions for runtime verification with partial observability (2018), submitted to RV2018, available from <https://www.disi.unige.it/person/MascardiV/Download/Probabilistic-Trace-Expressions.pdf>
12. Babaee, R., Gurfinkel, A., Fischmeister, S.: Prevent : A predictive run-time verification framework using statistical learning. In: SEFM. Lecture Notes in Computer Science, vol. 10886, pp. 205–220. Springer (2018)
13. Baldoni, M., Baroglio, C., Capuzzimati, F.: A commitment-based infrastructure for programming socio-technical systems. ACM Trans. Internet Techn. **14**(4), 23:1–23:23 (2014). <https://doi.org/10.1145/2677206>
14. Baldoni, M., Baroglio, C., Capuzzimati, F., Micalizio, R.: Exploiting social commitments in programming agent interaction. In: Proc. of PRIMA 2015. LNCS, vol. 9387, pp. 566–574. Springer (2015)
15. Bartocci, E.: Sampling-based decentralized monitoring for networked embedded systems. In: HAS. EPTCS, vol. 124, pp. 85–99 (2013)
16. Bartocci, E., Grosu, R., Katsaros, P., Ramakrishnan, C.R., Smolka, S.A.: Model repair for probabilistic systems. In: TACAS. Lecture Notes in Computer Science, vol. 6605, pp. 326–340. Springer (2011)
17. Bonakdarpour, B., Frgaignaud, P., Rajsbaum, S., Travers, C. (eds.): Bertinoro Seminar on Distributed Runtime Verification, May 2016, Available from <http://www.labri.fr/perso/travers/DRV2016/> (2016)
18. Bonakdarpour, B., Frgaignaud, P., Rajsbaum, S., Travers, C.: Challenges in fault-tolerant distributed runtime verification. In: Margaria, T., Steffen, B. (eds.) Leveraging Applications of Formal Methods, Verification and Validation: Discussion,

- Dissemination, Applications: 7th International Symposium, ISoLA 2016. Proceedings, Part II. pp. 363–370. Springer (2016)
19. Briola, D., Mascardi, V., Ancona, D.: Distributed runtime verification of JADE multiagent systems. In: IDC. Studies in Computational Intelligence, vol. 570, pp. 81–91. Springer (2014)
 20. Chopra, A.K., Christie, S., Singh, M.P.: Splee: A declarative information-based language for multiagent interaction protocols. In: Proc. of AAMAS 2017. pp. 1054–1063. ACM (2017)
 21. Chopra, A.K., Singh, M.P.: Cupid: Commitments in relational algebra. In: Proc. of AAAI 2015. pp. 2052–2059. AAAI Press (2015)
 22. Criado, N., Such, J.M.: Norm monitoring under partial action observability. *IEEE Trans. Cybernetics* **47**(2), 270–282 (2017)
 23. Criado Pacheco, N.: Resource-bounded norm monitoring in multi-agent systems. *Journal Artificial Intelligence Research* p. 1 (4 2018)
 24. Falcone, Y., Cornebize, T., Fernandez, J.C.: Efficient and generalized decentralized monitoring of regular languages. In: Ábrahám, E., Palamidessi, C. (eds.) Formal Techniques for Distributed Objects, Components, and Systems: 34th IFIP WG 6.1 International Conference, FORTE 2014. Proceedings. pp. 66–83. Springer (2014)
 25. Ferrando, A., Ancona, D., Mascardi, V.: Decentralizing MAS monitoring with decamon. In: Larson, K., Winikoff, M., Das, S., Durfee, E.H. (eds.) Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems, AAMAS 2017, São Paulo, Brazil, May 8-12, 2017. pp. 239–248. ACM (2017), <http://dl.acm.org/citation.cfm?id=3091164>
 26. Fraigniaud, P., Rajsbaum, S., Roy, M., Travers, C.: The opinion number of set-agreement. In: Aguilera, M.K., Querzoni, L., Shapiro, M. (eds.) Principles of Distributed Systems: 18th International Conference, OPODIS 2014. Proceedings. pp. 155–170. Springer (2014)
 27. Fraigniaud, P., Rajsbaum, S., Travers, C.: On the number of opinions needed for fault-tolerant run-time monitoring in distributed systems. In: Bonakdarpour, B., Smolka, S.A. (eds.) Runtime Verification: 5th International Conference, RV 2014. Proceedings. pp. 92–107. Springer (2014)
 28. Herlihy, M.: Wait-free synchronization. *ACM Trans. Program. Lang. Syst.* **13**(1), 124–149 (Jan 1991)
 29. Joshi, Y., Tchamgoue, G.M., Fischmeister, S.: Runtime verification of LTL on lossy traces. In: SAC. pp. 1379–1386. ACM (2017)
 30. Lamport, L.: Time, clocks, and the ordering of events in a distributed system. *Commun. ACM* **21**(7), 558–565 (Jul 1978). <https://doi.org/10.1145/359545.359563>, <http://doi.acm.org/10.1145/359545.359563>
 31. Mostafa, M., Bonakdarpour, B.: Decentralized runtime verification of LTL specifications in distributed systems. In: Parallel and Distributed Processing Symposium, IEEE International Conference, IPDPS 2015. Proceedings. pp. 494–503 (2015)
 32. Singh, M.P.: Information-driven interaction-oriented programming: BSPL, the blindingly simple protocol language. In: Proc. of AAMAS 2011. pp. 491–498. IFAAMAS (2011)
 33. Stoller, S.D., Bartocci, E., Seyster, J., Grosu, R., Havelund, K., Smolka, S.A., Zadok, E.: Runtime verification with state estimation. In: RV. Lecture Notes in Computer Science, vol. 7186, pp. 193–207. Springer (2011)
 34. Winikoff, M., Liu, W., Harland, J.: Enhancing commitment machines. In: Proc. of DALT 2004, Revised Selected Papers. LNCS, vol. 3476, pp. 198–220. Springer (2004)

35. Winikoff, M., Yadav, N., Padgham, L.: A new hierarchical agent protocol notation. *Autonomous Agents and Multi-Agent Systems* **32**(1), 59–133 (2018)
36. Wooldridge, M., Jennings, N.R.: Intelligent agents: theory and practice. *Knowledge Eng. Review* **10**(2), 115–152 (1995)
37. Yolum, P., Singh, M.P.: Commitment machines. In: *Proc. of ATAL 2001, Revised Papers*. vol. 2333, pp. 235–247. Springer (2002)