Activity Report 2015

Project-Team MEXICO

Modeling and Exploitation of Interaction and Concurrency

Table of contents

Project-Team MEXICO

Creation of the Team: 2009 March 01, updated into Project-Team: 2011 January 01

Keywords:

Computer Science and Digital Science:

2.3. - Embedded and cyber-physical systems
2.4. - Reliability, certification
4.5. - Formal methods for security
6.4.3. - Observability and Controlability
7.1. - Parallel and distributed algorithms
7.11. - Performance evaluation
7.2. - Discrete mathematics, combinatorics
7.3. - Operations research, optimization, game theory
7.4. - Logic in Computer Science
7.9. - Graph theory

Other Research Topics and Application Domains:

1.1.11. - Systems biology
1.1.12. - Synthetic biology
1.1.2. - Molecular biology
1.1.3. - Cellular biology
6.3.1. - Web
6.3.3. - Network services
7.1. - Traffic management
7.2.1. - Smart vehicles

1. Members

Research Scientists
Benedikt Bollig [CNRS, Researcher, HdR]
Stefan Haar [Team leader, Inria, Senior Researcher, HdR]

Faculty Members
Thomas Chatain [ENS Cachan, Associate Professor, HdR]
Paul Gastin [ENS Cachan, Professor, HdR]
Serge Haddad [ENS Cachan, Professor, HdR]
Claudine Picaronny [ENS Cachan, Associate Professor]
Stefan Schwoon [ENS Cachan, Associate Professor, HdR]

Engineer
Alban Linard [Inria, until Oct 2015]

PhD Students
Yann Duplouy [ENS Cachan, granted by Institut de Recherche SystemX]
Engel Lefaucheux [ENS Cachan/Inria Rennes]
Simon Theissing [Inria, granted by Institut de Recherche SystemX]

Administrative Assistant
Thida Iem [Inria]

2. Overall Objectives

2.1. Scientific Objectives

2.1.1. Introduction

In the increasingly networked world, reliability of applications becomes ever more critical as the number of users of, e.g., communication systems, web services, transportation etc., grows steadily. Management of networked systems, in a very general sense of the term, therefore is a crucial task, but also a difficult one.

MExICo strives to take advantage of distribution by orchestrating cooperation between different agents that observe local subsystems, and interact in a localized fashion.

The need for applying formal methods in the analysis and management of complex systems has long been recognized. It is with much less unanimity that the scientific community embraces methods based on asynchronous and distributed models. Centralized and sequential modeling still prevails.

However, we observe that crucial applications have increasing numbers of users, that networks providing services grow fast both in the number of participants and the physical size and degree of spatial distribution. Moreover, traditional *isolated* and *proprietary* software products for local systems are no longer typical for emerging applications.

In contrast to traditional centralized and sequential machinery for which purely functional specifications are efficient, we have to account for applications being provided from diverse and non-coordinated sources. Their distribution (e.g. over the Web) must change the way we verify and manage them. In particular, one cannot ignore the impact of quantitative features such as delays or failure likelihoods on the functionalities of composite services in distributed systems.

We thus identify three main characteristics of complex distributed systems that constitute research challenges:

- *Concurrency* of behavior;
- *Interaction* of diverse and semi-transparent components; and
- management of *Quantitative* aspects of behavior.

2.1.2. Concurrency

The increasing size and the networked nature of communication systems, controls, distributed services, etc. confront us with an ever higher degree of parallelism between local processes. This field of application for our work includes telecommunication systems and composite web services. The challenge is to provide sound theoretical foundations and efficient algorithms for management of such systems, ranging from controller synthesis and fault diagnosis to integration and adaptation. While these tasks have received considerable attention in the *sequential* setting, managing *non-sequential* behavior requires profound modifications for existing approaches, and often the development of new approaches altogether. We see concurrency in distributed systems as an opportunity rather than a nuisance. Our goal is to *exploit* asynchronicity and distribution as an advantage. Clever use of adequate models, in particular *partial order semantics* (ranging from Mazurkiewicz traces to event structures to MSCs) actually helps in practice. In fact, the partial order vision allows us to make causal precedence relations explicit, and to perform diagnosis and test for the dependency between events. This is a conceptual advantage that interleaving-based approaches cannot match. The two key features of our work will be *(i)* the exploitation of concurrency by using asynchronous models with partial order semantics, and *(ii)* distribution of the agents performing management tasks.

2.1.3. Interaction

Systems and services exhibit non-trivial *interaction* between specialized and heterogeneous components. A coordinated interplay of several components is required; this is challenging since each of them has only a limited, partial view of the system's configuration. We refer to this problem as *distributed synthesis* or *distributed control*. An aggravating factor is that the structure of a component might be semi-transparent, which requires a form of *grey box management*.

2.1.4. Quantitative Features

Besides the logical functionalities of programs, the *quantitative* aspects of component behavior and interaction play an increasingly important role.

- *Real-time* properties cannot be neglected even if time is not an explicit functional issue, since transmission delays, parallelism, etc, can lead to time-outs striking, and thus change even the logical course of processes. Again, this phenomenon arises in telecommunications and web services, but also in transport systems.

- In the same contexts, *probabilities* need to be taken into account, for many diverse reasons such as unpredictable functionalities, or because the outcome of a computation may be governed by race conditions.

- Last but not least, constraints on *cost* cannot be ignored, be it in terms of money or any other limited resource, such as memory space or available CPU time.

2.1.5. Evolution and Perspectives

Since the creation of *MExICo*, the weight of *quantitative* aspects in all parts of our activities has grown, be it in terms of the models considered (weighted automata and logics), be it in transforming verification or diagnosis verdict into probabilistic statements (probabilistic diagnosis, statistical model checking), or within the recently started SystemX cooperation on supervision in multi-modal transport systems. This trend is certain to continue over the next couple of years, along with the growing importance of diagnosis and control issues.

In another development, the theory and use of partial order semantics has gained momentum in the past four years, and we intend to further strengthen our efforts and contacts in this domain to further develop and apply partial-order based deduction methods.

As concerns the study of interaction, our progress has been thus far less in the domain of *distributed* approaches than in the analysis of *system composition*, such as in networks of untimed or timed automata. While continuing this line of study, we also intend to turn more strongly towards distributed *algorithms*, namely in terms of parametrized verification methods.

3. Research Program

3.1. Concurrency

Participants: Benedikt Bollig, Thomas Chatain, Paul Gastin, Stefan Haar, Serge Haddad, Stefan Schwoon.

Concurrency: Property of systems allowing some interacting processes to be executed in parallel.

Diagnosis: The process of deducing from a partial observation of a system aspects of the internal states or events of that system; in particular, *fault diagnosis* aims at determining whether or not some non-observable fault event has occurred.

Conformance Testing: Feeding dedicated input into an implemented system IS and deducing, from the resulting output of I, whether I respects a formal specification S.

3.1.1. Introduction

It is well known that, whatever the intended form of analysis or control, a *global* view of the system state leads to overwhelming numbers of states and transitions, thus slowing down algorithms that need to explore the state space. Worse yet, it often blurs the mechanics that are at work rather than exhibiting them. Conversely, respecting concurrency relations avoids exhaustive enumeration of interleavings. It allows us to focus on 'essential' properties of non-sequential processes, which are expressible with causal precedence relations. These precedence relations are usually called causal (partial) orders. Concurrency is the explicit absence of such a precedence between actions that do not have to wait for one another. Both causal orders and concurrency are in fact essential elements of a specification. This is especially true when the specification is constructed

in a distributed and modular way. Making these ordering relations explicit requires to leave the framework of state/interleaving based semantics. Therefore, we need to develop new dedicated algorithms for tasks such as conformance testing, fault diagnosis, or control for distributed discrete systems. Existing solutions for these problems often rely on centralized sequential models which do not scale up well.

3.1.2. *Diagnosis*

Participants: Benedikt Bollig, Stefan Haar, Serge Haddad, Stefan Schwoon.

Fault Diagnosis for discrete event systems is a crucial task in automatic control. Our focus is on *event oriented* (as opposed to *state oriented*) model-based diagnosis, asking e.g. the following questions:
given a - potentially large - *alarm pattern* formed of observations,

- what are the possible *fault scenarios* in the system that *explain* the pattern ?

- Based on the observations, can we deduce whether or not a certain - invisible - fault has actually occurred ?

Model-based diagnosis starts from a discrete event model of the observed system - or rather, its relevant aspects, such as possible fault propagations, abstracting away other dimensions. From this model, an extraction or unfolding process, guided by the observation, produces recursively the explanation candidates.

In asynchronous partial-order based diagnosis with Petri nets [63], [64], [68], one unfolds the *labelled product* of a Petri net model $\mathcal{N}$ and an observed alarm pattern $\mathcal{A}$, also in Petri net form. We obtain an acyclic net giving partial order representation of the behaviors compatible with the alarm pattern. A recursive online procedure filters out those runs *(configurations)* that explain *exactly* $\mathcal{A}$. The Petri-net based approach generalizes to dynamically evolving topologies, in dynamical systems modeled by graph grammars, see [46]

3.1.2.1. Observability and Diagnosability

Diagnosis algorithms have to operate in contexts with low observability, i.e., in systems where many events are invisible to the supervisor. Checking *observability* and *diagnosability* for the supervised systems is therefore a crucial and non-trivial task in its own right. Analysis of the relational structure of occurrence nets allows us to check whether the system exhibits sufficient visibility to allow diagnosis. Developing efficient methods for both verification of *diagnosability checking* under concurrency, and the *diagnosis* itself for distributed, composite and asynchronous systems, is an important field for *MExICo*.

3.1.2.2. Distribution

Distributed computation of unfoldings allows one to factor the unfolding of the global system into smaller *local* unfoldings, by local supervisors associated with sub-networks and communicating among each other. In [64], [48], elements of a methodology for distributed computation of unfoldings between several supervisors, underwritten by algebraic properties of the category of Petri nets have been developed. Generalizations, in particular to Graph Grammars, are still do be done.

Computing diagnosis in a distributed way is only one aspect of a much vaster topic, that of *distributed diagnosis* (see [60], [72]). In fact, it involves a more abstract and often indirect reasoning to conclude whether or not some given invisible fault has occurred. Combination of local scenarios is in general not sufficient: the global system may have behaviors that do not reveal themselves as faulty (or, dually, non-faulty) on any local supervisor's domain (compare [45], [51]). Rather, the local diagnosers have to join all *information* that is available to them locally, and then deduce collectively further information from the combination of their views. In particular, even the *absence* of fault evidence on all peers may allow to deduce fault occurrence jointly, see [77], [78]. Automatizing such procedures for the supervision and management of distributed and locally monitored asynchronous systems is a long-term goal to which *MExICo* hopes to contribute.

3.1.3. *Contextual nets*

Participant: Stefan Schwoon.

Assuring the correctness of concurrent systems is notoriously difficult due to the many unforeseeable ways in which the components may interact and the resulting state-space explosion. A well-established approach to alleviate this problem is to model concurrent systems as Petri nets and analyse their unfoldings, essentially an acyclic version of the Petri net whose simpler structure permits easier analysis [62].

However, Petri nets are inadequate to model concurrent read accesses to the same resource. Such situations often arise naturally, for instance in concurrent databases or in asynchronous circuits. The encoding tricks typically used to model these cases in Petri nets make the unfolding technique inefficient. Contextual nets, which explicitly do model concurrent read accesses, address this problem. Their accurate representation of concurrency makes contextual unfoldings up to exponentially smaller in certain situations. An abstract algorithm for contextual unfoldings was first given in [47]. In recent work, we further studied this subject from a theoretical and practical perspective, allowing us to develop concrete, efficient data structures and algorithms and a tool (Cunf) that improves upon existing state of the art. This work led to the PhD thesis of César Rodríguez in 2014 .

Contextual unfoldings deal well with two sources of state-space explosion: concurrency and shared resources. Recently, we proposed an improved data structure, called *contextual merged processes* (CMP) to deal with a third source of state-space explosion, i.e. sequences of choices. The work on CMP [79] is currently at an abstract level. In the short term, we want to put this work into practice, requiring some theoretical groundwork, as well as programming and experimentation.

Another well-known approach to verifying concurrent systems is *partial-order reduction*, exemplified by the tool SPIN. Although it is known that both partial-order reduction and unfoldings have their respective strengths and weaknesses, we are not aware of any conclusive comparison between the two techniques. Spin comes with a high-level modeling language having an explicit notion of processes, communication channels, and variables. Indeed, the reduction techniques implemented in Spin exploit the specific properties of these features. On the other side, while there exist highly efficient tools for unfoldings, Petri nets are a relatively general low-level formalism, so these techniques do not exploit properties of higher language features. Our work on contextual unfoldings and CMPs represents a first step to make unfoldings exploit richer models. In the long run, we wish raise the unfolding technique to a suitable high-level modelling language and develop appropriate tool support.

3.1.4. *Verification of Concurrent Recursive Programs*
Participants: Benedikt Bollig, Paul Gastin, Stefan Schwoon.

3.1.5. *Dynamic and parameterized concurrent systems*
Participants: Benedikt Bollig, Paul Gastin.

In the past few years, our research has focused on concurrent systems where the architecture, which provides a set of processes and links between them, is *static* and *fixed in advance*. However, the assumption that the set of processes is fixed somehow seems to hinder the application of formal methods in practice. It is not appropriate in areas such as mobile computing or ad-hoc networks. In concurrent programming, it is actually perfectly natural to design a program, and claim its correctness, independently of the number of processes that participate in its execution. There are, essentially, two kinds of systems that fall into this category. When the process architecture is static but unknown, it is a parameter of the system; we then call a system *parameterized*. When, on the other hand, the process architecure is generated at runtime (i.e., process creation is a communication primitive), we say that a system is *dynamic*. Though parameterized and dynamic systems have received increasing interest in recent years, there is, by now, no canonical approach to modeling and verifying such systems. Our research program aims at the development of *a theory of parameterized and dynamic concurrent systems*. More precisely, our goal is a *unifying* theory that lays algebraic, logical, and automata-theoretic foundations to support and facilitate the study of parameterized and dynamic concurrent systems. Such theories indeed exist in non-parameterized settings where the number of processes and the way they are connected are fixed in advance. However, parameterized and dynamic systems lack such foundations and often restrict to very particular models with specialized verification techniques.

3.1.6. *Testing*

Participants: Benedikt Bollig, Paul Gastin, Stefan Haar.

3.1.6.1. Introduction

The gap between specification and implementation is at the heart of research on formal testing. The general *conformance testing problem* can be defined as follows: Does an implementation $\mathcal{M}'$ conform a given specification $\mathcal{M}$? Here, both $\mathcal{M}$ and $\mathcal{M}'$ are assumed to have input and output channels. The formal model $\mathcal{M}$ of the specification is entirely known and can be used for analysis. On the other hand, the implementation $\mathcal{M}'$ is unknown but interacts with the environment through observable input and output channels. So the behavior of $\mathcal{M}'$ is partially controlled by input streams, and partially observable via output streams. The Testing problem consists in computing, from the knowledge of $\mathcal{M}$, *input streams* for $\mathcal{M}'$ such that observation of the resulting output streams from $\mathcal{M}'$ allows to determine whether $\mathcal{M}'$ conforms to $\mathcal{M}$ as intended.

In this project, we focus on distributed or asynchronous versions of the conformance testing problem. There are two main difficulties. First, due to the distributed nature of the system, it may not be possible to have a unique global observer for the outcome of a test. Hence, we may need to use *local* observers which will record only *partial views* of the execution. Due to this, it is difficult or even impossible to reconstruct a coherent global execution. The second difficulty is the lack of global synchronization in distributed asynchronous systems. Up to now, models were described with I/O automata having a centralized control, hence inducing global synchronizations.

3.1.6.2. Asynchronous Testing

Since 2006 and in particular during his sabbatical stay at the University of Ottawa, Stefan Haar has been working with Guy-Vincent Jourdan and Gregor v. Bochmann of UOttawa and Claude Jard of IRISA on asynchronous testing. In the synchronous (sequential) approach, the model is described by an I/O automaton with a centralized control and transitions labeled with individual input or output actions. This approach has known limitations when inputs and outputs are distributed over remote sites, a feature that is characteristic of , e.g., web computing. To account for concurrency in the system, they have developed in [70], [52] asynchronous conformance testing for automata with transitions labeled with (finite) partial orders of I/O. Intuitively, this is a "big step" semantics where each step allows concurrency but the system is synchronized before the next big step. This is already an important improvement on the synchronous setting. The non-trivial challenge is now to cope with fully asynchronous specifications using models with decentralized control such as Petri nets.

3.1.6.3. Near Future

Completion of asynchronous testing in the setting without any big-step synchronization, and an improved understanding of the relations and possible interconnections between local (i.e. distributed) and asynchronous (centralized) testing. This has been the objective of the *TECSTES* project (2011-2014), funded by a DIGITEO *DIM/LSC* grant, and which involved Hernán Ponce de Léon and Stefan Haar of *MExICo*, and Delphine Longuet at LRI, University Paris-Sud/Orsay. We have extended several well known conformance (ioco style) relations for sequential models to models that can handle concurrency (labeled event structures). Two semantics (interleaving and partial order) were presented for every relation. With the interleaving semantics, the relations we obtained boil down to the same relations defined for labeled transition systems, since they focus on sequences of actions. The only advantage of using labeled event structures as a specification formalism for testing remains in the conciseness of the concurrent model with respect to a sequential one. As far as testing is concerned, the benefit is low since every interleaving has to be tested. By contrast, under the partial order semantics, the relations we obtain allow to distinguish explicitly implementations where concurrent actions are implemented concurrently, from those where they are interleaved, i.e. implemented sequentially. Therefore, these relations will be of interest when designing distributed systems, since the natural concurrency between actions that are performed in parallel by different processes can be taken into account. In particular, the fact of being unable to control or observe the order between actions taking place on different processes will not be considered as an impediment for testing. We have developed a complete testing framework for concurrent systems, which included the notions of test suites and test cases. We studied what kind of systems are testable in such a framework, and we have proposed sufficient conditions for obtaining a complete test suite as well as an algorithm to construct a test suite with such properties.

A mid-to long term goal (which may or may not be addressed by *MExICo* depending on the availability of staff for this subject) is the comprehensive formalization of testing and testability in asynchronous systems with distributed architecture and test protocols.

3.2. Interaction

Participants: Benedikt Bollig, Thomas Chatain, Paul Gastin, Stefan Haar, Serge Haddad.

3.2.1. Introduction

Systems and services exhibit non-trivial *interaction* between specialized and heterogeneous components. This interplay is challenging for several reasons. On one hand, a coordinated interplay of several components is required, though each has only a limited, partial view of the system's configuration. We refer to this problem as *distributed synthesis* or *distributed control*. An aggravating factor is that the structure of a component might be semi-transparent, which requires a form of *grey box management*.

Interaction, one of the main characteristics of systems under consideration, often involves an environment that is not under the control of cooperating services. To achieve a common goal, the services need to agree upon a strategy that allows them to react appropriately regardless of the interactions with the environment. Clearly, the notions of opponents and strategies fall within *game theory*, which is naturally one of our main tools in exploring interaction. We will apply to our problems techniques and results developed in the domains of distributed games and of games with partial information. We will consider also new problems on games that arise from our applications.

3.2.2. Distributed Control

Participants: Benedikt Bollig, Thomas Chatain, Paul Gastin, Stefan Haar.

Program synthesis, as introduced by Church [59] aims at deriving directly an implementation from a specification, allowing the implementation to be correct by design. When the implementation is already at hand but choices remain to be resolved at run time then the problem becomes controller synthesis. Both program and controller synthesis have been extensively studied for sequential systems. In a distributed setting, we need to synthesize a distributed program or distributed controllers that interact locally with the system components. The main difficulty comes from the fact that the local controllers/programs have only a partial view of the entire system. This is also an old problem largely considered undecidable in most settings [76], [71], [74], [65], [67].

Actually, the main undecidability sources come from the fact that this problem was addressed in a synchronous setting using global runs viewed as sequences. In a truly distributed system where interactions are asynchronous we have recently obtained encouraging decidability results [66], [56]. This is a clear witness where concurrency may be exploited to obtain positive results. It is essential to specify expected properties directly in terms of causality revealed by partial order models of executions (MSCs or Mazurkiewicz traces). We intend to develop this line of research with the ambitious aim to obtain decidability for all natural systems and specifications. More precisely, we will identify natural hypotheses both on the architecture of our distributed system and on the specifications under which the distributed program/controller synthesis problem is decidable. This should open the way to important applications, e.g., for distributed control of embedded systems.

3.2.3. Adaptation and Grey box management

Participants: Stefan Haar, Serge Haddad.

Contrary to mainframe systems or monolithic applications of the past, we are experiencing and using an increasing number of services that are performed not by one provider but rather by the interaction and cooperation of many specialized components. As these components come from different providers, one can no longer assume all of their internal technologies to be known (as it is the case with proprietary technology). Thus, in order to compose e.g. orchestrated services over the web, to determine violations of specifications or contracts, to adapt existing services to new situations etc, one needs to analyze the interaction behavior of *boxes* that are known only through their public interfaces. For their semi-transparent-semi-opaque nature,

we shall refer to them as **grey boxes**. While the concrete nature of these boxes can range from vehicles in a highway section to hotel reservation systems, the tasks of *grey box management* have universal features allowing for generalized approaches with formal methods. Two central issues emerge:

- Abstraction: From the designer point of view, there is a need for a trade-off between transparency (no abstraction) in order to integrate the box in different contexts and opacity (full abstraction) for security reasons.

- Adaptation: Since a grey box gives a partial view about the behavior of the component, even if it is not immediately useable in some context, the design of an adaptator is possible. Thus the goal is the synthesis of such an adaptator from a formal specification of the component and the environment.

Our work on direct modeling and handling of "grey boxes" via modal models (see [61]) was halted when Dorsaf El-Hog stopped her PhD work to leave academia, and has not resumed for lack of staff. However, it should be noted that semi-transparent system management in a larger sense remains an active field for the team, witness in particular our work on diagnosis and testing.

3.3. Management of Quantitative Behavior

Participants: Benedikt Bollig, Thomas Chatain, Paul Gastin, Stefan Haar, Serge Haddad.

3.3.1. Introduction

Besides the logical functionalities of programs, the *quantitative* aspects of component behavior and interaction play an increasingly important role.

- *Real-time* properties cannot be neglected even if time is not an explicit functional issue, since transmission delays, parallelism, etc, can lead to time-outs striking, and thus change even the logical course of processes. Again, this phenomenon arises in telecommunications and web services, but also in transport systems.

- In the same contexts, *probabilities* need to be taken into account, for many diverse reasons such as unpredictable functionalities, or because the outcome of a computation may be governed by race conditions.

- Last but not least, constraints on *cost* cannot be ignored, be it in terms of money or any other limited resource, such as memory space or available CPU time.

Traditional mainframe systems were proprietary and (essentially) localized; therefore, impact of delays, unforeseen failures, etc. could be considered under the control of the system manager. It was therefore natural, in verification and control of systems, to focus on *functional* behavior entirely.

With the increase in size of computing system and the growing degree of compositionality and distribution, quantitative factors enter the stage:

- calling remote services and transmitting data over the web creates *delays*;

- remote or non-proprietary components are not "deterministic", in the sense that their behavior is uncertain.

Time and *probability* are thus parameters that management of distributed systems must be able to handle; along with both, the *cost* of operations is often subject to restrictions, or its minimization is at least desired. The mathematical treatment of these features in distributed systems is an important challenge, which *MExICo* is addressing; the following describes our activities concerning probabilistic and timed systems. Note that cost optimization is not a current activity but enters the picture in several intended activities.

3.3.2. Probabilistic distributed Systems

Participants: Stefan Haar, Serge Haddad, Claudine Picaronny.

3.3.2.1. Non-sequential probabilistic processes

Practical fault diagnosis requires to select explanations of *maximal likelihood*. For partial-order based diagnosis, this leads therefore to the question what the probability of a given partially ordered execution is. In Benveniste et al. [50], [43], we presented a model of stochastic processes, whose trajectories are partially ordered, based on local branching in Petri net unfoldings; an alternative and complementary model based on Markov fields is developed in [69], which takes a different view on the semantics and overcomes the first model's restrictions on applicability.

Both approaches abstract away from real time progress and randomize choices in *logical* time. On the other hand, the relative speed - and thus, indirectly, the real-time behavior of the system's local processes - are crucial factors determining the outcome of probabilistic choices, even if non-determinism is absent from the system.

In another line of research [54] we have studied the likelihood of occurrence of non-sequential runs under random durations in a stochastic Petri net setting. It remains to better understand the properties of the probability measures thus obtained, to relate them with the models in logical time, and exploit them e.g. in *diagnosis*.

3.3.2.2. Distributed Markov Decision Processes
Participant: Serge Haddad.

Distributed systems featuring non-deterministic and probabilistic aspects are usually hard to analyze and, more specifically, to optimize. Furthermore, high complexity theoretical lower bounds have been established for models like partially observed Markovian decision processes and distributed partially observed Markovian decision processes. We believe that these negative results are consequences of the choice of the models rather than the intrinsic complexity of problems to be solved. Thus we plan to introduce new models in which the associated optimization problems can be solved in a more efficient way. More precisely, we start by studying connection protocols weighted by costs and we look for online and offline strategies for optimizing the mean cost to achieve the protocol. We have been cooperating on this subject with the SUMO team at Inria Rennes; in the joint work [44]; there, we strive to synthesize for a given MDP a control so as to guarantee a specific stationary behavior, rather than - as is usually done - so as to maximize some reward.

3.3.3. Large scale probabilistic systems

Addressing large-scale probabilistic systems requires to face state explosion, due to both the discrete part and the probabilistic part of the model. In order to deal with such systems, different approaches have been proposed:

- Restricting the synchronization between the components as in queuing networks allows to express the steady-state distribution of the model by an analytical formula called a product-form [49].

- Some methods that tackle with the combinatory explosion for discrete-event systems can be generalized to stochastic systems using an appropriate theory. For instance symmetry based methods have been generalized to stochastic systems with the help of aggregation theory [58].

- At last simulation, which works as soon as a stochastic operational semantic is defined, has been adapted to perform statistical model checking. Roughly speaking, it consists to produce a confidence interval for the probability that a random path fulfills a formula of some temporal logic [80] .

We want to contribute to these three axes: (1) we are looking for product-forms related to systems where synchronization are more involved (like in Petri nets), see [9]; (2) we want to adapt methods for discrete-event systems that require some theoretical developments in the stochastic framework and, (3) we plan to address some important limitations of statistical model checking like the expressiveness of the associated logic and the handling of rare events.

3.3.4. Real time distributed systems

Nowadays, software systems largely depend on complex timing constraints and usually consist of many interacting local components. Among them, railway crossings, traffic control units, mobile phones, computer servers, and many more safety-critical systems are subject to particular quality standards. It is therefore becoming increasingly important to look at networks of timed systems, which allow real-time systems to operate in a distributed manner.

Timed automata are a well-studied formalism to describe reactive systems that come with timing constraints. For modeling distributed real-time systems, networks of timed automata have been considered, where the local clocks of the processes usually evolve at the same rate [73] [55]. It is, however, not always adequate to assume that distributed components of a system obey a global time. Actually, there is generally no reason to assume that different timed systems in the networks refer to the same time or evolve at the same rate. Any component is rather determined by local influences such as temperature and workload.

3.3.4.1. Implementation of Real-Time Concurrent Systems
Participants: Thomas Chatain, Stefan Haar, Serge Haddad.

This was one of the tasks of the ANR ImpRo.

Formal models for real-time systems, like timed automata and time Petri nets, have been extensively studied and have proved their interest for the verification of real-time systems. On the other hand, the question of using these models as specifications for designing real-time systems raises some difficulties. One of those comes from the fact that the real-time constraints introduce some artifacts and because of them some syntactically correct models have a formal semantics that is clearly unrealistic. One famous situation is the case of Zeno executions, where the formal semantics allows the system to do infinitely many actions in finite time. But there are other problems, and some of them are related to the distributed nature of the system. These are the ones we address here.

One approach to implementability problems is to formalize either syntactical or behavioral requirements about what should be considered as a reasonable model, and reject other models. Another approach is to adapt the formal semantics such that only realistic behaviors are considered.

These techniques are preliminaries for dealing with the problem of implementability of models. Indeed implementing a model may be possible at the cost of some transformation, which make it suitable for the target device. By the way these transformations may be of interest for the designer who can now use high-level features in a model of a system or protocol, and rely on the transformation to make it implementable.

We aim at formalizing and automating translations that preserve both the timed semantics and the concurrent semantics. This effort is crucial for extending concurrency-oriented methods for logical time, in particular for exploiting partial order properties. In fact, validation and management - in a broad sense - of distributed systems is not realistic *in general* without understanding and control of their real-time dependent features; the link between real-time and logical-time behaviors is thus crucial for many aspects of *MExICo*'s work.

3.3.5. Weighted Automata and Weighted Logics
Participants: Benedikt Bollig, Paul Gastin.

Time and probability are only two facets of quantitative phenomena. A generic concept of adding weights to qualitative systems is provided by the theory of weighted automata [42]. They allow one to treat probabilistic or also reward models in a unified framework. Unlike finite automata, which are based on the Boolean semiring, weighted automata build on more general structures such as the natural or real numbers (equipped with the usual addition and multiplication) or the probabilistic semiring. Hence, a weighted automaton associates with any possible behavior a weight beyond the usual Boolean classification of "acceptance" or "non-acceptance". Automata with weights have produced a well-established theory and come, e.g., with a characterization in terms of rational expressions, which generalizes the famous theorem of Kleene in the unweighted setting. Equipped with a solid theoretical basis, weighted automata finally found their way into numerous application areas such as natural language processing and speech recognition, or digital image compression.

What is still missing in the theory of weighted automata are satisfactory connections with verification-related issues such as (temporal) logic and bisimulation that could lead to a general approach to corresponding satisfiability and model-checking problems. A first step towards a more satisfactory theory of weighted systems was done in [53]. That paper, however, does not give definite answers to all the aforementioned problems. It identifies directions for future research that we will be tackling.

4. Application Domains

4.1. Telecommunications

Participants: Stefan Haar, Serge Haddad.

MExICo's research is motivated by problems on system management in several domains:
- In the domain of service oriented computing, it is often necessary to insert some Web service into an existing orchestrated business process, e.g. to replace another component after failures. This requires to ensure, often actively, conformance to the interaction protocol. One therefore needs to synthesize *adaptators* for every component in order to steer its interaction with the surrounding processes.
- Still in the domain of telecommunications, the supervision of a network tends to move from out-of-band technology, with a fixed dedicated supervision infrastructure, to in-band supervision where the supervision process uses the supervised network itself. This new setting requires to revisit the existing supervision techniques using control and diagnosis tools.

We have participated in the Univerself Projecton self-aware networks, and will be searching new cooperations.

4.2. Transport Systems

Participants: Stefan Haar, Serge Haddad, Yann Duplouy, Simon Theissing.

We participate in the IRT *System X*'s *system of systems* program *TMM*, in two projects:
- project MIC on multi-modal transport systems with academic partners UPMC, IFSTTAR and CEA, and several industrial partners including Alstom (project leader), COSMO and Renault. Transportation operators in an urban area need to plan, supervise and steer different means of transportation with respect to several criteria:
  - Maximize capacity;
  - guarantee punctuality and robustness of service;
  - minimize energy consumption.

  The systems must achieve these objectives not only under ideal conditions, but also be robust to perturbations (such as a major cultural or sport event creating additional traffic), modifications of routes (roadwork, accidents, demonstrations, ...) and tolerant to technical failures. Therefore, systems must be enabled to raise appropriate alarms upon detection of anomalies, diagnose the type of anomaly and select the appropriate response. While the above challenges belong already to the tasks of individual operators in the unimodal setting, the rise of and increasing demand for *multi-modal* transports forces to achieve these planning, optimization and control goals not in isolation, but in a cooperative manner, across several operators. The research task here is first to analyze the transportation system regarding the available means, capacities and structures, and so as to identify the impacting factors and interdependencies of the system variables. Based on this analysis, the task is to derive and implement robust planning, with tolerance to technical faults; diagnosis and control strategies that are optimal under several, possibly different, criteria (average case vs worst case performance, energy efficiency, etc.) and allow to adapt to changes e.g. from nominal mode to reduced mode, sensor failures, etc.
- the project SVA (*Simulation pour la Sécurité du Véhicule Autonome*), where the PhD Thesis of Yann Duplouy targets the application of formal methods to the development of embedded systems for autonomous vehicles.

4.3. Biological Systems

Participants: Stefan Haar, Serge Haddad, Stefan Schwoon, Thomas Chatain.

We have begun in 2014 to examine concurrency issues in systems biology, and are currently enlarging the scope of our research's applications in this direction. To see the context, note that in recent years, a considerable shift of biologists' interest can be observed, from the mapping of *static* genotypes to *gene expression*, i.e. the processes in which genetic information is used in producing functional products. These processes are far from being uniquely determined by the gene itself, or even jointly with static properties of the environment; rather, *regulation* occurs throughout the expression processes, with specific mechanisms increasing or decreasing the production of various products, and thus modulating the outcome. These regulations are central in understanding cell fate (how does the cell differenciate ? Do mutations occur ? etc), and progress there hinges on our capacity to analyse, predict, monitor and control complex and variegated processes. Our first step in this domain is related in the conference contribution [57], where we apply Petri net unfolding techniques for the efficient computation of *attractors* in a regulatory network; that is, to identify strongly connected reachability components that correspond to stable evolutions, e.g. of a cell that differentiates into a specific functionality (or mutation). This constitutes the starting point of a broader research with Petri net unfolding techniques in regulation. In fact, ,he use of *ordinary* Petri nets for capturing regulatory network (RN) dynamics overcomes the limitations of traditional RN models : those impose e.g. Monotonicity properties in the influence that one factor had upon another, i.e. always increasing or always decreasing, and were thus unable to cover all actual behaviours (see [75]). Rather, we follow the more refined model of boolean networks of automata, where the local states of the different factors jointly detemine which state transitions are possible. For these connectors, ordinary PNs constitute a first approximation, improving greatly over the literature but leaving room for improvement in terms of introducing more refined logical connectors. Future work thus involves transcending this class of PN models. Via unfoldings, one has access – provided efficient techniques are available – to all behaviours of the model, rather than over-or under-approximations as previously. This opens the way to efficiently searching in particular for determinants of the cell fate : which attractors are reachable from a given stage, and what are the factors that decide in favor of one or the other attractor, etc. The list of potential applications in biology and medicine of such a methodology would be too long to reproduce here.

5. Highlights of the Year

5.1. Highlights of the Year

5.1.1. *Major Results*

5.1.1.1. *Approaching the Coverability Problem Continuously*

The coverability problem for Petri nets plays a central role in the verification of concurrent shared-memory programs. However, its high EXPSPACE-complete complexity poses a challenge when encountered in real-world instances. In [39], we develop a new approach to this problem which is primarily based on applying forward coverability in continuous Petri nets as a pruning criterion inside a backward coverability framework. A cornerstone of our approach is the efficient encoding of a recently developed polynomial-time algorithm for reachability in continuous Petri nets into SMT. We demonstrate the effectiveness of our approach on standard benchmarks from the literature, which shows that our approach decides significantly more instances than any existing tool and is in addition often much faster, in particular on large instances.

5.1.1.2. *An Automata-Theoretic Approach to the Verification of Distributed Algorithms*

In [21] we introduce an automata-theoretic method for the verification of distributed algorithms running on ring networks. In a distributed algorithm, an arbitrary number of processes cooperate to achieve a common goal (e.g., elect a leader). Processes have unique identifiers (pids) from an infinite, totally ordered domain. An algorithm proceeds in synchronous rounds, each round allowing a process to perform a bounded sequence of actions such as send or receive a pid, store it in some register, and compare register contents w.r.t. the associated

total order. An algorithm is supposed to be correct independently of the number of processes. To specify correctness properties, we introduce a logic that can reason about processes and pids. Referring to leader election, it may say that, at the end of an execution, each process stores the maximum pid in some dedicated register. Since the verification of distributed algorithms is undecidable, we propose an underapproximation technique, which bounds the number of rounds. This is an appealing approach, as the number of rounds needed by a distributed algorithm to conclude is often exponentially smaller than the number of processes. We provide an automata-theoretic solution, reducing model checking to emptiness for alternating two-way automata on words. Overall, we show that round-bounded verification of distributed algorithms over rings is PSPACE-complete.

5.1.1.3. Unfolding-Based Process Discovery

In [33] we presents a novel technique for process discovery. In contrast to the current trend, which only considers an event log for discovering a process model, we assume two additional inputs: an independence relation on the set of logged activities, and a collection of negative traces. After deriving an intermediate net unfolding from them, we perform a controlled folding giving rise to a Petri net which contains both the input log and all independence-equivalent traces arising from it. Remarkably, the derived Petri net cannot execute any trace from the negative collection. The entire chain of transformations is fully automated. A tool has been developed and experimental results are provided.

6. New Software and Platforms

6.1. General Remark

The team's software and platform are the same as in 2014, namely

- COSMOS,
- MOLE,
- CosyVerif;

no major changes have occurred in 2015.

7. New Results

7.1. Highlights

Please note that three of our most important and novel results are given in the 'Highlights' section above.

7.2. Specifying and Verifying Concurrent C Programs with TLA+

Verifying software systems automatically from their source code rather than modelling them in a dedicated language gives more confidence in establishing their properties. In [37] we propose a formal specification and verification approach for concurrent C programs directly based on the semantics of C. We define a set of translation rules and implement it in a tool (C2TLA+) that automatically translates C code into a TLA+ specification. The TLC model checker can use this specification to generate a model, allowing to check the absence of runtime errors and dead code in the C program in a given configuration. In addition, we show how translated specifications interact with manually written ones to: check the C code against safety or liveness properties; provide concurrency primitives or model hardware that cannot be expressed in C; and use abstract versions of translated C functions to address the state explosion problem. All these verifications have been conducted on an industrial case study, which is a part of the microkernel of the PharOS real-time system.

7.3. Active Diagnosis with Observable Quiescence

Active diagnosis of a discrete-event system consists in controlling the system such that faults can be detected. In [27] we extend the framework of active diagnosis presented in [7] by introducing modalities for actions and states and a new capability for the controller, namely observing that the system is quiescent. We design a game-based construction for both the decision and the synthesis problems that is computationally optimal. Furthermore we prove that the size and the delay provided by the active diagnoser (when it exists) are almost optimal.

7.4. Test Case Generation for Concurrent Systems Using Event Structures

In [23] we deal with the test-case generation problem for concurrent systems that are specified by true-concurrency models such as Petri nets. We show that using true-concurrency models reduces both the size and the number of test cases needed for achieving certain coverage criteria. We present a test-case generation algorithm based on Petri net unfoldings and a SAT encoding for solving controllability problems in test cases. Finally, we evaluate our algorithm against traditional test-case generation methods under interleaving semantics.

7.5. State Space Reduction Strategie for Model Checking Concurrent C Programs

Model checking is an effective technique for uncovering subtle errors in concurrent systems. Unfortunately, the state space explosion is the main bottleneck in model checking tools. In [31] we propose a state space reduction technique for model checking concurrent programs written in C. The reduction technique consists in an analysis phase, which defines an approximate agglomeration predicate. This latter states whether a statement can be agglomerated or not. We implement this predicate using a syntactic analysis, as well as a semantic analysis based on abstract interpretation. We show the usefulness of using agglomeration technique to reduce the state space, as well as to generate an abstract TLA+ specification from a C program.

7.6. Simple Priced Timed Games Are Not That Simple

Priced timed games are two-player zero-sum games played on priced timed automata (whose locations and transitions are labeled by weights modeling the costs of spending time in a state and executing an action, respectively). The goals of the players are to minimise and maximise the cost to reach a target location, respectively. In [25] we consider priced timed games with one clock and arbitrary (positive and negative) weights and show that, for an important subclass of theirs (the so-called simple priced timed games), one can compute, in exponential time, the optimal values that the players can achieve, with their associated optimal strategies. As side results, we also show that one-clock priced timed games are determined and that we can use our result on simple priced timed games to solve the more general class of so-called reset-acyclic priced timed games (with arbitrary weights and one-clock).

7.7. A Hybrid-Dynamical Model for Passenger-flow in Transportation Systems

In a network with different transportation modes, or multimodal public transportation system (MPTS), modes are linked among one another not by resources or infrastructure elements—which are not shared, e.g., between different metro lines—but by the flow of passengers between them. Now, the movements of passengers are steered by the destinations that individual passengers have, and by which they can be grouped into trip profiles. To use the strength of fluid dynamics, introduce in [30] a multiphase hybrid Petri net model, in which the vehicle dynamics is rendered by individual tokens moving in an infrastructure net, while passenger quantities are given as vectors—whose components correspond to trip profiles—and evolve at stations according to fluid dynamics. This model is intended as a building block for obtaining supervisory control, via transport operator actions, to mitigate congestion.

7.8. An Algebraic View of Space/Belief and Extrusion/Utterance for Concurrency/Epistemic Logic

In [29] we enrich spatial constraint systems with operators to specify information and processes moving from a space to another. We shall refer to these news structures as spatial constraint systems with extrusion. We shall investigate the properties of this new family of constraint systems and illustrate their applications. From a computational point of view the new operators provide for process/information extrusion, a central concept in formalisms for mobile communication. From an epistemic point of view extrusion corresponds to a notion we shall call utterance; a piece of information that an agent communicates to others but that may be inconsistent with the agent's beliefs. Utterances can then be used to express instances of epistemic notions, which are common place in social media, such as hoaxes or intentional lies. Spatial constraint systems with extrusion can be seen as complete Heyting algebras equipped with maps to account for spatial and epistemic specifications.

7.9. Preserving Partial Order Runs in Parametric Time Petri Nets

Parameter synthesis for timed systems aims at deriving parameter valuations satisfying a given property. In [22] we target concurrent systems; it is well known that concurrency is a source of state-space explosion, and partial order techniques were defined to cope with this problem. Here we use partial order semantics for parametric time Petri nets as a way to significantly enhance the result of an existing synthesis algorithm. Given a reference parameter valuation, our approach synthesizes other valuations preserving, up to interleaving, the behavior of the reference parameter valuation. We show the applicability of our approach using acyclic asynchronous circuits.

7.10. Non-Atomic Transition Firing in Contextual Nets

The firing rule for Petri nets assumes instantaneous and simultaneous consumption and creation of tokens. In the context of ordinary Petri nets, this poses no particular problem because of the system's asynchronicity, even if token creation occurs later than token consumption in the firing. With read arcs, the situation changes, and several different choices of semantics are possible. The step semantics introduced by Janicki and Koutny can be seen as imposing a two-phase firing scheme: first, the presence of the required tokens is checked, then consumption and production of tokens happens. Pursuing this approach further, we develop in [28] a more general framework based on explicitly splitting the phases of firing, allowing to synthesize coherent steps. This turns out to define a more general non-atomic semantics, which has important potential for safety as it allows to detect errors that were missed by the previous semantics. Then we study the characterization of partial-order processes feasible under one or the other semantics.

8. Bilateral Contracts and Grants with Industry

8.1. Bilateral Contracts with Industry

At present, our industrial cooperations are centered in the IRT SystemX, see below; there are currently no *bilateral* agreements.

9. Partnerships and Cooperations

9.1. Regional Initiatives: IRT

9.1.1. *SystemX*

Participants: Simon Theissing, Yann Duplouy, Serge Haddad.

We participate in the projects

- MIC on multi-modal transport systems with in the IRT *System X*, with academic partners UPMC, IFSTTAR and CEA, and several industrial partners including Alstom (project leader), COSMO and Renault. MIC is scheduled to be completed late in 2016, and

- the project SVA (*Simulation pour la Sécurité du Véhicule Autonome*), where the PhD Thesis of Yann Duplouy targets the application of formal methods to the development of embedded systems for autonomous vehicles.

9.2. National Initiatives

We have not yet been notified about acceptance of our ANR submissions.

9.3. European Initiatives

In preparation.

9.3.1. FP7 & H2020 Projects

Serge Haddad is participating in the ERC *EQualIS*, 'Enhancing the Quality of Interacting Systems', directed by Patricia Bouyer.

9.4. International Initiatives

9.4.1. Inria International Labs projects

LIA INFORMEL with CMI, Chennai, India ; see below.

9.4.2. Inria International Partners

9.4.2.1. Informal International Partners

1. The CMI (Chennai Mathematical Institute) is a long-standing partner of our team. The project *Île de France/Inde* in the *ARCUS* program from 2008 to 2011 has allowed several exchange visits between Cachan and Chennai, organizations of ACTS workshops with french and indian researchers in Chennai, internships in Cachan, and two theses in *co-tutelle* (Akshay Sundararaman, defended in 2010) and Aiswarya Cyriac (defended in 2014).

   Currently, Paul Gastin is co-head (with Madhavan Mukund) of the CNRS International Associated Laboratory (LIA) INFORMEL (INdo-French FORmal Methods Lab, http://projects.lsv.ens-cachan.fr/informel/), see below.

2. We have been exchanging visits for several years between *MExICo* the computer science and electrical engineering departments at Newcastle University, UK , with visits in both directions; they involve in particular Maciej Koutny, Alex Yakovlev, Victor Khomenko and Andrey Mokhov, as well as Anil Wipat, co-director of the center for Synthetic Biology and the Bioeconomy at Newcastle University.

3. Exchanges are frequent with Rolf Hennicker from LMU and Javier Esparza at TUM, both in Munich, Germany.

9.5. International Research Visitors

9.5.1. Visits of International Scientists

- 5 – 31 March 2015: Prakash Saivasan (CMI) visits LSV to work with Paul Gastin on nested words for higher-order pushdown systems

- 19 May – 6 June 2015: S. Krishna and S. Akshay visit LSV to work with Paul Gastin on split-width techniques for the analysis of timed systems.

- 10 June – 4 July 2015: K. Narayan Kumar (CMI) visit France to pursue several collaborations: with Paul Gastin (LSV) on bounded time-stamping for message passing systems, with Ahmed Bouajjani (LIAFA) on analysis of multi-pushdown systems, and with Pascal Weil (LaBRI) on bounded reachability analysis for shared memory systems.

9.5.1.1. Internships

Georgios Christodoulis

Date: May 2015 - Jul 2015

Institution: National University Athens (Greece)

Supervisor: Stefan Haar

Sougata Bose

Date: May 2015 - Jul 2015

Institution: CMI (India)

Supervisor: Benedikt Bollig and Paul Gastin

9.5.2. Visits to International Teams

9.5.2.1. Short stays abroad

- In July 2015, Serge Haddad visited U of Turin, Italy, for a research cooperation with Prof. Giuliana Franceschinis.
- Stefan Haar visited Newcastle University (UK), TU of Eindhoven (NL) and University of Luxemburg for short visits.
- 29 November – 20 December 2015: Paul Gastin (LSV) visits S. Krishna and S. Akshay (IIT Bombay) to work on tree automata techniques for timed-systems.

10. Dissemination

10.1. Promoting Scientific Activities

10.1.1. LIA INFORMEL

The Indo-French Formal Methods Lab is an International Associated Laboratory (LIA) fostering the scientific collaboration between India and France in the domain of formal methods and applications to the verification of complex systems. Our research focuses on theoretical foundations of games, automata, and logics, three important tools in formal methods. We study applications to the verification of safety-critical systems, with an emphasis on quantitative aspects (time, cost, energy, etc.), concurrency, control, and security protocols. The Laboratory was founded in 2012 by a consortium of researchers from the French Centre for Scientific Research (CNRS), Ecole Normale Supérieure de Cachan (ENS Cachan), Université Bordeaux 1, the Institute of Mathematical Sciences Chennai (IMSc), the Chennai Mathematical Institute (CMI), and the Indian Institute of Science Bangalore (IISc). It is directed by Paul Gastin (ENS Cachan, MExICo team) and Madhavan Mukund (CMI). The LIA has been scientifically extremely active and productive since its creation. The LIA has supported numerous scientific exchanges and joint research papers, see http://projects.lsv.ens-cachan.fr/informel/. Among many other activities, the LIA organised another edition of the ACTS workshop, see here.

The renewal beyond 2015 has been approved by CNRS' national committee, the FSD's approval is pending.

10.1.2. Scientific events organisation

10.1.2.1. General chair, scientific chair

Serge Haddad is a member of the steering committee for ICATPN.

10.1.3. *Scientific events selection*

10.1.3.1. Chair of conference program committees

- Stefan Haar was co-chair of the program committee of the conference Application of Concurrency to System Design (ACSD) 2015 [38].

10.1.3.2. Member of the conference program committees

- Thomas Chatain was a member of the program committee of the conference Application of Concurrency to System Design (ACSD) 2015 [38].
- Paul Gastin was a member of the program committee of the workshop ACTS 2015.
- Stefan Haar was a member of the program committees of the ETFA 2015 workshop, the CIIA 2015 workshop, and of the SAFEPROCESS 2015 Symposium and the ICTAC 2015 conference.
- Stefan Schwoon was a member of the PCs for ICATPN 2015 and SPIN 2015.
- Serge Haddad was a member of the PCs of the workshops
    - FOR-MOVES, associated with ICSOC 2015, Goa, India;
    - VECOS 2015, Bejaia, Algeria;
    - ADECS and PNSE, both associated with ICATPN 2015, Brussels, Belgium.

10.1.3.3. Reviewer

- Stefan Haar was a reviewer for CSL 2015, ACC 2015
- Stefan Schwoon was a reviewer for the 2015 editions of SODA, TAMC, POPL, and TACAS.
- Paul Gastin,Benedikt Bollig and Serge Haddad are regularly reviewers for many international conferences.

10.1.4. *Journal*

10.1.4.1. Member of the editorial boards

- Paul Gastin is on the advisory boards of *Journal of Automata, Languages and Combinatorics* and of the EATCS Springer Book series *Monographs in Theoretical Computer Science* and *Texts in Theoretical Computer Science*.
- Serge Haddad was Editor of one edition of the TOPNOC journal (LNCS 8910).
- Stefan Haar is an associate editor for the *Journal of Discrete Event Dynamic systems*.

10.1.4.2. Reviewer - Reviewing activities

- Stefan Haar was a reviewer for *Automatica*, *Journal of Computer and System Sciences*, *IEEE Transactions on Automatic Control*, *Computer Journal* and *TOPNOC*. He is also a regular reviewer for *AMS Mathematical Reviews*.
- Thomas Chatain was a reviewer for reviewer for*ACM Transactions on Embedded Computing Systems*.
- Stefan Schwoon acted as reviewer for the journals TCS, Fundamenta, STTT, FAOC (Formal Aspects of Computing), and ToPNoC.
- Paul Gastin, Benedikt Bollig and Serge Haddad are regularly reviewers for many international journals.

10.1.5. *Invited talks*

- Serge Haddad gave invited talks
    - on 'Active Diagnosis', July 7, 2015, Turin University;
    - on 'Time and Stochastic Petri Nets', June 22, during the Petri Nets tutorials at ICATPN 2015, Brussels, Belgium;
    - on 'Polynomial Interrupt Timed Automata' at Inria Rennes, Sept. 17,

- on 'probabilistic automata' at *Ecole Jeunes Chercheurs EJCIM 2015*, April 1st, 2015, Orléans, France.

- Paul Gastin was an invited speaker for

  - 'Formal methods for the verification of distributed algorithms' Invited talk at INFINITY 2015, Bangalore (India), December 15, 2015.

  - 'Weighted Automata: Highlighted Excerpts.' Tutorial at HIGHLIGHTS'15, Prague, September 15, 2015.

  - Gossip: Maintening latest information beyond channel bounds.
 Invited talk at ALFA'15, Bordeaux (France), June 15 - 17, 2015, and talk at FRIDA'15, Grenoble (France,) June 5, 2015.

  - '10 years of weighted logics for weighted automata.' Invited talk at AutoMathA, Leipzig (Germany), May 6 - 9, 2015.

- Benedikt Bollig was an invited speaker

  - at the 22nd International Symposium on Temporal Representation and Reasoning (TIME 2015), Kassel, Germany

  - at the 20th International Conference on Implementation and Application of Automata (CIAA 2015), Umeå, Sweden

  - at the 2nd Workshop on Parameterized Verification, Madrid, Spain

  - at ALFA 15: Automata, Logic, Formal languages, Algebra, Bordeaux, France

  - at Automata, Concurrency and Timed Systems (ACTS 2015), Chennai, India

- Stefan Schwoon gave a talk in the CASSTING workshop in Brussels on May 18th, 2015.

10.1.6. Scientific expertise

- Serge Haddad was member of a selection committee for a professor's position at U. Paris 6, and an external expert for the evaluation of ONERA's DCPS department in November 2015.

10.1.7. Research administration

- Paul Gastin is director of the LIA INFORMEL.

- Stefan Haar is the Saclay Inria center's correspondent for european partnerships. He also was a member of DIGITEO's project selection committee, and will be directing the *SCILEX* axis of the *LABEX DIGICOSME*.

All members of the team reviewed numerous papers for numerous international conferences.

10.2. Teaching - Supervision - Juries

10.2.1. Teaching

Note: We only list here the teaching activities of researchers, not the courses of full-time teachers in the team.

Licence : Stefan Haar, *Théorie des Langages*, 22.3 h (EQTD) L2, Université Paris-René Descartes, France

Master : Benedikt Bollig, *Non-Sequential Theory of Distributed Systems*, 22.5 h (EQTD), M2, Parisian Master of Research in Computer Science (MPRI), France

We further note that as of September 2015, Serge Haddad is no longer the coordinator of the M1 level but has become the director of the computer science department of ENS Cachan, as the successor of Paul Gastin in this position; and Stefan Schwoon has become the coordinator of the L3 level.

10.2.2. Supervision

- PhD in progress :
  - Engel Lefaucheux, *controlling information in probabilistic systems*, Co-supervisors Serge Haddad and Nathalie Bertrand (Inria Rennes, SUMO team);
  - Yann Duplouy, *application of formal methods to the development of embedded systems for autonomous vehicles*, since Sept. 2015, Supervisor Serge Haddad (with Béatrice Berard, UPMC).
  - Simon Theissing, *Supervision for Multi-Modal Transport Systems*, since September 2013, Supervisor Stefan Haar
  - Salim Perchy (Ecole Polytechnique), *D-spaces*, since November 2013, co-supervisor Stefan Haar, Supervisor Franck Valencia (Note : S. Perchy belongs to the COMETE team, not MExICo).

10.2.3. Juries

- Benedikt Bollig was a member of the PhD committee of Mathieu Caralp, "Problèmes de bornes pour les automates et les transducteurs à pile visible", December 18, 2015 at Universié Marseille.
- Paul Gastin was
  - President of the PhD committee for Jad Hamza at University Paris-Diderot, on November 27, 2015;
  - Reviewer for the theses of Roy Mennicke, Universität Ilmenau (Germany), and of Vitaly Perevoshchikov, Universität Leipzig (Germany);
  - a member of the HdR committees of Nathalie Bertrand, Rennes, November 16 and of Benedikt Bollig, ENS Cachan, june 2, 2015;
  - reviewer of the HDR of Olivier Serre, University Paris-Diderot, March 10, 2015
- Serge Haddad was a member of the PhD committees of M. Amziani at Telecom SudParis in june 2015, and of D. Paquereau at INSA Lyon in March 2015.
- On May 27, 2015, *Thomas Chatain* was a member of the jury in the competition for recruitment of Master's level students at ENS Rennes.
- *Stefan Haar* was a reviewer of the thesis by Kari Kähkönen on *Automated Systematic Testing Methods for Multithreaded Programs*, defended on February 2, 2015, at Aalto University, Finland.
- *Stefan Haar* was a member of the PhD examination board of the Thesis of Houssame-Eddine Gougam on *Analyse de l'impact du temps sur la diagnosticabilité des systèmes à événements discrets*, defended on Sep 28, INSA Toulouse.

11. Bibliography

Major publications by the team in recent years

[1] S. BALAGUER, TH. CHATAIN. *Avoiding Shared Clocks in Networks of Timed Automata*, in "Logical Methods in Computer Science", November 2013, vol. 9, n[o] 4:13, http://www.lsv.ens-cachan.fr/Publis/PAPERS/PDF/BC-lmcs13.pdf

[2] P. BALDAN, A. BRUNI, A. CORRADINI, B. KÖNIG, C. RODRÍGUEZ, S. SCHWOON. *Efficient unfolding of contextual Petri nets*, in "Theoretical Computer Science", August 2012, vol. 449, n[o] 1, pp. 2-22, http://www.lsv.ens-cachan.fr/Publis/PAPERS/PDF/bbckrs-tcs12.pdf

[3] B. BARBOT, S. HADDAD, C. PICARONNY. *Coupling and Importance Sampling for Statistical Model Checking*, in "Proceedings of the 18th International Conference on Tools and Algorithms for Construction and Analysis of Systems (TACAS'12)", Tallinn, Estonia, C. FLANAGAN, B. KÖNIG (editors), Lecture Notes in Computer Science, Springer, March 2012, vol. 7214, pp. 331-346, http://www.lsv.ens-cachan.fr/Publis/PAPERS/PDF/BHP-tacas12.pdf

[4] B. BOLLIG. *Logic for Communicating Automata with Parameterized Topology*, in "Proceedings of the Joint Meeting of the 23rd EACSL Annual Conference on Computer Science Logic and the 29th Annual ACM/IEEE Symposium on Logic In Computer Science (CSL/LICS'14)", Vienna, Austria, ACM Press, July 2014, http://hal.inria.fr/hal-00872807/

[5] B. BOLLIG, P. GASTIN, B. MONMEGE, M. ZEITOUN. *Pebble Weighted Automata and Weighted Logics*, in "ACM Transactions on Computational Logic", April 2014, vol. 15, n⁰ 2:15, http://www.lsv.ens-cachan.fr/Publis/PAPERS/PDF/BGMZ-tocl13.pdf

[6] A. CYRIAC, P. GASTIN, K. NARAYAN KUMAR. *MSO Decidability of Multi-Pushdown Systems via Split-Width*, in "Proceedings of the 23rd International Conference on Concurrency Theory (CONCUR'12)", Newcastle, UK, M. KOUTNY, I. ULIDOWSKI (editors), Lecture Notes in Computer Science, Springer, September 2012, vol. 7454, pp. 547-561, http://www.lsv.ens-cachan.fr/Publis/PAPERS/PDF/CGN-concur12.pdf

[7] S. HAAR, S. HADDAD, T. MELLITI, S. SCHWOON. *Optimal Constructions for Active Diagnosis*, in "Proceedings of the 33rd Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS'13)", Guwahati, India, A. SETH, N. VISHNOI (editors), Leibniz International Proceedings in Informatics, Leibniz-Zentrum für Informatik, December 2013, vol. 24, pp. 527-539, http://www.lsv.ens-cachan.fr/Publis/PAPERS/PDF/HHMS13-fsttcs.pdf

[8] S. HAAR, C. KERN, S. SCHWOON. *Computing the Reveals Relation in Occurrence Nets*, in "Theoretical Computer Science", July 2013, vol. 493, pp. 66-79, http://www.lsv.ens-cachan.fr/Publis/PAPERS/PDF/HKS-tcs13.pdf

[9] S. HADDAD, J. MAIRESSE, H.-T. NGUYEN. *Synthesis and Analysis of Product-form Petri Nets*, in "Fundamenta Informaticae", 2013, vol. 122, n⁰ 1-2, pp. 147-172, https://hal.archives-ouvertes.fr/hal-00925774

[10] H. PONCE DE LEÓN, S. HAAR, D. LONGUET. *Model-Based Testing for Concurrent Systems with Labeled Event Structures*, in "Software Testing, Verification and Reliability", November 2014, vol. 24, n⁰ 7, pp. 558-590, http://www.lsv.ens-cachan.fr/Publis/PAPERS/PDF/PHL-stvr14.pdf

Publications of the year

Doctoral Dissertations and Habilitation Theses

[11] B. BOLLIG. *Automata and Logics for Concurrent Systems: Realizability and Verification*, ENS Cachan, June 2015, Habilitation à diriger des recherches, https://hal.archives-ouvertes.fr/tel-01242861

Articles in International Peer-Reviewed Journals

[12] S. AKSHAY, P. GASTIN, M. MUKUND, K. NARAYAN KUMAR. *Checking conformance for time-constrained scenario-based specifications*, in "Journal of Theoretical Computer Science (TCS)", 2015, vol. 549, pp. 24-43, https://hal.archives-ouvertes.fr/hal-01242876

[13] P. BALLARINI, B. BARBOT, M. DUFLOT, S. HADDAD, N. PEKERGIN. *HASL: A new approach for performance evaluation and model checking from concepts to experimentation*, in "Performance Evaluation", August 2015, vol. 90, pp. 53-77 [*DOI :* 10.1016/J.PEVA.2015.04.003], https://hal.inria.fr/hal-01221815

[14] B. BARBOT, S. HADDAD, M. HEINER, C. PICARONNY. *Rare Event Handling in Signalling Cascades*, in "International Journal On Advances in Systems and Measurements", 2015, vol. 8, n^o 1-2, pp. 69-79, https://hal.archives-ouvertes.fr/hal-01242981

[15] E. FRACA, S. HADDAD. *Complexity Analysis of Continuous Petri Nets*, in "Fundamenta Informaticae", 2015, vol. 137, n^o 1, pp. 1-28 [*DOI :* 10.3233/FI-2015-1168], https://hal.archives-ouvertes.fr/hal-01241119

[16] V. GERMANOS, S. HAAR, V. KHOMENKO, S. SCHWOON. *Diagnosability under Weak Fairness*, in "ACM Transactions in Embedded Computing Systems", 2015, vol. 14, n^o 4, https://hal.archives-ouvertes.fr/hal-01242989

[17] S. HAAR. *Cyclic Ordering through Partial Orders*, in "Journal of Multiple-Valued Logic and Soft Computing", 2015, forthcoming, https://hal.archives-ouvertes.fr/hal-01242990

[18] H. PONCE DE LEÓN, S. HAAR, D. LONGUET. *Model-based Testing for Concurrent Systems: Unfolding-based Test Selection*, in "International Journal on Software Tools for Technology Transfer", 2015, forthcoming, https://hal.archives-ouvertes.fr/hal-01242994

Invited Conferences

[19] B. BOLLIG. *Automata and Logics for Concurrent Systems: Five Models in Five Pages*, in "20th International Conference on Implementation and Application of Automata (CIAA'15)", Umeå, Sweden, SPRINGER (editor), 2015, vol. 9223, pp. 3-12 [*DOI :* 10.1007/978-3-319-22360-5_1], https://hal.archives-ouvertes.fr/hal-01241125

[20] B. BOLLIG. *Towards Formal Verification of Distributed Algorithms*, in "22nd International Symposium on Temporal Representation and Reasoning", Kassel, Germany, IEEE Computer Society Press, 2015, https://hal.archives-ouvertes.fr/hal-01242959

International Conferences with Proceedings

[21] C. AISWARYA, B. BOLLIG, P. GASTIN. *An Automata-Theoretic Approach to the Verification of Distributed Algorithms*, in "26th International Conference on Concurrency Theory (CONCUR'15)", Madrid, Spain, Leibniz-Zentrum für Informatik, 2015, vol. 42, pp. 340-353, https://hal.archives-ouvertes.fr/hal-01241122

[22] E. ANDRÉ, T. CHATAIN, C. RODRIGUEZ. *Preserving Partial Order Runs in Parametric Time Petri Nets*, in "15th International Conference on Application of Concurrency to System Design (ACSD'15)", Brussels, Belgium, IEEE Computer Society Press, 2015, https://hal.archives-ouvertes.fr/hal-01242847

[23] K. ATHANASIOU, H. PONCE DE LEÓN, S. SCHWOON. *Test Case Generation for Concurrent Systems Using Event Structures*, in "9th International Conference on Tests and Proofs (TAP'15)", L'Aquila, Italy, Springer, 2015, vol. 9154, https://hal.archives-ouvertes.fr/hal-01242923

[24] N. BERTRAND, S. HADDAD, E. LEFAUCHEUX. *Accurate approximate diagnosability of stochastic systems*, in "10th International Conference on Language and Automata Theory and Applications", Prague, Czech Republic, Springer, March 2016, https://hal.inria.fr/hal-01220954

[25] T. BRIHAYE, G. GEERAERTS, A. HADDAD, E. LEFAUCHEUX, B. MONMEGE. *Simple Priced Timed Games Are Not That Simple*, in "35th Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS'15)", Banaglore, India, Leibniz-Zentrum für Informatik, 2015, https://hal.archives-ouvertes.fr/hal-01242902

[26] B. BÉRARD, S. HADDAD, C. PICARONNY, M. SAFEY EL DIN, M. SASSOLAS. *Polynomial Interrupt Timed Automata*, in "The 9th Workshop on Reachability Problems (RP'15)", Warsaw, Poland, Lecture Notes in Computer Science, Springer, September 2015, vol. 9328, pp. 20-32 [*DOI : 10.1007/978-3-319-24537-9_3*], https://hal.archives-ouvertes.fr/hal-01222572

[27] S. BÖHM, S. HAAR, S. HADDAD, P. HOFMAN, S. SCHWOON. *Active Diagnosis with Observable Quiescence*, in "54th IEEE Conference on Decision and Control (CDC'15)", Osaka, Japan, IEEE Control System Society, 2015, https://hal.archives-ouvertes.fr/hal-01242934

[28] T. CHATAIN, S. HAAR, M. KOUTNY, S. SCHWOON. *Non-Atomic Transition Firing in Contextual Nets*, in "36th International Conference on Applications and Theory of Petri Nets (ICATPN'15)", Brussels, Belgium, SPRINGER (editor), 2015, vol. 9115, https://hal.archives-ouvertes.fr/hal-01241129

[29] S. HAAR, S. PERCHY, C. RUEDA, F. VALENCIA. *An Algebraic View of Space/Belief and Extrusion/Utterance for Concurrency/Epistemic Logic*, in "17th International Conference on Principles and Practice of Declarative Programming (PPDP'15)", Siena, Italy, ACM Press, 2015, pp. 161-172, https://hal.archives-ouvertes.fr/hal-01242971

[30] S. HAAR, S. THEISSING. *A Hybrid-Dynamical Model for Passenger-flow in Transportation Systems*, in "5th IFAC Conference on Analysis and Design of Hybrid Systems (ADHS'15)", Atlanta, Georgia, United States, 2015, https://hal.archives-ouvertes.fr/hal-01242903

[31] A. METHNI, M. LEMERRE, B. BEN HEDIA, S. HADDAD, K. BARKAOUI. *State Space Reduction Strategie for Model Checking Concurrent C Programs*, in "9th Workshop on Verification and Evaluation of Computer and Communication Systems (VECoS'15)", Bucharest, Romania, 2015, vol. 1431, pp. 65-76, https://hal.archives-ouvertes.fr/hal-01242951

[32] H. PONCE DE LEON, A. MOKHOV. *Building Bridges Between Sets of Partial Orders*, in "International Conference on Language and Automata Theory and Applications", Nice, France, March 2015, https://hal.inria.fr/hal-01060449

[33] H. PONCE DE LEÓN, C. RODRIGUEZ, J. CARMONA, K. HELJANKO, S. HAAR. *Unfolding-Based Process Discovery*, in "13th International Symposium on Automated Technology for Verification and Analysis (ATVA'15)", Shanghai, China, Springer, 2015, vol. 9364, https://hal.archives-ouvertes.fr/hal-01242885

Conferences without Proceedings

[34] S. HAAR, S. PERCHY, C. RUEDA, F. VALENCIA. *An Algebraic View of Space/Belief and Extrusion/Utterance for Concurrency/Epistemic Logic*, in "17th International Symposium on Principles and Practice of Declarative Programming (PPDP 2015)", Siena, Italy, E. ALBERT, M. FALASCHI (editors), ACM SIGPLAN, July 2015, pp. 161-172 [*DOI : 10.1145/2790449.2790520*], https://hal.inria.fr/hal-01256984

Scientific Books (or Scientific Book chapters)

[35] N. BERTRAND, S. HADDAD. *Contrôle, probabilités et observation partielle*, in "Informatique Mathématique. Une photographie en 2015", CNRS Édition, 2015, pp. 177-227, https://hal.archives-ouvertes.fr/hal-01242962

[36] M. KOUTNY, J. DESEL, S. HADDAD. *Transactions on Petri Nets and Other Models of Concurrency X*, Springer, 2015, vol. 9410, https://hal.archives-ouvertes.fr/hal-01242977

[37] A. METHNI, M. LEMERRE, B. BEN HEDIA, S. HADDAD, K. BARKAOUI. *Specifying and Verifying Concurrent C Programs with TLA+*, in "Formal Techniques for Safety-Critical Systems", Communications in Computer and Information Science, Springer, 2015, vol. 476, pp. 206-222, https://hal.archives-ouvertes.fr/hal-01242954

Books or Proceedings Editing

[38] S. HAAR, R. MEYER (editors). *Proceedings 15th International Conference on Application of Concurrency to System Design (ACSD), 2015* , IEEE, Brussels, Belgium, 2015 [*DOI :* 10.1109/ACSD.2015.23], https://hal.inria.fr/hal-01250369

Research Reports

[39] M. BLONDIN, A. FINKEL, C. HAASE, S. HADDAD. *Approaching the Coverability Problem Continuously*, Computing Research Repository, 2015, n^o 1510.05724, https://hal.archives-ouvertes.fr/hal-01242916

Other Publications

[40] M. GUZMAN, S. HAAR, S. PERCHY, C. RUEDA, F. VALENCIA. *Space and Extrusion in Epistemic Concurrent Constraint Systems*, January 2016, Submitted to the Journal of Logical and Algebraic Methods in Programming, https://hal.inria.fr/hal-01257113

[41] S. HAAR, S. THEISSING. *Forecasting Passenger Loads in Transportation Networks*, January 2016, working paper or preprint, https://hal.inria.fr/hal-01259585

References in notes

[42] W. KUICH, H. VOGLER, M. DROSTE (editors). *Handbook of Weighted Automata*, EATCS Monographs in Theoretical Computer Science, Springer, 2009

[43] S. ABBES, A. BENVENISTE, S. HAAR. *A Petri net model for distributed estimation*, in "Proc. MTNS 2004, Sixteenth International Symposium on Mathematical Theory of Networks and Systems, Louvain (Belgium), ISBN 90-5682-517-8", 2004

[44] S. AKSHAY, N. BERTRAND, S. HADDAD, L. HELOUET. *The steady-state control problem for Markov decision processes*, in "Qest 2013", Buenos Aires, Argentina, K. R. JOSHI, M. SIEGLE, M. STOELINGA, P. R. D'ARGENIO (editors), Springer, September 2013, vol. 8054, pp. 290-304, https://hal.inria.fr/hal-00879355

[45] R. ALUR, K. ETESSAMI, M. YANNAKAKIS. *Realizability and Verification of MSC Graphs*, in "Theor. Comput. Sci.", 2005, vol. 331, n^o 1, pp. 97–114

[46] P. BALDAN, TH. CHATAIN, S. HAAR, B. KÖNIG. *Unfolding-based Diagnosis of Systems with an Evolving Topology*, in "Information and Computation", October 2010, vol. 208, n^o 10, pp. 1169-1192, http://www.lsv.ens-cachan.fr/Publis/PAPERS/PDF/BCHK-icomp10.pdf

[47] P. BALDAN, A. CORRADINI, B. KÖNIG, S. SCHWOON. *McMillan's complete prefix for contextual nets*, in "Transactions on Petri Nets and Other Models of Concurrency", November 2008, vol. 1, pp. 199–220, Volume 5100 of Lecture Notes in Computer Science

[48] P. BALDAN, S. HAAR, B. KOENIG. *Distributed Unfolding of Petri Nets*, in "Proc.FOSSACS 2006", LNCS, Springer, 2006, vol. 3921, pp. 126-141, Extended version: Technical Report CS-2006-1. Department of Computer Science, University Ca' Foscari of Venice

[49] F. BASKETT, K. M. CHANDY, R. R. MUNTZ, F. G. PALACIOS. *Open, Closed, and Mixed Networks of Queues with Different Classes of Customers*, in "J. ACM", April 1975, vol. 22, pp. 248–260, http://doi.acm.org/10.1145/321879.321887

[50] A. BENVENISTE, É. FABRE, S. HAAR. *Markov Nets: Probabilistic Models for distributed and concurrent Systems*, in "IEEE Transactions on Automatic Control", 2003, vol. 48 (11), pp. 1936-1950, Extended version: IRISA Research Report 1538

[51] P. BHATEJA, P. GASTIN, M. MUKUND, K. NARAYAN KUMAR. *Local testing of message sequence charts is difficult*, in "Proceedings of the 16th International Symposium on Fundamentals of Computation Theory (FCT'07)", Budapest, Hungary, E. CSUHAJ-VARJÚ, Z. ÉSIK (editors), Lecture Notes in Computer Science, Springer, August 2007, vol. 4639, pp. 76-87 [*DOI :* 10.1007/978-3-540-74240-1_8], http://www.lsv.ens-cachan.fr/Publis/PAPERS/PDF/BGMN-fct07.pdf

[52] G. V. BOCHMANN, S. HAAR, C. JARD, G.-V. JOURDAN. *Testing Systems Specified as Partial Order Input/Output Automata*, in "Proc. TESTCOM/Fates 08, 20th IFIP International Conference on Testing of Communicating Systems and 8th International Workshop on Formal Approaches to Testing of Software", LNCS, Springer, 2008, vol. 5047, pp. 169-183

[53] B. BOLLIG, P. GASTIN. *Weighted versus Probabilistic Logics*, in "Proceedings of the 13th International Conference on Developments in Language Theory (DLT'09)", Stuttgart, Germany, V. DIEKERT, D. NOWOTKA (editors), Lecture Notes in Computer Science, Springer, June-July 2009, vol. 5583, pp. 18-38 [*DOI :* 10.1007/978-3-642-02737-6_2], http://www.lsv.ens-cachan.fr/Publis/PAPERS/PDF/BG-dlt09.pdf

[54] A. BOUILLARD, S. HAAR, S. ROSARIO. *Critical paths in the Partial Order Unfolding of a Stochastic Petri Net*, in "Proceedings of the 7th International Conference on Formal Modelling and Analysis of Timed Systems (FORMATS'09)", Budapest, Hungary, J. OUAKNINE, F. VAANDRAGER (editors), Lecture Notes in Computer Science, Springer, September 2009, vol. 5813, pp. 43-57 [*DOI :* 10.1007/978-3-642-04368-0_6], http://www.lsv.ens-cachan.fr/Publis/PAPERS/PDF/BHR-formats09.pdf

[55] P. BOUYER, S. HADDAD, P.-A. REYNIER. *Timed Unfoldings for Networks of Timed Automata*, in "Proceedings of the 4th International Symposium on Automated Technology for Verification and Analysis (ATVA'06)", Beijing, ROC, S. GRAF, W. ZHANG (editors), Lecture Notes in Computer Science, Springer, October 2006, vol. 4218, pp. 292-306, http://www.lsv.ens-cachan.fr/Publis/PAPERS/PDF/BHR-atva06.pdf

[56] TH. CHATAIN, P. GASTIN, N. SZNAJDER. *Natural Specifications Yield Decidability for Distributed Synthesis of Asynchronous Systems*, in "Proceedings of the 35th International Conference on Current Trends in Theory and Practice of Computer Science (SOFSEM'09)", Špindlerův Mlýn, Czech Republic, M. NIELSEN, A. KUČERA, P. BRO MILTERSEN, C. PALAMIDESSI, P. TŮMA, F. VALENCIA (editors), Lecture Notes in Computer Science, Springer, January 2009, vol. 5404, pp. 141-152 [*DOI :* 10.1007/978-3-540-95891-8_16], http://www.lsv.ens-cachan.fr/Publis/PAPERS/PDF/CGS-sofsem09.pdf

[57] T. CHATAIN, S. HAAR, L. JEZEQUEL, L. PAULEVÉ, S. SCHWOON. *Characterization of Reachable Attractors Using Petri Net Unfoldings*, in "CMSB 2014", Manchester, United Kingdom, P. MENDES, J. DADA, K. SMALLBONE (editors), LNCS/LNBI, Springer International Publishing, November 2014, vol. 8859, 14 p. [*DOI :* 10.1007/978-3-319-12982-2_10], https://hal.archives-ouvertes.fr/hal-01060450

[58] G. CHIOLA, C. DUTHEILLET, G. FRANCESCHINIS, S. HADDAD. *Stochastic Well-Formed Colored Nets and Symmetric Modeling Applications*, in "IEEE Transactions on Computers", November 1993, vol. 42, n^o 11, pp. 1343-1360, http://www.lsv.ens-cachan.fr/Publis/PAPERS/PS/CDFH-toc93.ps

[59] A. CHURCH. *Logic, arithmetics, and automata*, in "Proc. of Int. Congr. of Mathematicians", 1962, pp. 23–35

[60] R. DEBOUK, D. TENEKETZIS. *Coordinated decentralized protocols for failure diagnosis of discrete-event systems*, in "Journal of Discrete Event Dynamical Systems: Theory and Application", 2000, vol. 10, pp. 33–86

[61] D. EL HOG-BENZINA, S. HADDAD, R. HENNICKER. *Process Refinement and Asynchronous Composition with Modalities*, in "Proceedings of the 2nd International Workshop on Abstractions for Petri Nets and Other Models of Concurrency (APNOC'10)", Braga, Portugal, N. SIDOROVA, A. SEREBRENIK (editors), June 2010, http://www.lsv.ens-cachan.fr/Publis/PAPERS/PDF/EHH-apnoc10.pdf

[62] J. ESPARZA, K. HELJANKO. *Unfoldings - A Partial-Order Approach to Model Checking*, EATCS Monographs in Theoretical Computer Science, Springer, 2008

[63] É. FABRE, A. BENVENISTE, C. JARD, S. HAAR. *Diagnosis of Asynchronous Discrete Event Systems, a Net Unfolding Approach*, in "IEEE Trans. Aut. Control", 2003, vol. 48 (5), pp. 714-727

[64] É. FABRE, A. BENVENISTE, C. JARD, S. HAAR. *Distributed monitoring of concurrent and asynchronous systems*, in "Discrete Event Dynamic Systems: theory and application", 2005, vol. 15 (1), pp. 33-84, Preliminary version: Proc. CONCUR 2003, LNCS 2761, pp.1–28, Springer

[65] B. FINKBEINER, S. SCHEWE. *Uniform distributed synthesis*, in "Proc. of the 20th IEEE Annual Symposium on Logic in Computer Science (LICS'05)", IEEE Computer Society Press, 2005, pp. 321–330

[66] P. GASTIN, B. LERMAN, M. ZEITOUN. *Distributed games with causal memory are decidable for series-parallel systems*, in "Proceedings of the 24th Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS'04)", Chennai, India, K. LODAYA, M. MAHAJAN (editors), Lecture Notes in Computer Science, Springer, December 2004, vol. 3328, pp. 275-286, http://www.lsv.ens-cachan.fr/Publis/PAPERS/PDF/GLZ-fsttcs04.pdf

[67] P. GASTIN, N. SZNAJDER, M. ZEITOUN. *Distributed synthesis for well-connected architectures*, in "Proceedings of the 26th Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS'06)", Kolkata, India, N. GARG, S. ARUN-KUMAR (editors), Lecture Notes in Computer Science, Springer, December 2006, vol. 4337, pp. 321-332 [*DOI :* 10.1007/11944836_30], http://www.lsv.ens-cachan.fr/Publis/PAPERS/PDF/GSZ-fsttcs2006.pdf

[68] S. HAAR, A. BENVENISTE, É. FABRE, C. JARD. *Partial Order Diagnosability Of Discrete Event Systems Using Petri Net Unfoldings*, in "42nd IEEE Conference on Decision and Control (CDC)", 2003

[69] S. HAAR. *Probabilistic Cluster Unfoldings*, in "Fundamenta Informaticae", 2003, vol. 53 (3-4), pp. 281-314

[70] S. HAAR, C. JARD, G.-V. JOURDAN. *Testing Input/Output Partial Order Automata*, in "Proc. TESTCOM/FATES", LNCS, Springer, 2007, vol. 4581, pp. 171-185, LNCS 4581

[71] O. KUPFERMAN, M. Y. VARDI. *Synthesizing Distributed Systems*, in "Proc. of the 16th IEEE Annual Symposium on Logic in Computer Science (LICS'01)", IEEE Computer Society Press, 2001

[72] S. LAFORTUNE, Y. WANG, T.-S. YOO. *Diagnostic Decentralisé Des Systèmes A Evénements Discrets*, in "Journal Europeen des Systèmes Automatisés (RS-JESA)", August 2005, vol. 99, n⁰ 99, pp. 95–110

[73] K. G. LARSEN, P. PETTERSSON, W. YI. *Compositional and symbolic model-checking of real-time systems*, in "Proc. of RTSS 1995", IEEE Computer Society, 1995, pp. 76-89

[74] S. MOHALIK, I. WALUKIEWICZ. *Distributed Games*, in "Proc. of the 23rd Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS'03)", LNCS, Springer, 2003, vol. 2914, pp. 338–351

[75] M. NOUAL, D. REGNAULT, S. SENÉ. *About non-monotony in Boolean automata networks*, in "Theoretical Computer Science", 2012, vol. 504, pp. 12-25

[76] A. PNUELI, R. ROSNER. *Distributed reactive systems are hard to synthesize*, in "Proc. of the 31st Annual IEEE Symposium on Foundations of Computer Science (FOCS'90)", IEEE Computer Society Press, 1990, vol. II, pp. 746–757

[77] L. RICKER, K. RUDIE. *Know Means No: Incorporating Knowledge into Discrete-Event Control Systems*, in "IEEE Transactions on Automatic Control", September 2000, vol. 45, n⁰ 9, pp. 1656–1668

[78] L. RICKER, K. RUDIE. *Knowledge Is a Terrible Thing to Waste: Using Inference in Discrete-Event Control Problems*, in "IEEE Transactions on Automatic Control", MarchSeptember 2007, vol. 52, n⁰ 3, pp. 428–441

[79] C. RODRÍGUEZ, S. SCHWOON, V. KHOMENKO. *Contextual Merged Processes*, in "34th International Conference on Applications and Theory of Petri Nets (ICATPN'13)", Italy, Lecture Notes in Computer Science, Springer, 2013, vol. 7927, pp. 29-48 [*DOI :* 10.1007/978-3-642-38697-8_3], https://hal.archives-ouvertes.fr/hal-00926202

[80] H. L. S. YOUNES, R. G. SIMMONS. *Statistical probabilistic model checking with a focus on time-bounded properties*, in "Inf. Comput.", September 2006, vol. 204, pp. 1368–1409 [*DOI :* 10.1016/J.IC.2006.05.002], http://dl.acm.org/citation.cfm?id=1182767.1182770