
IN PARTNERSHIP WITH:
**CNRS**

**Ecole normale supérieure de Cachan**

Activity Report 2014

# Project-Team MEXICO

## Modeling and Exploitation of Interaction and Concurrency

IN COLLABORATION WITH: Laboratoire specification et vérification (LSV)

RESEARCH CENTER
**Saclay - Île-de-France**

THEME
**Proofs and Verification**

# Table of contents

# Project-Team MEXICO

**Keywords:** Concurrency, Discrete Event Systems, Distributed System, Formal Methods, Model Of Computation

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

# 1. Members

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

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

**Engineer**
Alban Linard [Inria, Engineer]

**PhD Students**
Benoît Barbot [ENS Cachan, PhD Student, until Nov 2014]
Hernán Ponce de León [Inria, PhD Student, until Nov 2014]
Simon Theissing [Inria, PhD Student with SystemX]
Aiswarya Cyriac [ENS Cachan, PhD Student, until Jan 2014]

**Post-Doctoral Fellow**
Loïg Jezequel [Inria, PostDoc, until Aug 2014]

**Administrative Assistant**
Thida Iem [Inria, Assistant]

**Others**
Konstantinos Athanasiou [Inria, Intern, from Apr 2014 until Aug 2014]
Francisco Andres Gimenez [CNRS, Intern, from Mar 2014 until Aug 2014]

# 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, Aiswarya Cyriac, Paul Gastin, Stefan Haar, Serge Haddad, Hernán Ponce de Léon, 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, Loig Jezequel, Hernán Ponce de Léon, 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 [47]

*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], [49], 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], [73]). 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 [46], [52]). 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 [78], [79]. 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 [48]. 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.

Contexutal 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 [81] 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, Aiswarya Cyriac, Paul Gastin, Stefan Schwoon.

In a DIGITEO PhD project, we will study logical specification formalisms for concurrent recursive programs. With the advent of multi-core processors, the analysis and synthesis of such programs is becoming more and more important. However, it cannot be achieved without more comprehensive formal mathematical models of concurrency and parallelization. Most existing approaches have in common that they restrict to the analysis of an over- or underapproximation of the actual program executions and do not focus on a behavioral semantics. In particular, temporal logics have not been considered. Their design and study will require the combination of prior works on logics for sequential recursive programs and concurrent finite-state programs.

### *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 restict to very particular models with specialized verification techniques.

### *3.1.6. Testing*

**Participants:** Benedikt Bollig, Paul Gastin, Stefan Haar, Hernán Ponce de Léon.

*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], [53] 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 developped 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 [77], [72], [75], [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], [57]. 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, Benjamin Monmege.

### *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. [51], [44], 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 [55] 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 [45]; 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 [50].
- 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 [83] .

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 [74] [56]. 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 [43]. 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 [54]. 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 Project (see below) on self-aware networks, and will be searching new cooperations.

## 4.2. Transport Systems

**Participants:** Stefan Haar, Simon Theissing.

We participate in the project 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. 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.

## 4.3. Biological Systems

**Participants:** Stefan Haar, Serge Haddad, Stefan Schwoon, Thomas Chatain, Loïg Jezequel.

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 [33], 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 [76]). 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. New Software and Platforms

## 5.1. Software

### *5.1.1. Software*

*5.1.1.1. Mole/Cunf: unfolders for Petri Nets*

**Participant:** Stefan Schwoon [correspondant].

Mole computes, given a safe Petri net, a finite prefix of its unfolding. It is designed to be compatible with other tools, such as PEP and the Model-Checking Kit, which are using the resulting unfolding for reachability checking and other analyses. The tool Mole arose out of earlier work on Petri nets. Details on Mole can be found at http://www.lsv.ens-cachan.fr/~schwoon/tools/mole/. Mole served as an experimentation platform for several of our papers in recent years, most recently [33].

In the context of MExICo, we have created a new tool called Cunf [82], which is able to handle contextual nets, i.e. Petri nets with read arcs [80].While in principle every contextual net can be transformed into an equivalent Petri net and then unfolded using Mole, Cunf can take advantage of their special features to do the job faster and produce a smaller unfolding. Cunf has recently been extended with a verification component that takes advantage of these features; More details can be found at http://www.lsv.ens-cachan.fr/~rodrigue/tools/cunf/. Moreover, Cunf has been integrated into the CosyVerif environment (see section 5.2.1). Cunf has also participated in the Model Checking Contest held at the Petri Nets conference in 2013 and 2014.

*5.1.1.2. TOURS: Testing On Unfolded Reactive Systems*

**Participant:** Hernán Ponce de León [correspondant].

The MOLE - based testing tool TOURS [42] has been developed in 2014 with the help of intern Konstantinos Athanasiou, jointly supervised by Hernán Ponce de León and Stefan Schwoon of the MExICo team at LSV); it has served successfully to experiment the partial-order based testing methodology on a scalable benchmark example (elevator control).

*5.1.1.3. COSMOS : a Statistical Model Checker for the Hybrid Automata Stochastic Logic*

**Participant:** Benoît Barbot [correspondant].

COSMOS is a statistical model checker for the Hybrid Automata Stochastic Logic (HASL). HASL employs Linear Hybrid Automata (LHA), a generalization of Deterministic Timed Automata (DTA), to describe accepting execution paths of a Discrete Event Stochastic Process (DESP), a class of stochastic models which includes, but is not limited to, Markov chains. As a result HASL verification turns out to be a unifying framework where sophisticated temporal reasoning is naturally blended with elaborate reward-based analysis. COSMOS takes as input a DESP (described in terms of a Generalized Stochastic Petri Net), an LHA and an expression Z representing the quantity to be estimated. It returns a confidence interval estimation of Z; recently, it has been equipped with functionalities for rare event analysis. COSMOS is written in C++ and is freely available to the research community.

Details on COSMOS can be found at http://www.lsv.ens-cachan.fr/~barbot/cosmos/

## 5.2. Platforms

### *5.2.1. Platform CosyVerif*

CosyVerif (http://www.cosyverif.org/) is a platform dedicated to the formal specification and verification of dynamic systems. It allows to specify systems using several formalisms (such as automata and Petri nets), and to run verification tools on these models. CosyVerif integrates several tools, that are mainly developed by researchers of the MeFoSyLoMa group (a Parisian verification group, http://www.mefosyloma.fr/).

The platform is client/server based. The modeler creates models on the client side, either programmatically, or in a dedicated graphical editor. Tools are then executed on the server side.

CosyVerif is available as installable bundles, that embed the client, the server, and also the tools. It is also usable through a public server hosted within the laboratory.

The platform offers a common language for the description of the models, in order to create interoperability between clients and tools. It also provides a way to define easily new formalisms within the platform, and to manipulate models that are instances of these formalisms. To the best of our knowledge, no other verification framework presents such a feature.

CosyVerif targets three different kinds of users:

- Students use this platform in two M2 courses in modeling and verification courses. *Citer les deux cours*
- Tool developers, that are usually researchers, use the platform to distribute their tools, and have a demonstration version easily available. They also use CosyVerif for tutorials in conferences or workshops *Citer Petri nets 2014*.
- Industrial case studies have used the platform since its creation to prove properties on systems in various fields, such as: transportation systems, scheduling, hardware, robotics, databases, banking systems, home automation...

The platform is managed by a steering committee consisting of researchers and engineers of three laboratories: LIP6, LIPN, LSV. This committee decides strategic orientations as well as technical choices.

This year, we have fully redesigned the platform, with two goals in mind: first, to use technologies that target better our users; and second, to provide more functionalities.

- We switched to lightweight web technologies, in order to ease the deployment and use of CosyVerif. For the users, it means that they can access a graphical editor within their web browser. They can also access the platform through an API, usable with any HTTP client.
- We improved the language for formalisms and models in order to allow the modular definition of new formalisms. We switched from a class/instance paradigm to a prototype one, that allows to represent complex models in a both efficient and usable way.
- We extended the server to handle not only executions. It is now primarily a repository of formalismes, models, services and executions, that belong to users or project. It also handles the tools executions, and the collaborative edition of models.
- We started working on a system to help building packages for the various components of the platform (client, server, tools, ...), to ease its installation. It is used to create the bundles of CosyVerif, that are available to download. Another team (Secsi) of the LSV laboratory is interested in this system, and will support its development in 2015.

All the developed software are open source and free software tools.

Two engineers have worked this year on CosyVerif:

- Francis Hulin-Hubard, part-time (CNRS engineer);
- Alban Linard, full-time (Inria engineer).

CosyVerif has been used for teaching in two master programs (Universities Paris 6 and Paris 13/Villetaneuse) It has been used in a tutorial in the Petri Nets 2014 conference.

We are currently in the process of giving a better visibility to the project, by transforming it into a consortium. Our goal is to identify industrial fields where the tools of the platform can be applied successfully, by proposing services to the industry. The strength of the platform relies on the variety of techniques offered by the tools, that adapt to a wide range of problems. In order to increase the number of techniques, we have been joined by another partner from Geneva.

# 6. New Results

## 6.1. Highlights of the Year

### *6.1.1. Active Diagnosis for Probabilistic Systems*

Diagnosis fits well with probabilistic systems since it is natural to model the uncertainty about the behaviour of a partially observed system by distributions. We had previously revisited the active diagnosis (which aims at controlling the system to make it diagnosable) in discrete event systems designing optimal decision and synthesis procedures [7]. This year, we have considered active diagnosis for probabilistic discrete event systems, obtaining again optimal procedures [26]. Furthermore we have refined the notion of active diagnosis by introducing the *safe active diagnosis* which ensures that after the control is applied, there is a positive probability that a fault never occurs. Interestingly this problem is undecidable but for finite memory controller we have shown that the problem becomes again decidable and we have designed optimal decision and synthesis procedures. Our approach has raised an issue that has not be observed by previous researchers: while in discrete event system, most variants of diagnosis are in fact equivalent, this is no more the case for probabilistic systems. So in [26], we have undertaken the task of classifying the different versions obtaining a complete landscape of the notions both in terms of relations and complexity. Furthermore we have proposed a new notion of diagnosis, the *prediagnosis* that combines the advantages of diagnosis and prediction.

### *6.1.2. Weighted automata and weighted logics*

Weighted automata are a conservative quantitative extension of finite automata that enjoys applications, e.g., in language processing and speech recognition. Their expressive power, however, appears to be limited, especially when they are applied to more general structures than words, such as graphs. To address this drawback, we have introduced weighted pebble walking automata, which allow to navigate freely in the graph and may use pebbles to mark some positions.

In [20], we have shown with examples from natural language modeling and quantitative model-checking that weighted expressions and automata with pebbles are more expressive and allow much more natural and intuitive specifications than classical ones. We have extended Kleene-Schu ¨tzenberger theorem showing that weighted expressions and automata with pebbles have the same expressive power. We focussed on an efficient translation from expressions to automata. We also proved that the evaluation problem for weighted automata can be done very efficiently if the number of reusable pebbles is low.

In [18], we have studied the expressive power of these automata on words. We have proved that two-way pebble weighted automata, one-way pebble weighted automata, and our weighted logic with transitive closure are expressively equivalent. We also gave new logical characterizations of standard recognizable series.

In [30], we addressed the more general case of graphs such as nested words, trees, pictures, Mazurkiewicz traces, ... We established that weighted pebble walking automata have the same expressive power as weighted first order logic with transitive closure logic, lifting a similar result by Engelfriet and Hoogeboom from the Boolean case to a quantitative setting.

### *6.1.3. Verification of concurrent recursive programs*

Distributed systems form a crucially important but particularly challenging domain. Designing correct distributed systems is demanding, and verifying its correctness is even more so. The main cause of difficulty here is concurrency and interaction (or communication) between various distributed components. Hence it is important to provide a framework that makes easy the design of systems as well as their analysis. There are two schools of thought on reasoning about distributed systems: one following the interleaving based semantics, and one following the visual partial-order/graph based semantics. In [23], we compare these two approaches and argue in favour of the latter. An introductory treatment of the split-width technique is also provided.

In [34], we develop a general technique based on split-width for the verification of networks of multi-threaded recursive programs communicating via reliable FIFO channels. We extend the approach of [6] to this setting. Split-width offers an intuitive visual technique to decompose our behaviour graphs such as MSCs and nested words. The decomposition is mainly a divide-and-conquer technique which naturally results in a tree decomposition. Every behaviour can now be interpreted over its decomposition tree. Properties over the behaviour naturally transfer into properties over the decomposition tree. This allows us to use tree-automata techniques to obtain decision procedures for a range of problems such as reachability, model checking against logical formalisms etc. In this way, we obtain simple, uniform and optimal decision procedures for various verification problems parametrised by split-width. Furthermore, the simple visual mechanism of split-width is as powerful as yardstick graph measures such as tree-width or clique-width. Hence it captures any class of distributed behaviours with a decidable MSO theory.

Multi-threaded recursive programs communicating via channels are turing powerful, hence their verification has focussed on under-approximation techniques. Any error detected in the under-approximation implies an error in the system. However the successful verification of the under-approximation is not as useful if the system exhibits unverified behaviours. In [24], we study controllers that observe/restrict the system so that it stays within the verified under-approximation. We identify some important properties that a good controller should satisfy. We consider an extensive under-approximation class, construct a distributed controller with the desired properties and also establish the decidability of verification problems for this class.

### *6.1.4. Regulation in Systems Biology*

*6.1.4.1. Rare events in Signalling Cascades*

The visit in 2013 of Professor Monika Heiner from Cottbus University has led to a fruitful collaboration related to statistical model checking of rare events in signalling cascades (a regulatory biological system) [25]. This work has received one of the five top paper awards of the conference. In addition, we have improved the statistical methods used in our tool Cosmos.

#### *6.1.4.2. Characterization of Reachable Attractors Using Petri Net Unfoldings*

Attractors of network dynamics represent the long-term behaviours of the modelled system. Their characterization is therefore crucial for understanding the response and differentiation capabilities of a dynamical biological system. In the scope of qualitative models of interaction networks, the computation of attractors reachable from a given state of the network faces combinatorial issues due to the state space explosion.

In [33], we have presented a new algorithm that exploits the concurrency between transitions of parallel acting components in order to reduce the search space. The algorithm relies on Petri net unfoldings that can be used to compute a compact representation of the dynamics. We have illustrated the applicability of the algorithm with Petri net models of cell signalling and regulation networks, boolean and multi-valued. The proposed approach aims at being complementary to existing methods for deriving the attractors of Boolean models, while being generic since it applies to any safe Petri net.

## 6.2. Diagnosis

### *6.2.1. Diagnosability under Weak Fairness*

In partially observed Petri nets, diagnosis is the task of detecting whether or not the given sequence of observed labels indicates that some unobservable fault has occurred. Diagnosability is an associated property of the Petri net, stating that in any possible execution an occurrence of a fault can eventually be diagnosed. In [35] we consider diagnosability under the weak fairness (WF) assumption, which intuitively states that no transition from a given set can stay enabled forever; it must eventually either fire or be disabled. Following our previous work [71] on how to perform *weak diagnosis* by exploiting the fact that weak fairness reveals faults in parallel with the current observation, sometimes even before their actual accurrence, we turn to the associated *diagnosability* problem in [35]. First, we show that a previous approach to WF-diagnosability in the literature has a major flaw, and present a corrected notion. Moreover, we present an efficient method for verifying WF-diagnosability based on a reduction to LTL-X model checking. An important advantage of this method is that the LTL-X formula is fixed ? in particular, the WF assumption does not have to be expressed as a part of it (which would make the formula length proportional to the size of the specification), but rather one exploits the ability of existing model checkers to handle weak fairness directly.

## 6.3. Asynchronous Testing

In the final year of the TECSTES project, we have extended and completed the co-ioco - based conformance and testing theory that we had developed thus far and published in [21], in several directions:

- The testing framework now provides a test generation algorithm [21] for concurrent systems specified with true concurrency models, such as Petri nets or networks of automata. The semantic model of computation of such formalisms are labeled event structures, which allow to represent concurrency explicitly.
- Our test generation algorithm based on Petri net unfolding is able to build a complete test suite w.r.t our co-ioco conformance relation [22]. In addition we propose several coverage criteria that allow to select finite prefixes of an unfolding in order to build manageable test suites.
- We propose an extension of the *ioco* conformance relation, a standard for labeled event structures, named co-ioco, allowing to deal with strong and weak concurrency. We extend the notions of test cases and test execution to labeled event structures, and give a test generation algorithm building a complete test suite for co-ioco. Further, we have introduced and exploited [21] the notions of *strong* and *weak* concurrency: strongly concurrent events must be concurrent in the implementation, while weakly concurrent ones may eventually be ordered, leading to refine *co-ioco* into the *wsc-ioco* relation accounting for weak and strong concurrency.

- The *co-ioco* relation assumes a global control and observation of the system under test, which is not usually realistic in the case of physically distributed systems. Such systems can be partially observed at each of their points of control and observation by the sequences of inputs and outputs exchanged with their environment. Unfortunately, in general, global observation cannot be reconstructed from local ones, so global conformance cannot be decided with local tests. We showed in [39] how appending time stamps to the observable actions of the system under test in order to regain global conformance, via vector clock information, from local testing.
- The MOLE - based testing tool TOURS [42] has been developed with the help of intern Konstantinos Athanasiou, jointly supervised by Hernán Ponce de León and Stefan Schwoon of the MExICo team at LSV), and successful experiments have been conducted with a scalable benchmark example (elevator control). The results show clearly how the true-concurrency approach leads to the test case required being not only smaller individually, but also that *fewer* such test cases are necessary. In addition to the conceptual and analytical enrichment, the results obtained in TECSTES thus also allow to obtain important speedups and reductions in storage space.

Hernán Ponce de León has completed his thesis [40] reporting on the above results, and very successfully defended on Nov. 7, 2014, at ENS Cachan, before the PhD committee consisting of reviewers Rob Hierons and Alex Yakovlev, examiners Thierry Jeron, Remi Morin and Pascal Poizat, and the two supervisors.

## 6.4. Reachability in MDPs

Markov decision process (MDP) provide the appropriate formalism for the control of fully observable probabilistic systems. There are three kinds of methods for their analysis: linear programming, policy iteration and value iteration. However for large scale systems, only value iteration is still available as it requires less memory than the other methods. For quantitative problems like optimal control for maximizing the discounted reward of an MDP, value iteration is equipped with a stopping criterion that ensures an error bound provided by the user. Value iteration algorithms have also been proposed for the central problem of reachability. However neither stopping criterion nor convergence rate were known for such algorithms. In [37], we have solved these two problems and based on it we have also improved the bound on the number of iterations in order to adapt the value iteration for an exact computation.

## 6.5. Parameterized Communicating Automata

As a part of our research program on concurrent systems with variable communication topology, we studied system models where the topology is *static* but *unknown*, so that it becomes a parameter of the system. In [28], we introduced parameterized communicating automata (PCAs), where finite-state processes exchange messages via rendez-vous or through bounded FIFO channels. Unlike classical communicating automata, a given PCA can be run on any network topology of bounded degree. We presented various Büchi-Elgot-Trakhtenbrot theorems for PCAs, which roughly read as follows: Let $\varphi$ be an existential MSO formula and $T$ be any of the following topology classes: pipelines, ranked trees, grids, or rings. There is a PCA that is equivalent to $\varphi$ on all topologies from $T$. In the case where each process executes a bounded number of contexts (each context restricting communication in a suitable way), we could show that PCAs are closed under complementation, are expressively equivalent to full MSO logic [29], and have a decidable emptiness problem [31]. The papers [29], [31] are a result of a collaboration with Akshay Kumar (IIT Kanpur) and Jana Schubert (TU Dresden).

## 6.6. Quantitative behaviours

Several measures have been proposed in literature for quantifying the information leaked by the public outputs of a program with secret inputs. In [32] we studied how to quantify the information leaked by a deterministic or probabilistic program when the measure of information is based on min-entropy or Shannon entropy. A direct computation of these quantities is often infeasible because of the state-explosion problem. In our paper, we model the program as a pushdown system equipped with multi-terminal decision diagrams (ADDs) and propose algorithms to compute said entropies.

The advantage of this approach is that the resulting algorithms can be easily implemented in any BDD-based model-checking tool that checks for reachability in deterministic non-recursive programs by computing program summaries. We demonstrate the validity of our approach by implementing these algorithms in a tool Moped-QLeak.

# 7. Bilateral Contracts and Grants with Industry

## 7.1. Bilateral Contracts and Grants with Industry

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

# 8. Partnerships and Cooperations

## 8.1. Regional Initiatives

### *8.1.1. DIM/LSC TECSTES - 2011-052D*

In this DIGITEO project (No. 6024), Hernán Ponce de Léon, Delphine Longuet (ParisSud) and Stefan Haar cooperate on the subject of conformance testing for concurrent systems, using Event Structures. The project started on September 1, 2011 and has ended on August 31, 2014.

## 8.2. IRT

### *8.2.1. SystemX*

**Participants:** Simon Theissing, Stefan Haar.

We participate in the project 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.

## 8.3. National Initiatives

### *8.3.1. ANR project IMPRO*

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

The Project ANR **ImpRo** ANR-2010-BLAN-0317 involves *IRCCyN* (Nantes), *IRISA* (Rennes), *LIP6*(Paris), *LSV* (Cachan), *LIAFA* (Paris) and *LIF* (Marseille). It addresses issues related to the practical implementation of formal models for the design of communication-enabled systems: such models abstract away from many complex features or limitations of the execution environment. The modeling of *time*, in particular, is usually idealized, with infinitely precise clocks, instantaneous tests or mode communications, etc. Our objective is thus to study to what extent the practical implementation of these models preserves their good properties. We aim at a generic mathematical framework to reason about and measure implementability, and then study the possibility to integrate implementability constraints in the models. A particular focus is on the combination of several sources of perturbation such as resource allocation, the distributed architecture of applications, etc. We also study implementability through control and diagnosis techniques, and apply the developed methods to a case study based on the AUTOSAR architecture, a standard in the automotive industry.

## 8.4. European Initiatives

### *8.4.1. FP7 & H2020 Projects*

*8.4.1.1. Hycon2*

Type: FP7 COOPERATION

Defi: Engineering of Networked Monitoring and Control Systems

Instrument: Network of Excellence

Objectif: Engineering of Networked Monitoring and Control systems

Duration: September 2010 - August 2014

Coordinator: CNRS

Partners: ETH Zürich, TU Berlin, TU Delft and many others.

Inria contact: C. Canudas de Wit

Abstract: Hycon2 aims at stimulating and establishing a long-term integration in the strategic field of control of complex, large-scale, and networked dynamical systems. It focuses in particular on the domains of ground and aerospace transportation, electrical power networks, process industries, and biological and medical systems.

## 8.5. International Initiatives

### *8.5.1. Inria International Partners*

*8.5.1.1. 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 (thesis in progress).

   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* and the DISCO team (Lucia Pomello and Luca Bernardinello) at University Milano-Bicocca, Italy.
3. Exchanges are frequent with Rolf Hennicker from LMU and Javier Esparza at TUM, both in Munich, Germany.
4. With the computer science and electrical engineering departments at Newcastle University, UK (Maciej Koutny, Alex Yakovlev, Victor Khomenko and Andrey Mokhov), with visits in both directions.

### *8.5.2. Participation In Other International Programs (non-Inria)*

*8.5.2.1. EGIDE: TAMTV*

Since October 2013, Benedikt Bollig has been the French coordinator of the EGIDE-Procope project TAMTV (2013/2014), which is a collaboration with LIAFA (Paris) and the University of Ilmenau (Germany).

*8.5.2.2. 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/

## 8.6. International Research Visitors

### *8.6.1. Visits of International Scientists*

- Maciej Koutny from Newcastle University came as an invited Profeessor (for ENS Cachan) from February 10 to 14 and from March 3 to 7, 2014.
- From May 12 to June 3rd, K. Narayan Kumar from CMI Chennai, India, visited to work with C. Aiswarya and Paul Gastin on controllers for distributed systems.
- From June 1 to 10, 2014, S. Akshay from IIT Bombay visited MExICo to work with Paul Gastin, on split-width techniques for timed systems.
- Stanislav Böhm from the Technical University of Ostrava visited the group from 7 october to 7 december 2014.

### *8.6.2. Internships hosted by MExICo*

Athanasiou Konstantinos - Athanasios

Date: Apr 2014 - Aug 2014

Institution: National University of Athens, Greece

Jana Schubert

Date: 30 Sept 2013 - 28 February 2014

Institution: Universität Dresden, Germany

Akshay Kumar

Date: May 10 to July 22, 2014

Institution: IIT Khanpur

### *8.6.3. Visits to International Teams*

*8.6.3.1. Shorter Visits*

- Paul Gastin visited S. Akshay at IIT Bombay twice, first January 11-17 to work on probabilistic timed systems , and then from December 7 to 19 to work on timed pushdown systems and to deliver an invited talk at FSTTCS in Delhi.
- Stefan Haar visited the PAIS lab at Higher School of Economics in Moscow from Sept. 15 to 23.

# 9. Dissemination

## 9.1. Promoting Scientific Activities

### *9.1.1. Scientific events organisation*

*9.1.1.1. General chair, Scientific chair*

- Serge Haddad was co-chair for workshop and tutorials for 34th Int. Conf. on Application and Theory of Petri nets (ATPN), Tunis, Tunisia

*9.1.1.2. Member of the organizing committee*

- Paul Gastin co-organized the Dagstuhl seminar on "Quantitative Models: Expressiveness, Analysis, and new applications", January 19 to 24, see http://www.dagstuhl.de/en/program/calendar/semhp/?semnr=14041.

### *9.1.2. Scientific events selection*

*9.1.2.1. Member of the conference program committee*

- Serge Haddad was a member of the program committees of
  - FOR-MOVES (associated with ICSOC 2014),
  - 22nd International Conference on Real Time Networks and Systems (RTNS 2014), Versailles, France,
  - 8th International Workshop on Verification and Evaluation of Computer and Communication Systems (VECOS 2014), Bejaia, Algeria,
  - PNSE 2014,
  - 12th IFAC - IEEE International Workshop on Discrete Event Systems (WODES), Cachan, France
- Stefan Haar was a member of the program committees of PNSE 2014 and ETFA 2014. He will be co-chairing the PC of ACSD 2015 (Brussels) and be a member of the PC for ICTAC 2015.
- Thomas Chatain was a member of the program committee of ACSD 2014 and will be a member of the PC of ACSD 2015.
- Stefan Schwoon was a member of the PC of ICATPN 2014 and will be in the PCs for ICATPN 2015 and SPIN 2015.
- Benedikt Bollig was a member of the scientific committee of the workshop INFINITY'14, co-located with FSTTCS'14.

*9.1.2.2. Reviewer*

- Benedikt Bollig was a reviewer for AFL'14, DLT'14, TACAS'14, FOSSACS'14, ICALP'14, CSL-LICS'14, CONCUR'14, MFCS'14, TIME'14, FSTTCS'14.
- Stefan Haar was a reviewer for CDC 2014 and ACC 2015.
- Stefan Schwoon was a reviewer for the conferences POST, CSL-LICS, and ICALP in 2014.

### *9.1.3. Journal*

*9.1.3.1. Member of the editorial board*

- 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*.

*9.1.3.2. Reviewer*

- Benedikt Bollig was a reviewer for *ACM Transactions on Computational Logic*, *Theoretical Computer Science*, and *Acta Informatica*.
- Stefan Haar was a reviewer for *Automatica*, *Journal of Computer and System Sciences*, *IEEE Transactions on Automatic Control*, and *Information Systems*
- Thomas Chatain was a reviewer for *International Journal of Foundations of Computer Science* and *ACM Transactions on Embedded Computing Systems*.
- Stefan Schwoon acted as reviewer for the journals Fundamenta Informaticae (several occasions), TOPLAS, and TECS in 2014.
- Paul Gastin and Serge Haddad are regularly reviewers for many international conferences and journals.

Serge Haddad was also a member of the AERES evaluation committee of the VERIMAG laboratory in 2014.

## 9.2. Teaching - Supervision - Juries

### 9.2.1. Teaching

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

**Master**

Benedikt Bollig and Paul Gastin, Non-Sequential Theory of Distributed Systems, 24 heures de cours, M2, ENS de Cachan

### 9.2.2. Supervision

PhD :

- Benoît Barbot, *Acceleration for Statistical Model Checking* [11]; ENS Cachan, Defence on November 20, 2014. Supervisor : Serge Haddad and Claudine Picaronny.
- Aiswarya Cyriac, *Verification of Communicating Recursive Programs via Split-width* [12], ENS Cachan, Defence on January 28, 2014; Supervisors: Paul Gastin and Benedikt Bollig
- Hernán Ponce de Léon, *Testing Concurrent Systems Through Event Structures* [13], ENS Cachan, Defence on November 7, 2014; Supervisor: Stefan Haar, Co-supervisor: Delphine Longuet (U Paris SUd)

PhD in progress :

- Simon Theissing, *Supervision for Multi-Modal Transport Systems*, since September 2013, Supervisor Stefan Haar
- Salim Perchy (Ecole Polytechnique), *D-spaces*, since November 2013, Supervisor Stefan Haar, Co-Supervisor Franck Valencia (Note : S. Perchy belongs to the COMETE team, not MExICo).

### 9.2.3. Juries

- Paul Gastin was the president of the PhD examination board of Vincent Carnino, U. Paris-Est, december 5, 2014.
- Stefan Haar was a reviewer of the theses of Sébastien CHEDOR at University of Rennes 1 and of Kari KÄHKÖNEN at Aalto University, Finland.
- Serge Haddad was a reviewer of the thesis of Ariane Piel defended in October 2014 at U. Paris-Nord (Villetaneuse), and of the thesis of Mouhamadou Tafsir Sakho at U. Orléans in December 2014. He also was a member of the PhD juries of Thomas Husja (UPMC, october 2014).
- Stefan Schwoon was a reviewer for the theses of Stanislav Böhm (University of Ostrava, Czechia) and Ala Eddine Ben Salem (Université Paris 6).

## 9.3. Popularization

- Stefan Haar gave a talk entitled "Revèle tes défauts" on fault diagnosis in the popularization series "Unithé ou café" of Inria Saclay-Idf, on February 7, 2014.

# 10. 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, n$^{o}$ 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, n$^{o}$ 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, n$^{o}$ 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. BARBOT. *Acceleration for Statistical Model Checking*, ENS Cachan, November 2014, https://hal.archives-ouvertes.fr/tel-01110159

[12] A. CYRIAC. *Verification of Communicating Recursive Programs via Split-width*, ENS Cachan, January 2014, https://hal.archives-ouvertes.fr/tel-01110177

[13] H. PONCE DE LEÓN. *Testing Concurrent Systems Through Event Structures*, Ecole Normale Superieure de Cachan, November 2014, https://hal.inria.fr/tel-01095412

### Articles in International Peer-Reviewed Journals

[14] S. AKSHAY, B. BOLLIG, P. GASTIN, M. MUKUND, K. NARAYAN KUMAR. *Distributed Timed Automata with Independently Evolving Clocks*, in "Fundamenta Informaticae", April 2014, vol. 130, n° 4, pp. 377-407, https://hal.archives-ouvertes.fr/hal-01089524

[15] M. BECCUTI, G. FRANCESCHINIS, D. CODETTA-RAITERI, S. HADDAD. *Computing Optimal Repair Strategies by Means of NdRFT Modeling and Analysis*, in "The Computer Journal", 2014, vol. 57, n° 12, pp. 1870-1892 [*DOI :* 10.1093/COMJNL/BXT134], https://hal.archives-ouvertes.fr/hal-01110130

[16] L. BERNARDINELLO, C. FERIGATO, S. HAAR, L. POMELLO. *Closed Sets in Occurrence Nets with Conflicts*, in "Fundamenta Informaticae", 2014, vol. 133, n° 4, pp. 323-344, https://hal.archives-ouvertes.fr/hal-01091152

[17] B. BOLLIG, A. CYRIAC, P. GASTIN, M. ZEITOUN. *Temporal Logics for Concurrent Recursive Programs: Satisfiability and Model Checking*, in "Journal of Applied Logic", 2014, vol. 12, n° 4, pp. 395-416 [*DOI :* 10.1016/J.JAL.2014.05.001], https://hal.archives-ouvertes.fr/hal-01005353

[18] B. BOLLIG, P. GASTIN, B. MONMEGE, M. ZEITOUN. *Pebble Weighted Automata and Weighted Logics*, in "ACM Transactions on Computational Logic", April 2014, vol. 15, n° 2, 15 p. , 34 pages [*DOI :* 10.1145/2579819], https://hal.archives-ouvertes.fr/hal-00964994

[19] T. CHATAIN, S. HAAR. *A Canonical Contraction for Safe Petri Nets*, in "Transactions on Petri Nets and Other Models of Concurrency", 2014, vol. LNCS 8910, pp. 83-98, https://hal.archives-ouvertes.fr/hal-01091171

[20] P. GASTIN, B. MONMEGE. *Adding Pebbles to Weighted Automata - Easy Specification & Efficient Evaluation*, in "Theoretical Computer Science", 2014, vol. 534, pp. 24-44, https://hal.archives-ouvertes.fr/hal-01091105

[21] H. PONCE DE LEÓN, S. HAAR, D. LONGUET. *Model Based Testing for Concurrent Systems with Labeled Event Structures*, in "Journal of Software Testing, Verification, and Reliability", August 2014, vol. 24, n° 7, pp. 558–590 [*DOI :* 10.1002/STVR.1543], https://hal.inria.fr/hal-00914796

[22] 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", November 2014, pp. 14-28 [*DOI :* 10.1007/S10009-014-0353-Y], https://hal.inria.fr/hal-00996000

### Invited Conferences

[23] C. AISWARYA, P. GASTIN. *Reasoning about distributed systems: WYSIWYG*, in "34th Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS'14)", New Delhi, India, December 2014, https://hal.archives-ouvertes.fr/hal-01091168

### International Conferences with Proceedings

[24] C. AISWARYA, P. GASTIN, K. NARAYAN KUMAR. *Controllers for the Verification of Communicating Multi-Pushdown Systems*, in "25th International Conference on Concurrency Theory (CONCUR)", Rome, Italy, Springer, 2014, vol. LNCS 8704, pp. 297-311, https://hal.archives-ouvertes.fr/hal-01057525

[25] B. BARBOT, S. HADDAD, M. HEINER, C. PICARONNY. *Rare Event Handling in Signalling Cascades*, in "6th International Conference on Advances in System Simulation (SIMUL'14)", Nice, France, 2014, pp. 126-131, https://hal.archives-ouvertes.fr/hal-01091144

[26] N. BERTRAND, E. FABRE, S. HAAR, S. HADDAD, L. HÉLOUËT. *Active diagnosis for probabilistic systems*, in "17th International Conference on Foundations of Software Science and Computation Structures (FoSSaCS'14)", Grenoble, France, A. MUSCHOLL (editor), Proceedings of FOSSACS 2014, Springer, April 2014 [*DOI :* 10.1007/978-3-642-54830-7_2], https://hal.inria.fr/hal-00930919

[27] N. BERTRAND, S. HADDAD, E. LEFAUCHEUX. *Foundation of Diagnosis and Predictability in Probabilistic Systems*, in "IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS'14)", New Delhi, India, December 2014, https://hal.inria.fr/hal-01088117

[28] B. BOLLIG. *Logic for Communicating Automata with Parameterized Topology*, in "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, July 2014, https://hal.archives-ouvertes.fr/hal-00872807

[29] B. BOLLIG, P. GASTIN, A. KUMAR. *Parameterized Communicating Automata: Complementation and Model Checking*, in "34th Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS'14)", New Delhi, India, December 2014, https://hal.archives-ouvertes.fr/hal-01030765

[30] B. BOLLIG, P. GASTIN, B. MONMEGE, M. ZEITOUN. *Logical Characterization of Weighted Pebble Walking Automata*, in "CSL-LICS '14", Vienna, Austria, Proceedings of the Joint Meeting of the Twenty-Third EACSL Annual Conference on Computer Science Logic (CSL) and the Twenty-Ninth Annual ACM/IEEE Symposium on Logic in Computer Science (LICS), July 2014 [*DOI :* 10.1145/2603088.2603118], https://hal.archives-ouvertes.fr/hal-01006125

[31] B. BOLLIG, P. GASTIN, J. SCHUBERT. *Parameterized Verification of Communicating Automata under Context Bounds*, in "8th Workshop on Reachability Problems in Computational Models (RP'14)", Oxford, United Kingdom, Springer, 2014, vol. LNCS 8762, pp. 45-57, https://hal.archives-ouvertes.fr/hal-00984421

[32] R. CHADHA, U. MATHUR, S. SCHWOON. *Computing Information Flow Using Symbolic Model-Checking*, in "34th Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS'14)", New Delhi, India, December 2014, pp. 505-516 [*DOI :* 10.4230/LIPICS.FSTTCS.2014.505], https://hal.archives-ouvertes.fr/hal-01110118

[33] 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 Berlin Heidelberg, November 2014, in press, https://hal.archives-ouvertes.fr/hal-01060450

[34] A. CYRIAC, P. GASTIN, K. NARAYAN KUMAR. *Verifying Communicating Multi-pushdown Systems*, in "12th International Symposium on Automated Technology for Verification and Analysis (ATVA'14)", Sydney, Australia, Springer, 2014, vol. LNCS 8837, pp. 1-17, https://hal.archives-ouvertes.fr/hal-00943690

[35] V. GERMANOS, S. HAAR, V. KHOMENKO, S. SCHWOON. *Diagnosability under Weak Fairness*, in "14th International Conference on Application of Concurrency to System Design (ACSD'14)", Tunis, Tunisia, IEEE Computer Society Press, 2014, https://hal.archives-ouvertes.fr/hal-01091162

[36] S. HADDAD, R. HENNICKER, M. H. MØLLER. *Specification of Asynchronous Component Systems with Modal I/O-Petri Nets*, in "8th Symposium on Trustworthy Global Computing (TGC'13)", Buenos Aires, Argentina, Revised Selected Papers of the 8th Symposium on Trustworthy Global Computing (TGC'13), Springer, 2014, vol. LNCS 8358, pp. 219-234, https://hal.archives-ouvertes.fr/hal-01091099

[37] S. HADDAD, B. MONMEGE. *Reachability in MDPs: Refining Convergence of Value Iteration*, in "8th Workshop on Reachability Problems in Computational Models (RP'14)", Oxford, United Kingdom, Proceedings of the 8th Workshop on Reachability Problems in Computational Models (RP'14), Springer, 2014, vol. LNCS 8762, pp. 125-137, https://hal.archives-ouvertes.fr/hal-01091122

[38] F. KORDON, F. HULIN-HUBARD. *BenchKit, a Tool for Massive Concurrent Benchmarking*, in "14th International Conference on Application of Concurrency to System Design (ACSD'14)", Tunis, Tunisia, IEEE Computer Society Press, 2014, https://hal.archives-ouvertes.fr/hal-01091157

[39] H. PONCE DE LEÓN, S. HAAR, D. LONGUET. *Distributed Testing of Concurrent Systems: Vector Clocks to the Rescue*, in "Theoretical Aspects of Computing", Bucharest, Romania, September 2014 [*DOI :* 10.1007/978-3-319-10882-7_22], https://hal.inria.fr/hal-00996002

[40] H. PONCE DE LEÓN, 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

### Other Publications

[41] L. BRANDAN-BRIONES, A. MADALINSKI, H. PONCE DE LEÓN. *Parallel Diagnosability Analysis with LTL-X Model Checking based on Petri Net Unfoldings*, September 2014, Workshop on Principles of Diagnosis, https://hal.inria.fr/hal-00915478

[42] H. PONCE DE LEÓN, S. SCHWOON, K. ATHANASIOU. *TOURS (Testing On Reactive Unfolded Systems)*, November 2014, tool, https://hal.inria.fr/hal-01097360

## References in notes

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

[44] 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

[45] 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

[46] 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

[47] 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

[48] 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

[49] 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

[50] 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

[51] 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

[52] 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

[53] 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

[54] 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

[55] 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

[56] 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

[57] 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

[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] S. HAAR, C. RODRÍGUEZ, S. SCHWOON. *Reveal Your Faults: It's Only Fair!*, in "Proceedings of the 13th International Conference on Application of Concurrency to System Design (ACSD'13)", Barcelona, Spain, M. PIETKIEWICZ-KOUTNY, M. T. LAZARESCU (editors), IEEE Computer Society Press, July 2013, pp. 120-129 [*DOI :* 10.1109/ACSD.2013.15], http://www.lsv.ens-cachan.fr/Publis/PAPERS/PDF/HRS-acsd13.pdf

[72] 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

[73] 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, n$^{o}$ 99, pp. 95–110

[74] 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

[75] 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

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

[77] 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

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

[79] 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, n$^{o}$ 3, pp. 428–441

[80] C. RODRÍGUEZ, S. SCHWOON, P. BALDAN. *Efficient contextual unfolding*, in "Proceedings of the 22nd International Conference on Concurrency Theory (CONCUR'11)", Aachen, Germany, J.-P. KATOEN, B. KÖNIG (editors), Lecture Notes in Computer Science, Springer, September 2011, vol. 6901, pp. 342-357 [*DOI :* 10.1007/978-3-642-23217-6_23], http://www.lsv.ens-cachan.fr/Publis/PAPERS/PDF/RSB-concur11.pdf

[81] 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

[82] C. RODRÍGUEZ, S. SCHWOON. *Cunf*, January 2011, n$^{o}$ N/A, https://hal.inria.fr/hal-00779948

[83] 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