RESEARCH CENTRE

**Saclay - Île-de-France**

**IN PARTNERSHIP WITH:**

**CNRS, Ecole normale supérieure de Cachan**

2020

# ACTIVITY REPORT

## Project-Team

# MEXICO

## Modeling and Exploitation of Interaction and Concurrency

**IN COLLABORATION WITH: Laboratoire specification et vérification (LSV)**

**DOMAIN**

**Algorithmics, Programming, Software and Architecture**

**THEME**

**Proofs and Verification**

# Contents

# Project-Team MEXICO

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

## Keywords

**Computer sciences and digital sciences**

A2.3. – Embedded and cyber-physical systems
A2.3.2. – Cyber-physical systems
A2.3.3. – Real-time systems
A2.4.1. – Analysis
A2.4.2. – Model-checking
A6.4.1. – Deterministic control
A6.4.3. – Observability and Controlability
A7.1. – Algorithms
A7.1.1. – Distributed algorithms
A7.2. – Logic in Computer Science
A7.3.1. – Computational models and calculability
A8.1. – Discrete mathematics, combinatorics
A8.2. – Optimization
A8.7. – Graph theory
A8.8. – Network science
A8.9. – Performance evaluation
A8.11. – Game Theory

**Other research topics and application domains**

B1.1.2. – Molecular and cellular biology
B1.1.7. – Bioinformatics
B1.1.10. – Systems and synthetic biology
B3.6. – Ecology
B7.1. – Traffic management
B7.2.1. – Smart vehicles

# 1 Team members, visitors, external collaborators

## Research Scientists

- Stefan Haar [Team leader, Inria, Senior Researcher, HDR]
- Matthias Fuegger [CNRS, Researcher]

## Faculty Members

- Thomas Chatain [École Normale Supérieure de Cachan, Associate Professor]
- Serge Haddad [École Normale Supérieure de Cachan, Professor, HDR]
- Stefan Schwoon [École Normale Supérieure de Cachan, Associate Professor, HDR]
- Lina Ye [Centrale-Supélec, Associate Professor, until Aug 2020]

## PhD Students

- Giann Karlo Aguirre Samboni [Inria, from Oct 2020]
- Mathilde Boltenhagen [CNRS]
- Igor Khmelnitsky [École Normale Supérieure de Cachan]
- Juraj Kolcak [Inria, until Jul 2020]

## Administrative Assistants

- Alexandra Merlin [Inria, from Oct 2020]
- Emmanuelle Perrot [Inria, until Aug 2020]

## External Collaborators

- Benoît Barbot [Univ Paris-Est Marne La Vallée]
- Juraj Kolcak [University of Southern Denmark,, from Aug 2020]
- Lina Ye [Centrale-Supélec, from Sep 2020]

# 2 Overall objectives

## 2.1 Scientific Objectives

**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.

## 2.2 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.

## 2.3 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*.

## 2.4 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.

## 2.5 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.

When no complete model of the underlying dynamic system is available, the analysis of logs may allow to reconstruct such a model, or at least to infer some properties of interest; this activity, which has emerged over the past 10 years on the international level, is referred to as **process mining**. In this emerging activity, we have contributed to unfolding-based process discovery [CI-146], and the study of process alignments [CI-121, CI-96, CI-83, CI-60, CI-33].

Finally, over the past years *biological* challenges have come to the center of our work, in two different directions:

1. **(Re-)programming in discrete concurrent models.** Cellular regulatory networks exhibit highly complex concurrent behaviours that is influenced by a high number of perturbations such as mutations. We are in particular investigating discrete models, both in the form of boolean networks and of Petri nets, to harness this complexity, and to obtain viable methods for two interconnected and central challenges:

   - find *attractors*, i.e. long-run stable states or sets of states, that indicate possible phenotypes of the organism under study, and
   - determine *reprogramming* strategies that apply perturbations in such a way as to steer the cell's long-run behaviour into some desired phenotype, or away from an undesired one.

2. **Process mining @ MExICo** The use of process models has increased in the last decade due to the advent of the process mining field. Process mining techniques aim at discovering, analyzing and enhancing formal representations of the real processes executed in any digital environment. These processes can only be observed by the footprints of their executions, stored in form of *event logs*. An event log is a collection of traces and is the input of process mining techniques. The derivation of an accurate formalization of an underlying process opens the door to the continuous improvement and analysis of the processes within an information system.

   Process models often use true concurrency to represent actions that appear in logs with different permutations.

   Among the important challenges in process mining, *conformance checking* is a crucial one: to assess the quality of a model (automatically discovered or manually designed) in describing the observed behavior, i.e., the event log.

   MExICo contributes to process mining, a field which discovers and manipulates true concurrency models and questions about their conformance to recorded event logs.

3. **Distributed Algorithms in wild or synthetic biological systems.** Since the arrival of Matthias Fuegger in the team, we also work, on the multi-cell level, with a distributed algorithms' view on microbiological systems, both with the goal to model and analyze existing microbiological systems as distributed systems, and to design and implement distributed algorithms in synthesized microbiological systems. Major long-term goals are drug production and medical treatment via synthesized bacterial colonies.

# 3 Research program

## 3.1 Concurrency

**Keywords:** Concurrency; Semantics; Automatic Control ; Diagnosis ; Verification.

**Participants** Thomas Chatain, Philippe Dague, Stefan Haar, Serge Haddad, Stefan Schwoon.

### Glossary

**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 $S$.

**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.

### 3.1.1 Diagnosis

**Participants** Stefan Haar, Serge Haddad, Stefan Schwoon, Philippe Dague, Lina Ye.

*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 [1] 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.

**Active Diagnosis.** Depending on the possible observations, a discrete-event system may be diagnosable or not. Active diagnosis aims at controlling the system to render it diagnosable. We have established in [5] a memory-optimal diagnoser whose delay is at most twice the minimal delay, whereas the memory required to achieve optimal delay may be highly greater. We have also provided solutions for parametrized active diagnosis, where we automatically construct the most permissive controller respecting a given delay. Further, we introduced four variants of diagnosability (FA, IA, FF, IF) in (finite) probabilistic systems (pLTS) depending whether one considers (1) finite or infinite runs and (2) faulty or all runs. The corresponding decision problems are PSPACE-complete. A key ingredient of the decision procedures was a characterisation of diagnosability by the fact that a random run almost surely lies in an open set whose specification only depends on the qualitative behaviour of the pLTS. For infinite pLTS, this characterisation still holds for FF-diagnosability but with a $G_\delta$ set instead of an open set and also for IF-and IA-diagnosability when pLTS are finitely branching. Surprisingly, FA-diagnosability cannot be characterised in this way even in the finitely branching case. Further extensions are under way, in particular in passing to *prediction* and *prevention* of faults prior to their occurrence.

**Asynchronous Diagnosis.** In asynchronous partial-order based diagnosis with Petri nets, 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 [41].

**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 the team. In 2019, a new property, manifestability, weaker than diagnosability (dual in some sense to opacity) has been studied in the context of automata and timed automata.

**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 [50, 53], 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 [48, 51]). 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 [40, 44]). 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 [54, 55]. 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.

### Hybrid Systems

**Participants** Philippe Dague, Lina Ye, Serge Haddad.

Hybrid systems constitute a model for cyber-physical systems which integrates continuous-time dynamics (modes) governed by differential equations, and discrete transitions which switch instantaneously

from one mode to another. Thanks to their ease of programming, hybrid systems have been integrated to power electronics systems, and more generally in cyber-physical systems. In order to guarantee that such systems meet their specifications, classical methods consist in finitely abstracting the systems by discretization of the (infinite) state space, and deriving automatically the appropriate mode control from the specification using standard graph techniques.

Diagnosability of hybrid systems has also been studied through an abstraction / refinement process in terms of timed automata.

#### Contextual Nets

**Participants** 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 [49].

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 [42]. In recent work, we further studied this subject from a theoretical and practical perspective, allowing us to develop concrete, efficient data structures and algorithms and a tool (Cunf) that improves upon existing state of the art. This work led to the PhD thesis of César Rodríguez in 2014 .

Contextual unfoldings deal well with two sources of state-space explosion: concurrency and shared resources. Recently, we proposed an improved data structure, called *contextual merged processes* (CMP) to deal with a third source of state-space explosion, i.e. sequences of choices. The work on CMP [56] 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.

### 3.1.2 Process Mining

MExICo introduced *anti-alignments* as a tool for conformance checking. The idea of anti-alignment is to search, for a model $N$ and a log $L$, what are the runs of $N$ which differ as much as possible from all the runs in $L$. Among other uses, anti-alignments serve as witnesses for imprecisions of the model, therefore, they are used to measure precision. MExICo designed and implemented several algorithms to compute and approximate anti-alignments.

MExICo has also been contributing to clustering of log traces.

Perspectives about process mining in MExICo include model repair, i.e. design and implementation of techniques to incrementally improve models in order to make them fit better to observed logs, including when the log itself grows continuously.

Another direction is to handle models which manipulate data and real time, in order to propose more accurate representation of the log traces when the events carry some additional information (time stamps, identifiers, quantities, costs…)

## 3.2 Management of Quantitative Behavior

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

**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 Probabilistic distributed Systems

**Participants** Stefan Haar, Serge Haddad.

### 3.3.1 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. [38, 57], 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 [58], 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 [45] 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 Distributed Markov Decision Processes

**Participants** 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 [39]; 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.

## 3.4 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 [43].
- 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 [47].
- 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 [59] .

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 [6]); (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.

## 3.5 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 [52] [46]. 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.

# 4 Application domains

## 4.1 Telecommunications

**Participants** Stefan Haar, Serge Haddad.

MExICo's research is motivated by problems of system management in several domains, such as:

- 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.

Currently, we have no active cooperation on these subjects.

## 4.2 Biological Regulation Networks

**Participants** Thomas Chatain, Matthias Fuegger, Stefan Haar, Serge Haddad, Juraj Kolcak, Hugues Mandon, Stefan Schwoon.

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. We have applied 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, the 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. 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. Our current research focusses cellular reprogramming on the one hand, and distributed algorithms in wild or synthetic biological systems on the other. The latter is a distributed algorithms' view on microbiological systems, both with the goal to model and analyze existing microbiological systems as distributed systems, and to design and implement distributed algorithms in synthesized microbiological systems. Envisioned major long-term goals are drug production and medical treatment via synthesized bacterial colonies. We are approaching our goal of a distributed algorithm's view of microbiological systems from several directions: (i) Timing plays a crucial role in microbiological systems. Similar to modern VLSI circuits, dominating loading effects and noise render classical delay models unfeasible. In previous work we showed limitations of current delay models and presented a class of new delay models, so called involution channels. In [26] we showed that involution channels are still in accordance with Newtonian physics, even in presence of noise. (ii) In [7] we analyzed metastability in circuits by a three-valued Kleene logic, presented a general technique to build circuits that can tolerate a certain degree of metastability at its inputs, and showed the presence of a computational hierarchy. Again, we expect metastability to play a crucial role in microbiological systems, as similar to modern VLSI circuits, loading effects are pronounced. (iii) We studied agreement problems in highly dynamic networks without stability guarantees [28], [27]. We expect such networks to occur in bacterial cultures where bacteria communicate by producing and sensing small signal molecules like AHL. Both works also have theoretically relevant implications: The work in [27] presents the first approximate agreement protocol in a multidimensional space with time complexity independent of the dimension, working also in presence of Byzantine faults. In [28] we proved a tight lower bound on convergence rates and time complexity of asymptotic and approximate agreement in dynamic and classical static fault models. (iv) We are currently working with Manish Kushwaha (INRA), and Thomas Nowak (LRI) on biological infection models for E. coli colonies and M13 phages.

In the context of the ESCAPE project (PhD thesis of G.K. Aguirre Samboni, started in October 2020) we are now extending our research on causal analysis of complex biological networks to the domain of *ecosystems*.

### 4.3 Transportation Systems

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

- **Autonomous Vehicles.** The validation of safety properties is a crucial concern for the design of computer guided systems, in particular for automated transport systems. Our approach consists in analyzing the interactions of a randomized environment (roads, cross-sections, etc.) with a vehicle controller.
- **Multimodal Transport Networks.** We are interested in predicting and harnessing the propagation of perturbations across different transportation modes.

## 5 Social and environmental responsibility

### 5.1 Footprint of research activities

The carbon footprint of our activities is generic for office work, and probably strongest in traveling. While the latter came essential to a halt in 2020 because of the COVID pandemic, we believe that even in the

future, intelligent use of online cooperation and communication can help limit the inevitable footprint of travel to the crucial activities of cooperation and networking, avoiding physical meetings when possible.

### 5.2 Impact of research results

With our Project *ESCAPE*, we are hoping for a strong impact on **ecosystem analysis and management.** Further, the research on biological regulation networks has the potential for enabling e.g. evaluation and design of medical therapies in epigenetic contexts.

## 6 Highlights of the year

While the pandemics slowed down our activity in almost all fields, the team managed to produce new results that compare well in terms of originality and quantity with the output of previous years. To single out the most unusual fact; our first publication in *Nature communications* on the novel, *most permissive* semantics of boolean Networks as models for biological networks [8, 18].

## 7 New software and platforms

### 7.1 New software

#### 7.1.1 COSMOS

**Keyword:** Model Checker

**Functional Description:** 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.

It is easy to generate and use a C code for discrete Simulink models (using only discrete blocks, which are sampled at fixed intervals) using MathWorks tools. However, it limits the expressivity of the models. In order to use more diverse Simulink models and control the flow of a multi-model simulation (with Discrete Event Stochastic Processes) we developed a Simulink Simulation Engine embedded into Cosmos.

COSMOS is written in C++

**URL:** http://www.lsv.ens-cachan.fr/~barbot/cosmos/

**Authors:** Hilal Djafri, Paolo Ballarini

**Contacts:** Benoît Barbot, Serge Haddad

**Participants:** Benoît Barbot, Hilal Djafri, Marie Duflot-Kremer, Paolo Ballarini, Serge Haddad

#### 7.1.2 CosyVerif

**Functional Description:** CosyVerif 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.

**URL:** http://www.cosyverif.org/

**Contact:** Serge Haddad

**Participants:** Alban Linard, Fabrice Kordon, Laure Petrucci, Serge Haddad

**Partners:** LIP6, LSV, LIPN (Laboratoire d'Informatique de l'Université Paris Nord)

#### 7.1.3 Mole

**Functional Description:** 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.

**URL:** http://www.lsv.ens-cachan.fr/~schwoon/tools/mole/

**Contact:** Stefan Schwoon

**Participant:** Stefan Schwoon

# 8 New results

## 8.1 Active Prediction for Discrete Event Systems

A central task in partially observed controllable system is to detect or prevent the occurrence of certain events called faults. Systems for which one can design a controller avoiding the faults are called actively safe. Otherwise, one may require that a fault is eventually detected, which is the task of diagnosis. Systems for which one can design a controller detecting the faults are called actively diagnosable; we can build here on our past work in [5]. An intermediate requirement is prediction, which consists in determining that a fault will occur whatever the future behaviour of the system. When a system is not predictable, one may be interested in designing a controller to make it so. Here we study the latter problem, called active prediction, and its associated property, active predictability. In other words, we investigate in [24] how to determine whether or not a system enjoys the active predictability property, i.e., there exists an active predictor for the system. Our contributions are threefold. From a semantical point of view, we refine the notion of predictability by adding two quantitative requirements: the minimal and maximal delay before the occurence of the fault, and we characterize the requirements fulfilled by a controller that performs predictions. Then we show that active predictability is EXPTIME-complete where the upper bound is obtained via a game-based approach. Finally we establish that active predictability is equivalent to active safety when the maximal delay is beyond a threshold depending on the size of the system, and we show that this threshold is accurate by exhibiting a family of systems fulfilling active predictability but not active safety.

## 8.2 Philosophers may dine - definitely !

In [32], we refine and extend the theory of Communicating Sequential Processes (CSP) of Hoare and Roscoe, whose denotational semantics of the Failure/Divergence Model was first formalized in Isabelle/HOL in 1997 to cope with infinite alphabets. We analyse a family of refinement notions, including some new ones. Better definitions allow us to clarify a number of obscure points in the classical literature, for example concerning the relationship between deadlock freeness and livelock freeness. As a result, we have a modern environment for formal proofs of concurrent systems that allow to combine general infinite processes with locally finite ones in a logically safe way. We demonstrate a number of verification-techniques for two examples: the Copy Buffer and Dijkstra's Dining Philosopher Problem of an arbitrary size.

## 8.3 A Coloured Petri Nets Based Attack Tolerance Framework

It is well-known that web services become very vulnerable when being attacked, especially in the situation where service continuity is one of the most important requirements. In [33], we propose a Coloured Petri Nets based method for attack tolerance by modelling and analysing basic behaviours of attack-network

interaction, attack detectors and their tolerance solutions. Furthermore, for complex attacks composed from basic ones, their corresponding tolerance solutions can be constructed from the corresponding basic solutions. The validity of our method is demonstrated through a case study on attack tolerance in cloud-based medical information storage.

## 8.4 Property-Directed Verification of Recurrent Neural Networks

We present in [35] a property-directed approach to verifying recurrent neural networks (RNNs). To this end, we learn a deterministic finite automaton as a surrogate model from a given RNN using active automata learning. This model may then be analyzed using model checking as verification technique. The term property-directed reflects the idea that our procedure is guided and controlled by the given property rather than performing the two steps separately. We show that this not only allows us to discover small counterexamples fast, but also to generalize them by pumping towards faulty flows hinting at the underlying error in the RNN

## 8.5 Guarded Autonomous Transitions Increase Conciseness and Expressiveness of Timed Automata

Timed Automata (TA) are an appropriate model for specifying timed requirements for Continuous Time Markov Chains (CTMC). However in order to keep tractable the model checking of a TA over a CTMC, temporal logics based on TA, like CSLTA, restrict TA to have a single clock and to be deterministic (DTA). Different variants of DTAs have been proposed to address the issue of their expressiveness and conciseness. In [30] we study the effect of two possible features: (1) autonomous transitions which are triggered by time elapsing in addition to synchronized transitions and (2) transitions guarded by propositional formulas instead of propositional formulas guarding locations. We first show that autonomous guarded transitions increase the expressiveness of DTAs (as already shown for guarded locations). Then we identify a hierarchy of DTAs subclasses all equivalent to DTAs without guarded autonomous transitions and we analyze their respective conciseness. In particular we show that eliminating resets in autonomous transitions implies an exponential blow-up, while eliminating autonomous transitions without reset can be performed in polynomial time if decision diagrams are used. Finally we compare TA with guarded transitions to TA with guarded locations showing that the former model is exponentially more concise than the latter one.

## 8.6 Dynamic Recursive Petri Nets

n the early two-thousands, Recursive Petri nets (RPN) have been introduced in order to model distributed planning of multi-agent systems for which counters and recursivity were necessary. While having a great expressive power, RPN suffer two limitations: (1) they do not include more general features for transitions like reset arcs, transfer arcs, etc. (2) the initial marking associated the recursive "call" only depends on the calling transition and not on the current marking of the caller. Here we introduce Dynamic Recursive Petri nets (DRPN) which address these issues. We show in [27] that the standard extensions of Petri nets for which decidability of the coverability problem is preserved are particular cases of DPRN. Then we establish that w.r.t. coverability languages, DRPN are strictly more expressive than RPN. Finally we prove that the coverability problem is still decidable for DRPN.

## 8.7 Concurrency in Boolean networks

Boolean networks (BNs) are widely used to model the qualitative dynamics of biological systems. Besides the logical rules determining the evolution of each component with respect to the state of its regulators, the scheduling of component updates can have a dramatic impact on the predicted behaviours. In [3, 14], we explore the use of Read (contextual) Petri Nets (RPNs) to study dynamics of BNs from a concurrency theory perspective. After showing bi-directional translations between RPNs and BNs and analogies between results on synchronism sensitivity, we illustrate that usual updating modes for BNs can miss plausible behaviours, i.e., incorrectly conclude on the absence/impossibility of reaching specific configurations. We propose an encoding of BNs capitalizing on the RPN semantics enabling more behaviour

than the generalized asynchronous updating mode. The proposed encoding ensures a correct abstraction of any multivalued refinement, as one may expect to achieve when modelling biological systems with no assumption on its time features.

### 8.8 Minimal coverability tree construction made complete and efficient.

Downward closures of Petri net reachability sets can be finitely represented by their set of maximal elements called the minimal coverability set or Clover. Many properties (coverability, boundedness, ...) can be decided using Clover, in a time proportional to the size of Clover. So it is crucial to design algorithms that compute it efficiently. We present in [23] a simple modification of the original but incomplete Minimal Coverability Tree algorithm (MCT), computing Clover, which makes it complete: it memorizes accelerations and fires them as ordinary transitions. Contrary to the other alternative algorithms for which no bound on the size of the required additional memory is known, we establish that the additional space of our algorithm is at most doubly exponential. Furthermore we have implemented a prototype MinCov which is already very competitive: on benchmarks it uses less space than all the other tools and its execution time is close to the one of the fastest tool.

### 8.9 Expressiveness and Conciseness of Timed Automata for the Verification of Stochastic Models.

Timed Automata are a well-known formalism for specifying timed behaviours. In [26] we are concerned with Timed Automata for the specification of timed behaviour of Continuous Time Markov Chains (CTMC), as used in the stochastic temporal logic CSLTA. A timed path formula of CSLTA is specified by a Deterministic Timed Automaton (DTA) that features two kinds of transitions: synchronizing transitions (triggered by CTMC transitions) and autonomous transitions (triggered when a clock reaches a given threshold). Other definitions of CSLTA are based on DTAs that do not include autonomous transitions. This raises the natural question: do autonomous transitions enhance expressiveness and/or conciseness of DTAs? We prove in [26] that this is the case and we provide a syntactical characterization of DTAs for which autonomous transitions do not add expressive power, but allow one to define exponentially more concise DTAs.

### 8.10 Commodification of accelerations for the Karp and Miller Construction.

Karp and Miller's algorithm is based on an exploration of the reachability tree of a Petri net where, the sequences of transitions with positive incidence are accelerated. The tree nodes of Karp and Miller are labeled with $\omega$-markings representing (potentially infinite) coverability sets. This set of $\omega$-markings allows us to decide several properties of the Petri net, such as whether a marking is coverable or whether the reachability set is finite. The edges of the Karp and Miller tree are labeled by transitions but the associated semantic is unclear which yields to a complex proof of the algorithm correctness. In this work we introduce three concepts: abstraction, acceleration and exploration sequence. In particular, we generalize the definition of transitions to $\omega$-transitions in order to represent accelerations by such transitions. The notion of abstraction makes it possible to greatly simplify the proof of the correctness. On the other hand, for an additional cost in memory, which we theoretically evaluated, we propose in [15] n "accelerated" variant of the Karp and Miller algorithm with an expected gain in execution time. Based on a similar idea we have accelerated (and made complete) the minimal coverability graph construction, implemented it in a tool and performed numerous promising benchmarks issued from realistic case studies and from a random generator of Petri nets.

### 8.11 Diagnosis and Degradation Control for Probabilistic Systems

Systems prone to faults are often equipped with a controller whose aim consists in restricting the behaviour of the system in order to perform a diagnosis. Such a task is called active diagnosis. However to avoid that the controller degrades the system in view of diagnosis, a second objective in terms of quality of service is usually assigned to the controller. In the framework of stochastic systems, a possible specification, called safe active diagnosis requires that the probability of correctness of the infinite (random)

run is non null. We introduce and study in [9] two alternative specifications that are in many contexts more realistic. The notion of ($\gamma$,v)-fault freeness associates with each run a value depending on the discounted length of its correct prefix where the discounting factor is $\gamma$. The controller has to ensure that the average of this value is above the threshold v. The notion of $\alpha$-resiliency requires that asymptotically, at every time step, a proportion greater than $\alpha$ of correct runs remain correct. From a semantic point of view, we determine the equivalences and (non) implications between the three notions of degradations both for finite and infinite systems. From an algorithmic point of view, we establish the border between decidability and undecidability of the diagnosability problems. Furthermore in the positive case, we exhibit their precise complexity and propose a synthesis of the controller which may require an infinite memory.

## 8.12 Synchronizer-Free Digital Link Controller

In [12], we present a producer-consumer link between two independent clock domains. The link allows for metastability-free, low-latency, high-throughput communication by slight adjustments to the clock frequencies of the producer and consumer domains steered by a controller circuit. Any such controller cannot deterministically avoid, detect, nor resolve metastability. Typically, this is addressed by synchronizers, incurring a larger dead time in the control loop. We follow the approach of Friedrichs et al. (TC 2018) who proposed metastability-containing circuits. The result is a simple control circuit that may become metastable, yet deterministically avoids buffer underrun or overflow. More specifically, the controller output may become metastable, but this may only affect oscillator speeds within specific bounds. In contrast, communication is guaranteed to remain metastability-free. We formally prove correctness of the producer-consumer link and a possible implementation that has only small overhead. With SPICE simulations of the proposed implementation we further substantiate our claims. The simulation uses 65nm process running at roughly 2GHz.

## 8.13 Optimized SAT encoding of conformance checking artefacts

Conformance checking is a growing discipline that aims at assisting organizations in monitoring their processes. On its core, conformance checking relies on the computation of particular artefacts which enable reasoning on the relation between observed and modeled behavior. It is widely acknowledge that the computation of these artifacts is the lion's share of conformance checking techniques. In [11], we show how important conformance artefacts like alignments, anti-alignments or multi-alignments, defined over the Levenshtein edit distance, can be efficiently computed by encoding the problem as an optimized SAT instance. From a general perspective, the work advocates for a unified family of techniques that can compute conformance artefacts in the same way. The implementation of the techniques presented in this paper show capabilities for dealing with both synthetic and real-life instances, which may open the door for a fresh way of applying conformance checking in the near future.

## 8.14 Anti-alignments – Measuring the precision of process models and event logs

Processes are a crucial artifact in organizations, since they coordinate the execution of activities so that products and services are provided. The use of models to analyze the underlying processes is a well-known practice. However, due to the complexity and continuous evolution of their processes, organizations need an effective way of analyzing the relation between processes and models. Conformance checking techniques assess the suitability of a process model in representing an underlying process, observed through a collection of real executions. One important metric in conformance checking is to assess the precision of the model with respect to the observed executions, i.e., characterize the ability of the model to produce behavior unrelated to the one observed. In [13], we present the notion of anti-alignment as a concept to help unveiling runs in the model that may deviate significantly from the observed behavior. Using anti-alignments, a new metric for precision is proposed. The proposed anti-alignment based precision metric satisfies most of the required axioms highlighted in a recent publication. Moreover, a complexity analysis of the problem of computing anti-alignments is provided, which sheds light into the practicability of using anti-alignment to estimate precision. Experiments are provided that witness the validity of the concepts introduced.

### 8.15 Model-based trace variant analysis of event logs

The comparison of trace variants of business processes opens the door for a fine-grained analysis of the distinctive features inherent in the executions of a process in an organization. The current approaches for trace variant analysis do not consider the situation where a process model is present, and therefore, it can guide the derivation of the trace variants by considering high-level structures present in the process model. In [10], we propose a fresh alternative to trace variant analysis, which proposes a generalized notion of trace variant that incorporates concurrency and iteration. This way, the analyst may be relieved from analyzing trace variants that are essentially the same, if these aspects are disregarded. We propose a general algorithm for model based trace variant analysis which is grounded in encoding the problem into SAT, and a family of heuristic alternatives including a very light sampling technique that represents a good trade-off between quality of the trace variants identified, and the complexity of the analysis. All the techniques of the paper are implemented in two open-source tools, and experiments with publicly available benchmarks are reported.

### 8.16 Reconciling Qualitative, Abstract, and Scalable Modeling of Biological Networks

Predicting biological systems' behaviors requires taking into account many molecular and genetic elements for which limited information is available past a global knowledge of their pairwise interactions. Logical modeling, notably with Boolean Networks (BNs), is a well-established approach that enables reasoning on the qualitative dynamics of networks. Several dynamical interpretations of BNs have been proposed. The synchronous and (fully) asynchronous ones are the most prominent, where the value of either all or only one component can change at each step. In [8, 18], we prove that, besides being costly to analyze, these usual interpretations can preclude the prediction of certain behaviors observed in quantitative systems. We introduce an execution paradigm, the Most Permissive Boolean Networks (MPBNs), which offers the formal guarantee not to miss any behavior achievable by a quantitative model following the same logic. Moreover, MPBNs significantly reduce the complexity of dynamical analysis, enabling to model genome-scale networks.

### 8.17 The Involution Tool for Accurate Digital Timing and Power Analysis

In [17], we introduce the prototype of a digital timing simulation and power analysis tool for integrated circuits that supports the involution delay model (Függer et al. 2019). Unlike the pure and inertial delay models typically used in digital timing analysis tools, the involution model faithfully captures short pulse propagation and related effects. Our Involution Tool facilitates experimental accuracy evaluation of variants of involution models, by comparing their timing and power predictions to those from SPICE and standard timing analysis tools. The tool is easily customizable w.r.t. instances of the involution model and circuits, and supports automatic test case generation and parameter sweeping.

We demonstrate the capabilities of the Involution Tool by providing timing and power analysis results for three different circuits, namely, an inverter tree, the clock tree of an open-source processor, and a combinational circuit that involves multi-input NAND gates. Our evaluation uses two different technologies (15 nm and 65 nm CMOS), and three different variants of involution channels (Exp, Hill and SumExp-channels). It turns out that the timing and power predictions of all involution models are significantly better than the predictions obtained by standard digital simulations for the inverter tree and the clock tree, with the SumExp-channel channel clearly outperforming the others. For the NAND circuit, the performance of any involution model is generally comparable but not significantly better than that of standard models, however, which reveals some shortcomings of the existing involution channels for modeling multi-input gates.

### 8.18 On the Radius of Nonsplit Graphs and Information Dissemination in Dynamic Networks

In [16], we A nonsplit graph is a directed graph where each pair of nodes has a common incoming neighbor. In [16], we show that the radius of such graphs is in O(log log n), where n is the number of nodes. This is an exponential improvement on the previously best known upper bound of O(log n).

We then generalize the result to products of nonsplit graphs. The analysis of nonsplit graph products has direct implications in the context of distributed systems, where processes operate in rounds and communicate via message passing in each round: communication graphs in several distributed systems naturally relate to nonsplit graphs and the graph product concisely represents relaying messages in such networks. Applying our results, we obtain improved bounds on the dynamic radius of such networks, i.e., the maximum number of rounds until all processes have received a message from a common process, if all processes relay messages in each round. We finally connect the dynamic radius to lower bounds for achieving consensus in dynamic networks.

## 8.19 An Alignment Cost-Based Classification of Log Traces Using Machine-Learning

Conformance checking is an important aspect of process mining that identifies the differences between the behaviors recorded in a log and those exhibited by an associated process model. Machine learning and deep learning methods perform extremely well in sequence analysis. In [20], we successfully apply both a Recurrent Neural Network and a Random Forest classifiers to the problem of evaluating whether the alignment cost of a log trace to a process model is below an arbitrary threshold, and provide a lower bound for the fitness of the process model based on the classification.

## 8.20 Synthesis in the Presence of Dynamic Links

. The problem of distributed synthesis is to automatically generate a distributed algorithm, given a target communication network and a specification of the algorithm's correct behavior. Previous work has focused on static networks with an apriori fixed message size. This approach has two shortcomings: Recent work in distributed computing is shifting towards dynamically changing communication networks rather than static ones, and an important class of distributed algorithms are so-called full-information protocols, where nodes piggy-pack previously received messages onto current messages. In [19], we consider the synthesis problem for a system of two nodes communicating in rounds over a dynamic link whose message size is not bounded. Given a network model, i.e., a set of link directions, in each round of the execution, the adversary choses a link from the network model, restricted only by the specification, and delivers messages according to the current link's directions. Motivated by communication buses with direct acknowledge mechanisms we further assume that nodes are aware of which messages have been delivered. We show that the synthesis problem is decidable for a network model if and only if it does not contain the empty link that dismisses both nodes' messages.

## 8.21 Distributed Computation with Continual Population Growth

Computing with synthetically engineered bacteria is a vibrant and active field with numerous applications in bio-production, bio-sensing, and medicine. Motivated by the lack of robustness and by resource limitation inside single cells, distributed approaches with communication among bacteria have recently gained in interest. In [22], we focus on the problem of population growth happening concurrently, and possibly interfering, with the desired bio-computation. Specifically, we present a fast protocol in systems with continuous population growth for the majority consensus problem and prove that it correctly identifies the initial majority among two inputs with high probability if the initial difference is $\Omega(\sqrt{n \log n})$ where n is the total initial population. We also present a fast protocol that correctly computes the NAND of two inputs with high probability. We demonstrate that combining the NAND gate protocol with the continuous-growth majority consensus protocol, using the latter as an amplifier, it is possible to implement circuits computing arbitrary Boolean functions.

## 8.22 Drawing the Line: Basin Boundaries in Safe Petri Nets

Attractors of network dynamics represent the long-term behaviours of the modelled system. Understanding the basin of an attrac-tor, comprising all those states from which the evolution will eventually lead into that attractor, is therefore crucial for understanding the response and differentiation capabilities of a dynamical system. Building on our previous results [2] allowing to find attractors via Petri net Un-foldings,

we exploit in [25] further the unfolding technique for a backward exploration of the state space, starting from a known attractor, and show how all strong or weak basins of attractions can be explicitly computed.

### 8.23 PALS: Plesiochronous and Locally Synchronous Systems

Consider an arbitrary network of communicating modules on a chip, each requiring a local signal telling it when to execute a computational step. There are three common solutions to generating such a local clock signal: (i) by deriving it from a single, central clock source, (ii) by local, free-running oscillators, or (iii) by handshaking between neighboring modules. Conceptually, each of these solutions is the result of a perceived dichotomy in which (sub)systems are either clocked or fully asynchronous, suggesting that the designer's choice is limited to deciding where to draw the line between synchronous and asynchronous design. In contrast, we take the view in [21] that the better question to ask is how synchronous the system can and should be. Based on a distributed clock synchronization algorithm, we present a novel design providing modules with local clocks whose frequency bounds are almost as good as those of corresponding free-running oscillators, yet neighboring modules are guaranteed to have a phase offset substantially smaller than one clock cycle. Concretely, parameters obtained from a 15 nm ASIC implementation running at 2 GHz yield mathematical worst-case bounds of 30 ps on phase offset for a $32 \times 32$ node grid network.

## 9 Partnerships and cooperations

### 9.1 National initiatives

- Thomas Chatain, Stefan Haar, Serge Haddad and Stefan Schwoon are participating in the ANR Project ALGORECELL.
- Matthias Függer participates in the ANR project FREDDA on verification and synthesis of distributed algorithms.

### 9.2 Regional initiatives

- Matthias Függer co-organizes the DIGICOSME working group HICDIESMEUS.
- Matthias Függer co-leads the CARE & U. Paris-Saclay project **ETSHI** on efficient test strategies for SARS-CoV-2 in healthcare institutions.
- Stefan Haar co-organizes the DIGICOSME working group THEOBIOR2, and leads the DIGICOSME DOCTORAL RESEARCH PROJECT **ESCAPE** with Franck Pommereau.

## 10 Dissemination

### 10.1 Promoting scientific activities

#### 10.1.1 Scientific events: organisation

- Matthias Függer has co-organized the CELLS'20 workshop on Computing among Cells at DISC'20.

**General chair, scientific chair**

- Serge Haddad is member of the steering committe of the *Internation Conference on Theory and Applications of Petri Nets (ICATPN)*.
- Matthias Függer is the topic chair for Digital Design at the forthcoming IEEE DDECS 2021.

**Member of the organizing committees**

- Matthias Függer was member of the steering committee of IEEE ASYNC 2020.

**Member of Conference Program Committees**

- Thomas Chatain was a member of the Program Committees of International Conference on Application and Theory of Petri Nets and Concurrency (Petri Nets'20) and International Conference on Process Mining (ICPM'20).
- Matthias Függer was a member of the Program Committees of IEEE DDECS 2020 and IEEE ASYNC 2020.
- Lina Ye was a member of the Program Committee for the Thirty-Fifth AAAI Conference on Artificial Intelligence (AAAI-21).
- Stefan Haar was a member of the program committees for *International Conference on Application and Theory of Petri Nets and Concurrency (Petri Nets'20)* and the associated workshop *Algorithms and Theories for the Analysis of Event Data 2020 (ATAED 2020)*.

**Reviewer**

- Matthias Függer was a reviewer for conferences FSTTCS'20, ASYNC'20, DDECS'20, ICALP'20, DISC'20,

#### 10.1.2 Journal

**Reviewer**

- Thomas Chatain was a reviewer for journals *Theoretical Computer Science* and *Discrete Events Dynamic Systems*.
- Matthias Függer was a reviewer for journals Microelectronics Reliability, IEEE Transactions on Parallel and Distributed Systems, Biochemical Society Transactions, and BioDesign Research.
- Stefan Haar was a reviewer for the journals *Automatica, Transactions on Software Engineering*, and *IEEE Transactions on Automatic Control*.
- Lina Ye was a reviewer for conferences CDC'20, ACC'21, and journals Science of Computer Programming, Discrete Event Dynamic Systems and IEEE Transactions on Automatic Control.

**Member of the editorial boards**

- Stefan Haar is an associate editor for *Journal of Discret Event Dynamic Systems: Theory and Application*

### 10.2 Teaching - Supervision - Juries

#### 10.2.1 Teaching

- MATHIAS FUEGGER; Master's level:
  - *Initiation à la recherche*¸ 10 h EQTD, M1, ENS Paris-Saclay, France
  - *How To Clock Your Computer* (remote lecture, 4h per week) at MPI-INF, Germany together with Christoph Lenzen, Moti Medina, Andreas Steininger, Danny Dolev, Ian Jones, and Milos Krstic.
- STEFAN HAAR , Master :
  - *Analyse de la dynamique des systèmes biologiques*¸ 10 h EQTD, M1, Université PAris-Saclay, France
- SERGE HADDAD is head of the Computer Science department of ENS Paris-Saclay. He teaches basic and advanced algorithmics (L3) and probabilistic features of computer science (M1).

- STEFAN SCHWOON
  - Responsable L3 Informatique, ENS Paris-Saclay
  - Enseignement au M1 MPRI : cours *Initiation à la Vérification* (22,5h)
  - Enseignement au L3 Info : cours *Architecture et Système* (45h), projet *Programmation orienté objet* (15h), TD *Langages Formels* (22,5h)
  - Enseignement à l'Aggrétation Maths Option Informatique: cours *Algorithmique* (22,5h)

#### 10.2.2 Supervision

- SERGE HADDAD is supervising with Alain Finkel the PhD thesis of Igor Khmelnitsky on Verification of infinite-state systems and machine learning.
- STEFAN HAAR has been supervising, with Co-superviser Loic Paulevé at LABRI, the PhD thesis of JURAJ KOLČÁK on *Parametric Logical Regulatory Networks*, PhD research started in March 2017. He is currently supervising, with Franck Pommereau of University Evry, the PhD thesis of GIANN KARLO AGUIRRE SAMBONÌ on ***Eco**System **C**ausal **A**nalysis using **PE**tri Net Unfoldings*, started in October 2020.
- THOMAS CHATAIN has been supervising, with co-superviser Josep Carmona at Universitat Politècnica de Catalunya (Barcelona, Spain), the PhD thesis of MATHILDE BOLTENHAGEN, *Optimization Techniques for Conformance Checking and Model Repair in Process Mining*, PhD research started in November 2018.
- LINA YE has been supervising, with Co-superviser Philippe Dague at LRI, the PhD these of LULU HE, *Robustness Analysis of Real-Time Systems*, PhD research started in February 2019.
- Matthias Függer has been co-supervising Corbin Hopper (MSc) with Thomas Nowak (LRI) and Manish Kushwaha (INRAE). He has been co-supervising Amit Pathania (Postdoc) with Thomas Nowak (LRI) and Manish Kushwaha (INRAE). He is currently (within an informal arrangement) co-supervising Bilal Manssouri and Victoria Andaur (both, BSc and now MSc) with Thomas Nowak and Janna Burmann (LRI) and Manish Kushwaha (INRAE).

#### 10.2.3 Juries

- THOMAS CHATAIN
- PHILIPPE DAGUE
- STEFAN HAAR was examiner and jury president for the PhD defence of Marco Romanelli on *Machine learning methods for privacy protection: leakage measurement and mechanism design* at École Polytechnique in the fall of 2020.

# 11 Scientific production

## 11.1 Major publications

[1] B. Bérard, S. Haar, S. Schmitz and S. Schwoon. 'The Complexity of Diagnosability and Opacity Verification for Petri Nets'. In: *Fundamenta Informaticae* 161.4 (2018), pp. 317–349. DOI: 10.3233/FI-2018-1706.

[2] T. Chatain, S. Haar, L. Jezequel, L. Paulevé and S. Schwoon. 'Characterization of Reachable Attractors Using Petri Net Unfoldings'. In: *CMSB 2014*. Ed. by P. Mendes, J. Dada and K. Smallbone. Vol. 8859. LNCS/LNBI. Manchester, United Kingdom: Springer International Publishing, Nov. 2014, p. 14. DOI: 10.1007/978-3-319-12982-2_10. URL: https://hal.archives-ouvertes.fr/hal-01060450.

[3] T. Chatain, S. Haar, J. Kolčák, L. Paulevé and A. Thakkar. 'Concurrency in Boolean networks'. In: *Natural Computing* (2019).

[4] S. Friedrichs, M. Függer and C. Lenzen. 'Metastability-Containing Circuits'. In: *IEEE Transactions on Computers* 67.8 (2018). DOI: 10.1109/TC.2018.2808185.

[5] S. Haar, S. Haddad, T. Melliti and S. Schwoon. 'Optimal constructions for active diagnosis'. In: *Journal of Computer and System Sciences* 83.1 (2017), pp. 101–120.

[6] S. Haddad, J. Mairesse and H.-T. Nguyen. 'Synthesis and Analysis of Product-form Petri Nets'. In: *Fundamenta Informaticae* 122.1-2 (2013), pp. 147–172.

[7] J. Kolčák, D. Šafránek, S. Haar and L. Paulevé. 'Parameter Space Abstraction and Unfolding Semantics of Discrete Regulatory Networks'. In: *Theoretical Computer Science* 765 (2019), pp. 120–144. URL: https://hal.archives-ouvertes.fr/hal-01734805.

[8] L. Paulevé, J. Kolčák, T. Chatain and S. Haar. 'Reconciling Qualitative, Abstract, and Scalable Modeling of Biological Networks'. In: *Nature Communications* 11 (2020). DOI: 10.1038/s41467-020-18112-5. URL: https://hal.archives-ouvertes.fr/hal-02518582.

## 11.2 Publications of the year

**International journals**

[9] N. Bertrand, S. Haddad and E. Lefaucheux. 'Diagnosis and Degradation Control for Probabilistic Systems'. In: *Discrete Event Dynamic Systems* 30.4 (Dec. 2020), pp. 695–723. DOI: 10.1007/s10626-020-00320-2. URL: https://hal.inria.fr/hal-03095652.

[10] M. Boltenhagen, T. Chatain and J. Carmona. 'Model-based trace variant analysis of event logs'. In: *Information Systems* (Nov. 2020), p. 101675. DOI: 10.1016/j.is.2020.101675. URL: https://hal.inria.fr/hal-03132606.

[11] M. Boltenhagen, T. Chatain and J. Carmona. 'Optimized SAT encoding of conformance checking artefacts'. In: *Computing* 103.1 (Jan. 2021), pp. 29–50. DOI: 10.1007/s00607-020-00831-8. URL: https://hal.inria.fr/hal-03132554.

[12] J. Bund, M. Függer, C. Lenzen and M. Medina. 'Synchronizer-free Digital Link Controller'. In: *IEEE Transactions on Circuits and Systems I: Regular Papers* 67.10 (Oct. 2020), pp. 3562–3573. DOI: 10.1109/TCSI.2020.2989552. URL: https://hal.archives-ouvertes.fr/hal-03070312.

[13] T. Chatain, M. Boltenhagen and J. Carmona. 'Anti-alignments—Measuring the precision of process models and event logs'. In: *Information Systems* 98 (May 2021), p. 101708. DOI: 10.1016/j.is.2020.101708. URL: https://hal.inria.fr/hal-03132544.

[14] T. Chatain, S. Haar, J. Kolčák, L. Paulevé and A. Thakkar. 'Concurrency in Boolean networks'. In: *Natural Computing* 19.1 (2020), pp. 91–109. DOI: 10.1007/s11047-019-09748-4. URL: https://hal.inria.fr/hal-01893106.

[15] A. Finkel, S. Haddad and I. Khmelnitsky. 'Commodification of accelerations for the Karp and Miller Construction.' In: *Discrete Event Dynamic Systems* (2021). URL: https://hal.inria.fr/hal-03137054.

[16] M. Függer, T. Nowak and K. Winkler. 'On the Radius of Nonsplit Graphs and Information Dissemination in Dynamic Networks'. In: *Discrete Applied Mathematics* (2020). URL: https://hal.archives-ouvertes.fr/hal-02946849.

[17] D. Öhlinger, J. Maier, M. Függer and U. Schmid. 'The Involution Tool for Accurate Digital Timing and Power Analysis'. In: *Systems Integration* 76 (Jan. 2021), pp. 87–98. DOI: 10.1016/j.vlsi.2020.09.007. URL: https://hal.archives-ouvertes.fr/hal-03070269.

[18] L. Paulevé, J. Kolčák, T. Chatain and S. Haar. 'Reconciling Qualitative, Abstract, and Scalable Modeling of Biological Networks'. In: *Nature Communications* 11 (2020). DOI: 10.1038/s41467-020-18112-5. URL: https://hal.archives-ouvertes.fr/hal-02518582.

**International peer-reviewed conferences**

[19] B. Bérard, B. Bollig, P. Bouyer, M. Függer and N. Sznajder. 'Synthesis in Presence of Dynamic Links'. In: *Proceedings of the 11th International Symposium on Games, Automata, Logics, and Formal Verification (GandALF'20)*. GandALF'20 - 11th International Symposium on Games, Automata, Logics, and Formal Verification. Brussels (on line), Belgium, 23rd Sept. 2020. URL: https://hal.archives-ouvertes.fr/hal-02917542.

[20] M. Boltenhagen, B. Chetioui and L. Huber. 'An Alignment Cost-Based Classification of Log Traces Using Machine-Learning'. In: ML4PM2020 - First International Workshop on Leveraging Machine Learning in Process Mining. Padua/ Virtual, Italy, 5th Oct. 2020. URL: https://hal.inria.fr/hal-03134114.

[21] J. Bund, M. Függer, C. Lenzen, M. Medina and W. Rosenbaum. 'PALS: Plesiochronous and Locally Synchronous Systems'. In: ASYNC 2020 - 26th IEEE International Symposium on Asynchronous Circuits and Systems. 26th IEEE International Symposium on Asynchronous Circuits and Systems (ASYNC). Salt Lake City, United States, 17th May 2020, pp. 36–43. DOI: 10.1109/ASYNC49171.2020.00013. URL: https://hal.archives-ouvertes.fr/hal-03070326.

[22] D.-J. Cho, M. Függer, C. Hopper, M. Kushwaha, T. Nowak and Q. Soubeyran. 'Distributed Computation with Continual Population Growth'. In: DISC 2020 - 34th International Symposium on DIStributed Computing. Vol. 179. 34th International Symposium on Distributed Computing (DISC 2020). virtual, Germany, 13th Oct. 2020, 7:1–7:17. DOI: 10.4230/LIPIcs.DISC.2020.6. URL: https://hal.archives-ouvertes.fr/hal-02946883.

[23] A. Finkel, S. Haddad and I. Khmelnitsky. 'Minimal Coverability Tree Construction Made Complete and Efficient'. In: FoSSaCS 2020 - 23rd International Conference on Foundations of Software Science and Computation Structures. Dublin, Ireland, 25th Apr. 2020. URL: https://hal.inria.fr/hal-02479879.

[24] S. Haar, S. Haddad, S. Schwoon and L. Ye. 'Active Prediction for Discrete Event Systems'. In: FSTTCS 2020 - 40th IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science. Goa / Virtual, India, 14th Dec. 2020. URL: https://hal.archives-ouvertes.fr/hal-02951944.

[25] S. Haar, L. Paulevé and S. Schwoon. 'Drawing the Line: Basin Boundaries in Safe Petri Nets'. In: CMSB 2020 - 18th International Conference on Computational Methods in Systems Biology. Konstanz / Online, Germany: https://cmsb2020.uni-saarland.de/, 2020. DOI: 10.1007/978-3-030-60327-4_17. URL: https://hal.archives-ouvertes.fr/hal-02898841.

[26] S. Haddad and S. Donatelli. 'Expressiveness and Conciseness of Timed Automata for the Verification of Stochastic Models'. In: Proceedings of the 14th International Conference on Language and Automata Theory and Applications (LATA'20). Milan, Italy, 20th Sept. 2021. URL: https://hal.inria.fr/hal-03150821.

[27] S. Haddad and I. Khmelnitsky. 'Dynamic Recursive Petri Nets'. In: PETRI NETS 2020 - 41st International Conference on Application and Theory of Petri Nets and Concurrency. Paris, France, 24th June 2020. URL: https://hal.inria.fr/hal-02511321.

**Conferences without proceedings**

[28] S. Mohammadreza Fani, M. Boltenhagen and A. Wil van der. 'Prototype Selection using Clustering and Conformance Metrics for Process Discovery'. In: BPI'20 - 16th International Workshop on Business Process Intelligence. Sevilla, Spain, 13th Sept. 2020. URL: https://hal.inria.fr/hal-03134093.

**Scientific books**

[29] R. Janicki, N. Sidorova and T. Chatain. *Application and Theory of Petri Nets and Concurrency - 41st International Conference, PETRI NETS 2020, Paris, France, June 24-25, 2020, Proceedings*. 30th June 2020. DOI: 10.1007/978-3-030-51831-8. URL: https://hal.inria.fr/hal-03132586.

**Scientific book chapters**

[30] S. Donatelli and S. Haddad. 'Guarded Autonomous Transitions Increase Conciseness and Expressiveness of Timed Automata'. In: *FORMATS 2020: Formal Modeling and Analysis of Timed Systems*. 2020, pp. 215–230. DOI: 10.1007/978-3-030-57628-8_13. URL: https://hal.inria.fr/hal-03136066.

[31] M. Függer, M. Kushwaha and T. Nowak. 'Digital Circuit Design for Biological and Silicon Computers'. In: *Advances in Synthetic Biology*. 14th Apr. 2020, pp. 153–171. DOI: 10.1007/978-981-15-0081-7_9. URL: https://hal.inrae.fr/hal-02549707.

[32] S. Taha, B. Wolff and L. Ye. 'Philosophers may Dine - Definitively!' In: *16th International Conference on Integrated Formal Methods*. 17th Nov. 2020. DOI: 10.1007/978-3-030-63461-2_23. URL: https://hal.archives-ouvertes.fr/hal-03134972.

[33] W. Zhou, P. Dague, L. Liu, L. Ye and F. Zaïdi. 'A Coloured Petri Nets Based Attack Tolerance Framework'. In: *the 27th Asia-Pacific Software Engineering Conference*. 2021. URL: https://hal.archives-ouvertes.fr/hal-03133790.

**Reports & preprints**

[34] T. Chatain, S. Haar, J. Kolčák and L. Paulevé. *Most Permissive Semantics of Boolean Networks*. Univ. Bordeaux, Bordeaux INP, CNRS, LaBRI, UMR5800, F-33400 Talence, France; LSV, ENS Cachan, CNRS, INRIA, Université Paris-Saclay, Cachan (France), 2020. URL: https://hal.archives-ouvertes.fr/hal-01864693.

[35] I. Khmelnitsky, D. Neider, R. Roy, B. Barbot, B. Bollig, A. Finkel, S. Haddad, M. Leucker and L. Ye. *Property-Directed Verification of Recurrent Neural Networks*. 22nd Sept. 2020. URL: https://hal.archives-ouvertes.fr/hal-03134999.

[36] A. Le Coënt, L. Fribourg, J. Vacher and R. Wisniewski. *Probabilistic reachability and control synthesis for stochastic switched systems using the tamed Euler method*. 4th Nov. 2020. URL: https://hal.archives-ouvertes.fr/hal-02987885.

[37] A. Saoud, A. Girard and L. Fribourg. *Assume-guarantee contracts for continuous-time systems*. 16th Feb. 2021. URL: https://hal.archives-ouvertes.fr/hal-02196511.

## 11.3 Cited publications

[38] A. Benveniste, É. Fabre and S. Haar. 'Markov Nets: Probabilistic Models for distributed and concurrent Systems'. In: *IEEE Transactions on Automatic Control* 48 (11) (2003). Extended version: IRISA Research Report 1538, pp. 1936–1950.

[39] S. Akshay, N. Bertrand, S. Haddad and L. Helouet. 'The steady-state control problem for Markov decision processes'. In: *Qest 2013*. Ed. by K. R. Joshi, M. Siegle, M. Stoelinga and P. R. D'Argenio. Vol. 8054. Buenos Aires, Argentina: Springer, Sept. 2013, pp. 290–304. URL: https://hal.inria.fr/hal-00879355.

[40] R. Alur, K. Etessami and M. Yannakakis. 'Realizability and Verification of MSC Graphs'. In: *Theor. Comput. Sci.* 331.1 (2005), pp. 97–114.

[41] P. Baldan, T. Chatain, S. Haar and B. König. 'Unfolding-based Diagnosis of Systems with an Evolving Topology'. In: *Information and Computation* 208.10 (Oct. 2010), pp. 1169–1192. URL: http://www.lsv.ens-cachan.fr/Publis/PAPERS/PDF/BCHK-icomp10.pdf.

[42] P. Baldan, A. Corradini, B. König and S. Schwoon. 'McMillan's complete prefix for contextual nets'. In: *Transactions on Petri Nets and Other Models of Concurrency* 1 (Nov. 2008). Volume 5100 of Lecture Notes in Computer Science, pp. 199–220.

[43] F. Baskett, K. M. Chandy, R. R. Muntz and F. G. Palacios. 'Open, Closed, and Mixed Networks of Queues with Different Classes of Customers'. In: *J. ACM* 22 (2 Apr. 1975), pp. 248–260. DOI: http://doi.acm.org/10.1145/321879.321887. URL: http://doi.acm.org/10.1145/321879.321887.

[44] P. Bhateja, P. Gastin, M. Mukund and 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)*. Ed. by E. Csuhaj-Varjú and Z. Ésik. Vol. 4639. Lecture Notes in Computer Science. Budapest, Hungary: Springer, Aug. 2007, pp. 76–87. DOI: 10.1007/978-3-540-74240-1_8. URL: http://www.lsv.ens-cachan.fr/Publis/PAPERS/PDF/BGMN-fct07.pdf.

[45] A. Bouillard, S. Haar and 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)*. Ed. by J. Ouaknine and F. Vaandrager. Vol. 5813. Lecture Notes in Computer Science. Budapest, Hungary: Springer, Sept. 2009, pp. 43–57. DOI: 10.1007/978-3-642-04368-0_6. URL: http://www.lsv.ens-cachan.fr/Publis/PAPERS/PDF/BHR-formats09.pdf.

[46] P. Bouyer, S. Haddad and 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)*. Ed. by S. Graf and W. Zhang. Vol. 4218. Lecture Notes in Computer Science. Beijing, ROC: Springer, Oct. 2006, pp. 292–306. URL: http://www.lsv.ens-cachan.fr/Publis/PAPERS/PDF/BHR-atva06.pdf.

[47] G. Chiola, C. Dutheillet, G. Franceschinis and S. Haddad. 'Stochastic Well-Formed Colored Nets and Symmetric Modeling Applications'. In: *IEEE Transactions on Computers* 42.11 (Nov. 1993), pp. 1343–1360. URL: http://www.lsv.ens-cachan.fr/Publis/PAPERS/PS/CDFH-toc93.ps.

[48] R. Debouk and D. Teneketzis. 'Coordinated decentralized protocols for failure diagnosis of discrete-event systems'. In: *Journal of Discrete Event Dynamical Systems: Theory and Application* 10 (2000), pp. 33–86.

[49] J. Esparza and K. Heljanko. *Unfoldings - A Partial-Order Approach to Model Checking*. EATCS Monographs in Theoretical Computer Science. Springer, 2008.

[50] É. Fabre, A. Benveniste, C. Jard and S. Haar. 'Distributed monitoring of concurrent and asynchronous systems'. In: *Discrete Event Dynamic Systems: theory and application* 15 (1) (2005). Preliminary version: Proc. CONCUR 2003, LNCS 2761, pp.1–28, Springer, pp. 33–84.

[51] S. Lafortune, Y. Wang and T.-S. Yoo. 'Diagnostic Decentralisé Des Systèmes A Evénements Discrets'. In: *Journal Europeen des Systèmes Automatisés (RS-JESA)* 99.99 (Aug. 2005), pp. 95–110.

[52] K. G. Larsen, P. Pettersson and W. Yi. 'Compositional and symbolic model-checking of real-time systems'. In: *Proc. of RTSS 1995*. IEEE Computer Society, 1995, pp. 76–89.

[53] P. Baldan, S. Haar and B. Koenig. 'Distributed Unfolding of Petri Nets'. In: *Proc.FOSSACS 2006*. Vol. 3921. LNCS. Extended version: Technical Report CS-2006-1. Department of Computer Science, University Ca' Foscari of Venice. Springer, 2006, pp. 126–141.

[54] L. Ricker and K. Rudie. 'Know Means No: Incorporating Knowledge into Discrete-Event Control Systems'. In: *IEEE Transactions on Automatic Control* 45.9 (Sept. 2000), pp. 1656–1668.

[55] L. Ricker and K. Rudie. 'Knowledge Is a Terrible Thing to Waste: Using Inference in Discrete-Event Control Problems'. In: *IEEE Transactions on Automatic Control* 52.3 (Mar. 2007), pp. 428–441.

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

[57] S. Abbes, A. Benveniste and 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.

[58] S. Haar. 'Probabilistic Cluster Unfoldings'. In: *Fundamenta Informaticae* 53 (3-4) (2003), pp. 281–314.

[59] H. L. S. Younes and R. G. Simmons. 'Statistical probabilistic model checking with a focus on time-bounded properties'. In: *Inf. Comput.* 204 (9 Sept. 2006), pp. 1368–1409. DOI: 10.1016/j.ic.2006.05.002. URL: http://dl.acm.org/citation.cfm?id=1182767.1182770.