

RESEARCH CENTRE  
Saclay - Île-de-France

IN PARTNERSHIP WITH:  
CNRS, Ecole normale supérieure de  
Cachan

2021  
ACTIVITY REPORT

Project-Team  
MEXICO

## Modeling and Exploitation of Interaction and Concurrency

IN COLLABORATION WITH: Laboratoire de Méthodes Formelles

### DOMAIN

Algorithmics, Programming, Software  
and Architecture

### THEME

Proofs and Verification

# Contents

<b>Project-Team MEXICO</b>	<b>1</b>
<b>1 Team members, visitors, external collaborators</b>	<b>2</b>
<b>2 Overall objectives</b>	<b>2</b>
2.1 Scientific Objectives . . . . .	2
2.2 Concurrency . . . . .	3
2.3 Interaction . . . . .	3
2.4 Quantitative Features . . . . .	3
2.5 Evolution and Perspectives . . . . .	4
<b>3 Research program</b>	<b>5</b>
3.1 Concurrency . . . . .	5
3.1.1 Diagnosis . . . . .	5
3.1.2 Process Mining . . . . .	7
3.2 Management of Quantitative Behavior . . . . .	8
3.3 Probabilistic distributed Systems . . . . .	8
3.3.1 Non-sequential probabilistic processes . . . . .	8
3.3.2 Distributed Markov Decision Processes . . . . .	9
3.4 Large scale probabilistic systems . . . . .	9
3.5 Real time distributed systems . . . . .	9
<b>4 Application domains</b>	<b>10</b>
4.1 Telecommunications . . . . .	10
4.2 Biological Regulation Networks . . . . .	10
4.3 Transportation Systems . . . . .	11
<b>5 Social and environmental responsibility</b>	<b>11</b>
5.1 Footprint of research activities . . . . .	11
5.2 Impact of research results . . . . .	12
<b>6 Highlights of the year</b>	<b>12</b>
6.1 QONFEST Multiconference . . . . .	12
6.2 AI and Process Mining . . . . .	12
<b>7 New results</b>	<b>12</b>
7.1 Optimized SAT encoding of conformance checking artefacts . . . . .	12
7.2 Anti-Alignments – Measuring The Precision of Process Models and Event Logs . . . . .	12
7.3 Computing Discounted Alignments and Anti-Alignments in Process Mining. . . . .	13
7.4 Extracting Context-Free Grammars from Recurrent Neural Networks using Tree-Automata Learning and A* Search . . . . .	13
7.5 Commodification of accelerations for the Karp and Miller Construction . . . . .	13
7.6 Synthesis in presence of dynamic links . . . . .	14
7.7 Distributed computation with continual population growth . . . . .	14
7.8 Generation of a fault-tolerant clock through redundant crystal oscillator . . . . .	15
7.9 Tight Bounds for Asymptotic and Approximate Consensus. . . . .	15
7.10 A Composable Glitch-Aware Delay Model . . . . .	16
<b>8 Partnerships and cooperations</b>	<b>16</b>
8.1 International initiatives . . . . .	16
8.1.1 Participation in International Programs . . . . .	16
8.1.2 Visits to international teams . . . . .	16
8.2 National initiatives . . . . .	16
8.3 Regional initiatives . . . . .	16

<b>9 Dissemination</b>	<b>17</b>
9.1 Promoting scientific activities	17
9.1.1 Scientific events: organisation	17
9.1.2 Scientific events: selection	17
9.1.3 Journal	18
9.1.4 Invited talks	18
9.1.5 Research administration	18
9.2 Teaching - Supervision - Juries	19
9.2.1 Teaching	19
9.2.2 Supervision	19
9.2.3 Juries	19
<b>10 Scientific production</b>	<b>20</b>
10.1 Major publications	20
10.2 Publications of the year	20
10.3 Cited publications	22

# Project-Team MEXICO

*Creation of the 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 [Ecole normale supérieure Paris-Saclay, Associate Professor]
- Philippe Dague [Univ Paris-Saclay, Emeritus, from Apr 2021]
- Serge Haddad [Ecole normale supérieure Paris-Saclay, Professor, HDR]
- Stefan Schwoon [Ecole normale supérieure Paris-Saclay, Associate Professor, HDR]

### PhD Students

- Gianni Karlo Aguirre Samboni [Inria]
- Mathilde Boltenhagen [CNRS, until Oct 2021]
- Igor Khmelnitsky [Ecole normale supérieure Paris-Saclay, until Nov 2021]

### Administrative Assistant

- Alexandra Merlin [Inria]

### External Collaborators

- Benoît Barbot [Univ Paris-Est Marne La Vallée]
- Juraj Kolcak [University Southern Denmark de Odense, until Jun 2021]
- Lina Ye [Centrale-Supélec, until Aug 2021]

## 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 [30].

**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 [39, 42], 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 to be done.

Computing diagnosis in a distributed way is only one aspect of a much vaster topic, that of *distributed diagnosis* (see [37, 40]). 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 [29, 33]). 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 [43, 44]. Automating 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 [38].

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 [31]. 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 [45] 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. [27, 46], 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 [47], 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 [34] 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 [28]; 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 [32].
- 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 [36].
- 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 [48]

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 [41] [35]. 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, Jura 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 differentiate ? 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 determine 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

### 6.1 QONFEST Multiconference

**Participants:** Benoît Barbot, Serge Haddad, Stefan Schwoon.

The **QONFEST2021** Multiconference involved **BENOÎT BARBOT** as general chair, **STEFAN SCHWOON** in the organisation and as webmaster, and **SERGE HADDAD** as program co-chair of **CONCUR**.

### 6.2 AI and Process Mining

MEXICo has made substantial scientific progress in the domains of

- learning techniques to support formal analyses, in particular extraction of context-free grammars, see point 7.4 in the 'nexwresults' sections, and
- Measuring The Precision of process models in process mining via anti-alignments, see 'new results' 7.2, 7.2 and 7.3.

## 7 New results

### 7.1 Optimized SAT encoding of conformance checking artefacts

**Participants:** Mathilde Boltenhagen, Thomas Chatain.

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 acknowledged that the computation of these artifacts is the lion's share of conformance checking techniques. Our paper [11] shows 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.

### 7.2 Anti-Alignments – Measuring The Precision of Process Models and Event Logs

**Participants:** Mathilde Boltenhagen, Thomas Chatain.

Processes are a crucial artefact in organizations, since they coordinate the execution of activities so that products and services are provided. The use of models to analyse 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 analysing 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 [12], 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. In contrast to existing metrics, anti-alignment based precision metrics satisfy 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.

### 7.3 Computing Discounted Alignments and Anti-Alignments in Process Mining.

**Participants:** Mathilde Boltenhagen, Thomas Chatain.

Alignments are a central notion in conformance checking. They establish the best possible connection between an observed trace and a process model, exhibiting the closest model run to the trace. Computing these alignments for huge amounts of traces, coming from big logs, is a computational bottleneck. We show in [21] that, for a slightly modified version of the distance function between traces and model runs, we significantly improve the execution time of an A\*-based search algorithm. We show experimentally that the alignments found with our modified distance approximate very nicely the optimal alignments for the classical distance. In [22], by introducing a discount factor in the edit distance used for the search of anti-alignments, we obtain the first efficient algorithm to approximate them. We show how this approximation is quite accurate in practice, by comparing it with the optimal results for small instances where the optimal algorithm can also compute anti-alignments. Finally, we compare the obtained precision metric with respect to the state-of-the-art metrics in the literature for real-life examples.

### 7.4 Extracting Context-Free Grammars from Recurrent Neural Networks using Tree-Automata Learning and A\* Search

**Participants:** Benoît Barbot, Serge Haddad, Igor Khmelnitsky.

In [20], we present (i) an active learning algorithm for visibly pushdown grammars and (ii) show its applicability for learning surrogate models of recurrent neural networks (RNNs) trained on context-free languages. Such surrogate models may be used for verification or explainability. Our learning algorithm makes use of the proximity of visibly pushdown languages and regular tree languages and builds on an existing learning algorithm for regular tree languages. Equivalence tests between a given RNN and a hypothesis grammar rely on a mixture of A\* search and random sampling. An evaluation of our approach on a set of RNNs from the literature shows good preliminary results.

### 7.5 Commodification of accelerations for the Karp and Miller Construction

**Participants:** Serge Haddad, Igor Khmelnitsky.

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 [16], 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 an "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.

## 7.6 Synthesis in presence of dynamic links

**Participants:** Matthias Függer.

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 a priori 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 [9], 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 chooses an arbitrary 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 the network model does not contain the empty link that dismisses both nodes' messages. We then extend the characterization to sequences of communication links that may contain empty links. We show that the synthesis problem is decidable in this case if and only if the number of consecutive empty links in all possible sequences is uniformly bounded from above.

## 7.7 Distributed computation with continual population growth

**Participants:** Matthias Függer, Da-Jung Cho.

Computing via 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 [13], 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 grows with the square root of  $n \log n$ , where  $n$  is the total initial population. We also present a fast protocol that correctly computes the Nand of two inputs with high probability. By combining Nand gates with the majority consensus protocol as an amplifier, it is possible to compute arbitrary Boolean functions. Finally, we extend the protocols to several biologically relevant settings. We simulate a plausible implementation

of a noisy Nand gate with engineered bacteria. In the context of continuous cultures with a constant outflow and a constant inflow of fresh media, we demonstrate that majority consensus is achieved only if the flow is slower than the maximum growth rate. Simulations suggest that flow increases consensus time over a wide parameter range. The proposed protocols help set the stage for bio-engineered distributed computation that directly addresses continuous stochastic population growth.

## 7.8 Generation of a fault-tolerant clock through redundant crystal oscillator

**Participants:** Matthias Függer.

Having a precise and stable clock that is still fault tolerant is a fundamental prerequisite in safety critical real-time systems. However, combining redundant independent clock sources to form a unified fault-tolerant clock supply is non-trivial, especially when redundant clock outputs are required-e.g., for supplying the replicated nodes within a TMR architecture through a clock network that does not suffer from a single point of failure. Having these outputs fail independent but still keeping them tightly synchronized is highly desirable, as it substantially eases the design of the overall architecture. In [15] we address exactly this challenge. Our approach extends an existing, ring-oscillator like distributed clock generation scheme by augmenting each of its constituent nodes with a stable clock reference. We introduce the appropriately modified algorithm and illustrate its operation by simulation experiments. These experiments further demonstrate that the four clock outputs of our circuit do not share a single point of failure, have small and bounded skew, remain stabilized to one crystal source during normal operation, do not propagate glitches from one failed clock to a correct one, and only exhibit slightly extended clock cycles during a short stabilization period after a component failure. In addition we give a rigorous formal proof for the correctness of the algorithm on an abstraction level that is close to the implementation.

## 7.9 Tight Bounds for Asymptotic and Approximate Consensus.

**Participants:** Matthias Függer.

Agreeing on a common value among a set of agents is a fundamental problem in distributed computing, which occurs in several variants: In contrast to exact consensus, approximate variants are studied in systems where exact agreement is not possible or required, e.g., in man-made distributed control systems and in the analysis of natural distributed systems, such as bird flocking and opinion dynamics. In [18], we study the time complexity of two classical agreement problems: non-terminating asymptotic consensus and terminating approximate consensus. Asymptotic consensus, requires agents to repeatedly set their outputs such that the outputs converge to a common value within the convex hull of initial values; approximate consensus requires agents to eventually stop setting their outputs, which must then lie within a predefined distance of each other. We prove tight lower bounds on the contraction ratios of asymptotic consensus algorithms subject to oblivious message adversaries, from which we deduce bounds on the time complexity of approximate consensus algorithms. In particular, the obtained bounds show optimality of asymptotic and approximate consensus algorithms presented by Charron-Bost et al. (ICALP'16) for certain systems, including the strongest oblivious message adversary in which asymptotic and approximate consensus are solvable. As a corollary we also obtain asymptotically tight bounds for asymptotic consensus in the classical asynchronous model with crashes. Central to the lower-bound proofs is an extended notion of valency, the set of reachable limits of an asymptotic consensus algorithm starting from a given configuration. We further relate topological properties of valencies to the solvability of exact consensus, shedding some light on the relation of these three fundamental problems in dynamic networks.

## 7.10 A Composable Glitch-Aware Delay Model

**Participants:** Matthias Függer.

In [25], we introduce the Composable Involution Delay Model (CIDM) for fast and accurate digital simulation. It is based on the Involution Delay Model (IDM) [Függer et al., IEEE TCAD 2020], which has been shown to be the only existing candidate model for faithful glitch propagation. The IDM, however, has shortcomings that limit its applicability. Our CIDM thus reduces the characterization effort by allowing independent discretization thresholds, improves composability and increases the modeling power by exposing canceled pulse trains at the gate interconnect. We formally show that, despite these improvements, the CIDM still retains the IDM's faithfulness.

## 8 Partnerships and cooperations

### 8.1 International initiatives

#### 8.1.1 Participation in International Programs

##### *LearNNiFy*

**Title:** New Challenges for Recurrent Neural Networks and Grammatical Inference

**Partner Institution(s):** • CNRS (B. Bollig, leader)

- MPI Kaiserslautern (D. Neider)
- Univ. Lübeck (M. Leucker)

**Participants:** Serge Haddad.

#### 8.1.2 Visits to international teams

- STEFAN HAAR visited the **Discrete Biomathematics** group of Prof. HEIKE SIEBERT at *Freie Universität Berlin*, Germany, from November 15 to 19.

### 8.2 National initiatives

- Thomas Chatain, Stefan Haar, Serge Haddad, Juraj Kolčák and Stefan Schwoon have participated in the ANR Project **ALGORECELL**.
- Matthias Függer has participated in the ANR project FREDDA on verification and synthesis of distributed algorithms and led the INS2I/CNRS funded *Emergence* project *Bacterial CONsensus (BACON)*, 2020-2021. He now leads the ANR project **Distributed Algorithms for Microbiological Systems** that has started on Oct.1st, 2021, and participates in the CARE (UPSaclay-ENS PSaclay)-funded project on *Efficient Test Strategies for SARS-CoV-2 in Healthcare Institutions (ETSHI)*.

### 8.3 Regional initiatives

- GIANN KARLO AGUIRRE SAMBONI and *Stefan Haar* are participating in the *DIGICOSME* project *EcoSystem Causal Analysis using PEtri net unfoldings (ESCAPE)* with Franck Pommereau at Univ. Evry.

## 9 Dissemination

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

### 9.1 Promoting scientific activities

#### 9.1.1 Scientific events: organisation

- PHILIPPE DAGUE is the head of the *GDR-IA*'s working group *BIOSS-IA* on the interface between AI and Bioinformatics, and has been organizing monthly online seminars as well as an on-location workshop (Paris, December 7 and 8, 2021). He also co-organized the *DATAIA* research days on AI and Security at Saclay, Dec 13, 2021.
- MATTHIAS FÜGGER organised the *CELLS Workshop on Computing among Cells: co-located with the International Symposium on DIStributed Computing (DISC)*, Oct. 2021.

#### General chair, scientific chair

- BENOÎT BARBOT was general chair of the *QONFEST2021* multiconference.
- MATTHIAS FÜGGER was Topic chair for Digital/System Design at *IEEE DDECS 2021* and a member of the steering committee for *IEEE ASYNC*.

#### Member of the organizing committees

- Stefan Schwoon participated in the organisation of the *QONFEST 2021* multiconference, in particular as webmaster.

#### 9.1.2 Scientific events: selection

##### Chair of conference program committees

- SERGE HADDAD was co-chair of the *13th Int. Conf. on Concurrency Theory (CONCUR 2021)*.

##### Member of the conference program committees

- THOMAS CHATAIN was a member of the program committees for the *3rd Int. Conf. on Process Mining (ICPM) 2021*, the *42nd Int. Conf. on Application and Theory of Petri Nets and Concurrency 2021* and *FORMATS 2021*.
- MATTHIAS FÜGGER was a member of the program committees of *IEEE DDECS 2021* and *IEEE ASYNC*.
- STEFAN HAAR was a member of the program committee for the *42nd Int. Conf. on Application and Theory of Petri Nets and Concurrency 2021*.
- SERGE HADDAD was a member of the program committees for the *PNSE workshop on Petri nets in Software Engineering*, of the *École d'Été Temps Réel*, and of the *15th Int. Conf. on Verification and Evaluation of Computer and Communication Systems VECOS 2021*.
- LINA YE was a member of the program committees of *American Control Conference (ACC 21)* and *Conf. on Decision and Control (CDC 21)*.

## Reviewer

- THOMAS CHATAIN was a reviewer for the conferences *ICPM 2021*, *FORMATS* and for the *42nd Int. Conf. on Application and Theory of Petri Nets and Concurrency 2021*.
- MATTHIAS FÜGGER was a reviewer for the conferences *FSTTCS*, *ICALP 21*, *[IEEE ASYNC, IEEE DDECS 2021 and DISC*
- STEFAN HAAR was a reviewer for the conferences *ICALP 21* and *ICPM 2021*.
- SERGE HADDAD provided reviews for *PNSE* and *VECOS 2021*).
- STEFAN SCHWOON was a reviewer for the conference *ICALP 21*.

### 9.1.3 Journal

#### Member of the editorial boards

- STEFAN HAAR is an associate editor for the *Journal of Discrete Events Dynamic Systems*, and was co-editor (with Susanna Donatelli) of the upcoming special issue on Petri nets of *Fundamenta Informaticae*.

#### Reviewer - reviewing activities

- THOMAS CHATAIN was a reviewer for *Information Systems, Logical Methods in Computer Science, Petri Nets and Other Models of Concurrency (ToPNoC)*, and *ACM Transactions on Privacy and Security*.
- MATTHIAS FÜGGER was a reviewer for *Biochemical Society Transactions, IEEE Transactions on Parallel and Distributed Systems, Microelectronics Reliability, BioDesign Research and Theoretical Computer Science*.
- STEFAN HAAR was a reviewer for *AMS Mathematical Reviews, Journal of Computer and System Sciences, Fundamenta Informaticae, IEEE Transactions on Automatic Control, Journal of Discrete Events Dynamic Systems, Automatica and Petri Nets and Other Models of Concurrency (ToPNoC)*.
- LINA YE was a reviewer for *Automatica*, *Discrete Event Dynamic Systems* and *IEEE Transactions on Automatic Control*.
- STEFAN SCHWOON was a reviewer for *Engineering, Theoretical Computer Science and Journal of Discrete Events Dynamic Systems*.

### 9.1.4 Invited talks

- PHILIPPE DAGUE gave an invited video talk at the *Metabolic Pathway Analysis Convergence (MPA 2021)*.
- STEFAN HAAR gave a talk entitled "Update on discrete networks" in the *Discrete Biomathematics* seminar at *Freie Universität Berlin*, Germany, on November 16.

### 9.1.5 Research administration

- THOMAS CHATAIN was a member of the jury for the *SIF-Gilles Kahn* award for outstanding PhD theses, and is the head of the *Interactions* pole at the LMF lab.
- PHILIPPE DAGUE is a member of the selection board for the *UDOPIA 2021* PhD thesis funding scheme.
- MATTHIAS FÜGGER is the head of the the LMF group *Concurrent and Distributed Systems*, elected member of the laboratory council for LMF and in charge of European contracts at LMF.
- STEFAN HAAR is *adjunct director for research* of the *Graduate School Computer Science* at *Université Paris-Saclay*, and head of the *Computer Science and Biology* axis at the LMF lab.

## 9.2 Teaching - Supervision - Juries

### 9.2.1 Teaching

- MATHIAS FUEGGER; Master's level:
  - *Initiation à la recherche*, 10 h EQTD, M1, ENS Paris-Saclay, France
  - Natural algorithms at Université Paris-Saclay. With Thomas Nowak.
  - Metastability-containing Synchronization Circuits. With Christoph Lenzen, Danny Dolev, Moti Medina, Ian W. Jones, and Johannes Bund. At CISP, Germany.
- STEFAN HAAR, Master :
  - *Analyse de la dynamique des systèmes biologiques*, 24 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
  - Teaching in M1 MPRI : cours *Initiation à la Vérification* (22,5h)
  - L3 Info : *Architecture et Système* (45h), project *Programmation orienté objet* (15h), TD *Langages Formels* (22,5h)
  - Agrégation Maths Option Informatique: *Algorithmique* (22,5h)

### 9.2.2 Supervision

- THOMAS CHATAIN was a supervisor (75 %, with Josep Carmona, Barcelona) of the thesis of MATHILDE BOLTENHAGEN on *Process Instance Clustering based on Conformance Checking Artefacts*, successfully defended on October 21, 2021 at ENS Paris-Saclay.
- MATTHIAS FÜGGER co-supervises, with Janna Burman and Thomas Nowak) the PhD thesis of FABRICIO CRAVO on *Distributed Circuit Design for Bacterial Systems*, started on Oct 1st, 2021 (ED STIC funded). He also supervised the master's internship of Baudouin Pilet on *Non-Markovian Growth Models of Bacteria*.
- SERGE HADDAD is co-supervising, with ALAIN FINKEL, the PhD thesis of IGOR KHMELNITSKY on *Verification of infinite-state systems and machine learning*, whose defence is scheduled for Jan 27, 2022.
- STEFAN HAAR co-supervised (with Loic Paulevé) the PhD the thesis of JURAJ KOLČÁK on *Unfoldings and Abstract Interpretation for Parametric Biological Regulatory Networks* [26], successfully defended at ENS Paris-Saclay on July 6, 2021. He currently supervises, with Franck Pommereau at Univ. Evry, the thesis of GIANN KARLO AGUIRRE SAMBONÍ, started on Oct1, 2020, on *Ecosystem causal analysis using Petri net unfoldings*.
- LINA YE (70 %) and PHILIPPE DAGUE (30 %) co-supervise the thesis of LULU HE on manifestability analysis using SMT solvers; started in 2019, the thesis is scheduled for defence in 2022.

### 9.2.3 Juries

- THOMAS CHATAIN was a reviewer of the PhD Thesis of GOVIND RAJANBABU on *Partial order reduction for timed systems* at Université de Bordeaux, on June 16, 2021.

## 10 Scientific production

### 10.1 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](https://doi.org/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](https://doi.org/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](https://doi.org/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](https://doi.org/10.1038/s41467-020-18112-5). URL: <https://hal.archives-ouvertes.fr/hal-02518582>.

### 10.2 Publications of the year

#### International journals

- [9] B. Bérard, B. Bollig, P. Bouyer, M. Függer and N. Sznajder. ‘Synthesis in presence of dynamic links’. In: *Information and Computation* (Dec. 2021), p. 104856. DOI: [10.1016/j.ic.2021.104856](https://doi.org/10.1016/j.ic.2021.104856). URL: <https://hal.inria.fr/hal-03518879>.
- [10] B. Bérard, S. Haddad, C. Picaronny, M. Safey El Din and M. Sassolas. ‘Polynomial interrupt timed automata: Verification and expressiveness’. In: *Information and Computation* 277 (Apr. 2021), p. 104580. DOI: [10.1016/j.ic.2020.104580](https://doi.org/10.1016/j.ic.2020.104580). URL: <https://hal.sorbonne-universite.fr/hal-02862629>.
- [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](https://doi.org/10.1007/s00607-020-00831-8). URL: <https://hal.inria.fr/hal-03132554>.
- [12] T. Chatain, M. Boltenhagen and J. Carmona. ‘Anti-Alignments – Measuring The Precision of Process Models and Event Logs’. In: *Information Systems* (May 2021). URL: <https://hal.inria.fr/hal-02383546>.
- [13] D.-J. Cho, M. Függer, C. Hopper, M. Kushwaha, T. Nowak and Q. Soubeyran. ‘Distributed computation with continual population growth’. In: *Distributed Computing* (23rd Oct. 2021). DOI: [10.1007/s00446-021-00404-8](https://doi.org/10.1007/s00446-021-00404-8). URL: <https://hal.inria.fr/hal-03519002>.
- [14] P. Dague. ‘Metabolic Pathway Analysis in the Presence of Biological Constraints’. In: *Computation* 9.10 (Oct. 2021), p. 111. DOI: [10.3390/computation9100111](https://doi.org/10.3390/computation9100111). URL: <https://hal.archives-ouvertes.fr/hal-03467719>.

- [15] W. Dür, M. Függer and A. Steininger. ‘Generation of a fault-tolerant clock through redundant crystal oscillators’. In: *Microelectronics Reliability* 120 (May 2021), p. 114088. DOI: [10.1016/j.microrel.2021.114088](https://doi.org/10.1016/j.microrel.2021.114088). URL: <https://hal.inria.fr/hal-03329844>.
- [16] 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>.
- [17] M. Függer, A. Kinali, C. Lenzen and B. Wiederhake. ‘Fast All-Digital Clock Frequency Adaptation Circuit for Voltage Droop Tolerance’. In: *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* (2022), pp. 1–1. DOI: [10.1109/TCAD.2021.3097599](https://doi.org/10.1109/TCAD.2021.3097599). URL: <https://hal.archives-ouvertes.fr/hal-03536906>.
- [18] M. Függer, T. Nowak and M. Schwarz. ‘Tight Bounds for Asymptotic and Approximate Consensus’. In: *Journal of the ACM (JACM)* (28th Oct. 2021). DOI: [10.1145/3485242](https://doi.org/10.1145/3485242). URL: <https://hal.archives-ouvertes.fr/hal-03408731>.
- [19] 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](https://doi.org/10.1016/j.vlsi.2020.09.007). URL: <https://hal.archives-ouvertes.fr/hal-03070269>.

#### International peer-reviewed conferences

- [20] B. Barbot, B. Bollig, A. Finkel, S. Haddad, I. Khmelnitsky, M. Leucker, D. Neider, R. Roy and L. Ye. ‘Extracting Context-Free Grammars from Recurrent Neural Networks using Tree-Automata Learning and A\* Search’. In: *ICGI 2021 - 15th International Conference on Grammatical Inference. Proceedings of the 15th International Conference on Grammatical Inference (ICGI 2020/21)*. New York City / Virtual, United States, Aug. 2021. URL: <https://hal.archives-ouvertes.fr/hal-03285433>.
- [21] M. Boltenhagen, T. Chatain and J. Carmona. ‘A Discounted Cost Function for Fast Alignments of Business Processes’. In: *BPM 2021 - 19th International Conference on Business Process Management*. Rome, France, 6th Sept. 2021. DOI: [10.1007/978-3-030-85469-0\\_17](https://doi.org/10.1007/978-3-030-85469-0_17). URL: <https://hal.inria.fr/hal-03362277>.
- [22] M. Boltenhagen, T. Chatain and J. Carmona. ‘An A\*-Algorithm for Computing Discounted Anti-Alignments in Process Mining’. In: *3rd International Conference on Process Mining*. 3rd International Conference on Process Mining. Eindhoven, Netherlands, 31st Oct. 2021. URL: <https://hal.archives-ouvertes.fr/hal-03229744>.
- [23] 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)*. LATA 2020 2021 - 14th et 15th International Conference on Language and Automata Theory and Applications. Milan, Italy, 20th Sept. 2021. URL: <https://hal.inria.fr/hal-03150821>.
- [24] I. Khmelnitsky, D. Neider, R. Roy, X. Xie, B. Barbot, B. Bollig, A. Finkel, S. Haddad, M. Leucker and L. Ye. ‘Property-Directed Verification and Robustness Certification of Recurrent Neural Networks’. In: *19th International Symposium on Automated Technology for Verification and Analysis (ATVA 2021)*. Proceedings of the 19th International Symposium on Automated Technology for Verification and Analysis (ATVA 2021). Gold Coast (Online), Australia, Oct. 2021. URL: <https://hal.archives-ouvertes.fr/hal-03285418>.
- [25] J. Maier, D. Öhlinger, U. Schmid, M. Függer and T. Nowak. ‘A Composable Glitch-Aware Delay Model’. In: *GLSVLSI '21: Great Lakes Symposium on VLSI 2021*. Virtual Event, United States: ACM, 22nd June 2021, pp. 147–154. DOI: [10.1145/3453688.3461519](https://doi.org/10.1145/3453688.3461519). URL: <https://hal.archives-ouvertes.fr/hal-03451070>.

#### Doctoral dissertations and habilitation theses

- [26] J. Kolčák. ‘Unfoldings and Abstract Interpretation for Parametric Biological Regulatory Networks’. ENS Paris-Saclay, 6th July 2021. URL: <https://hal.inria.fr/tel-03528592>.

### 10.3 Cited publications

- [27] 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.
- [28] 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>.
- [29] R. Alur, K. Etessami and M. Yannakakis. ‘Realizability and Verification of MSC Graphs’. In: *Theor. Comput. Sci.* 331.1 (2005), pp. 97–114.
- [30] 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-icom10.pdf>.
- [31] P. Baldan, A. Corradini, B. König and S. Schwoon. ‘McMillan’s complete prefix for contextual nets (long vers)’. In: *Transactions on Petri Nets and Other Models of Concurrency* 1 (Nov. 2008). Volume 5100 of Lecture Notes in Computer Science, pp. 199–220.
- [32] 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>.
- [33] 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](https://doi.org/10.1007/978-3-540-74240-1_8). URL: <http://www.lsv.ens-cachan.fr/Publis/PAPERS/PDF/BGMN-fct07.pdf>.
- [34] 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](https://doi.org/10.1007/978-3-642-04368-0_6). URL: <http://www.lsv.ens-cachan.fr/Publis/PAPERS/PDF/BHR-formats09.pdf>.
- [35] 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>.
- [36] 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>.
- [37] 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.
- [38] J. Esparza and K. Heljanko. *Unfoldings - A Partial-Order Approach to Model Checking*. EATCS Monographs in Theoretical Computer Science. Springer, 2008.
- [39] 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.
- [40] S. Lafortune, Y. Wang and T.-S. Yoo. ‘Diagnostic Décentralisé Des Systèmes A Événements Discrets’. In: *Journal Européen des Systèmes Automatisés (RS-JESA)* 99.99 (Aug. 2005), pp. 95–110.
- [41] 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.

- [42] 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.
- [43] 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.
- [44] 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.
- [45] 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](https://doi.org/10.1007/978-3-642-38697-8_3). URL: <https://hal.archives-ouvertes.fr/hal-00926202>.
- [46] 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.
- [47] S. Haar. ‘Probabilistic Cluster Unfoldings’. In: *Fundamenta Informaticae* 53 (3-4) (2003), pp. 281–314.
- [48] 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](https://doi.org/10.1016/j.ic.2006.05.002). URL: <http://dl.acm.org/citation.cfm?id=1182767.1182770>.