

RESEARCH CENTRE

**Inria Saclay Center
at Université Paris-Saclay**

IN PARTNERSHIP WITH:

CNRS, Université Paris-Saclay

2022

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

Inria

Contents

Project-Team MEXICO	1
1 Team members, visitors, external collaborators	2
2 Overall objectives	2
2.1 Concurrency	3
2.2 Interaction	3
2.3 Quantitative Features	3
2.4 Evolution and Perspectives	3
3 Research program	4
3.1 Concurrency	4
3.2 Management of Quantitative Behavior	7
3.3 Large scale probabilistic systems	9
3.4 Real time distributed systems	9
4 Application domains	9
4.1 Telecommunications	9
4.2 Biological Networks	10
4.3 Transportation Systems	11
5 Social and environmental responsibility	11
5.1 Footprint of research activities	11
5.2 Impact of research results	11
6 Highlights of the year	11
6.1 Ecosystem Modeling	11
6.2 Awards	12
7 New software and platforms	12
7.1 New software	12
7.1.1 COSMOS	12
7.1.2 CosyVerif	12
7.1.3 Mole	13
7.1.4 Ecofolder	13
7.2 New platforms	13
8 New results	13
8.1 Unfoldings for reset Petri nets	13
8.2 Avoid One's Doom: Finding Cliff-Edge Configurations in Petri Nets	13
8.3 Fast All-Digital Clock Frequency Adaptation Circuit for Voltage Droop Tolerance.	14
8.4 Analysis of Recurrent Neural Networks via Property-Directed Verification of Surrogate Models.	14
8.5 MobsPy: A Meta-species Language for Chemical Reaction Networks.	15
8.6 Analyzing Robustness of Angluin's L* Algorithm in Presence of Noise.	15
8.7 Timed Alignments in Process Mining	15
9 Bilateral contracts and grants with industry	16
10 Partnerships and cooperations	16
10.1 International research visitors	16
10.1.1 Visits of international scientists	16
10.2 National initiatives	16
10.3 Regional initiatives	16

11 Dissemination	17
11.1 Promoting scientific activities	17
11.1.1 Scientific events: selection	17
11.1.2 Journal	17
11.1.3 Invited talks	17
11.1.4 Research administration	18
11.2 Teaching - Supervision - Juries	18
11.2.1 Teaching	18
11.2.2 Supervision	18
11.2.3 Juries	19
11.2.4 Internal or external Inria responsibilities	19
12 Scientific production	19
12.1 Major publications	19
12.2 Publications of the year	20
12.3 Cited publications	21

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
- B3.4. – Risks
- 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]
- Philippe Dague [UNIV PARIS SACLAY, Emeritus]
- Matthias Fuegger [CNRS, Researcher]

Faculty Members

- Thomas Chatain [ENS PARIS-SACLAY, Associate Professor, HDR]
- Serge Haddad [ENS PARIS-SACLAY, Professor, HDR]
- Stefan Schwoon [ENS PARIS-SACLAY, Associate Professor, HDR]

PhD Students

- Gianni Karlo Aguirre Samboni [INRIA]
- Fabricio Cravo [Université Paris-Saclay]

Administrative Assistant

- Aissatou-Sadio Diallo [INRIA, from May 2022]

Visiting Scientist

- Nick Würdemann [University of Oldenburg, Germany, from Sep 2022 until Oct 2022]

External Collaborator

- Benoît Barbot [UNIV PARIS EST]

2 Overall 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.

2.1 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.

2.2 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*.

2.3 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.

2.4 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 three 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. More recently, but with a related dynamical systems approach and methods that intersect strongly with those in systems biology, we are addressing the modelling, prediction and control of *ecosystems*.
3. **Distributed Algorithms in wild or synthetic biological systems.** Since the arrival of Matthias Fuegger in the team, we have also been working, on the multi-cell level, with a distributed algorithms' view on microbiological systems. This approach has two goals: model and analyze existing microbiological systems as distributed systems on the one hand, and design and implement distributed algorithms in synthesized microbiological systems on the other. 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, Matthias Fuegger, 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.

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_δ 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 [22].

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 [29, 32], 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 [27, 30]). 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 [21, 24]). 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 [33, 34]. 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 [28].

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. 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 [35] 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.

Process Mining @ MExICo.

Participants: Thomas Chatain.

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. 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.

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 [20]; there, we strive to synthesize for a given MDP a control so as to guarantee a specific stationary behavior, rather than - as is usually done - so as to maximize some reward.

3.3 Large scale probabilistic systems

Participants: Serge Haddad.

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 [23].
- 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 [26].
- 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 [36]

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.

3.4 Real time distributed systems

Participants: Benoît Barbot, Matthias Fuegger, Thomas Chatain, Philippe Dague, Serge Haddad.

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 [31] [25]. 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 Networks

Participants: Thomas Chatain, Philippe Dague, Matthias Fuegger, Stefan Haar, Serge Haddad, 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*, in cooperation with INRAE researcher Cédric Gaucherel. The cooperation with INRAE has been intensifying in 2022 with the start of the AMI INRAE-Project SMART, jointly with INRAE RECOVER (Corinne Kurt, Franck Taillader, PhD student Souhila FOUNAS) in the fall, on modeling and assessing environmental multi-risks.

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.

Currently, no active contracts or projects in this field.

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 has been slowed down 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 Ecosystem Modeling

A framework for discrete, nondeterministic modelling of ecosystems has been given in terms of *reset Petri Nets*, for which the unfolding into ordinary occurrence nets has been formalized in [12] and implemented

in the ECOFOLDER tool, see below. The study of bifurcation in an Ecosystem towards a fatal set of states was developed in [16]

6.2 Awards

- MATHILDE BOLTENHAGEN, PhD 2021 in the MEXICO team, received the Best Process Mining PhD Dissertation Award 2022 during the Fourth International Conference on Process Mining (ICPM 2022) for her thesis entitled "Process Instance Clustering based on Conformance Checking Artefacts".

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/>

Contact: Benoît Barbot

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

7.1.4 Ecofolder

Name: Ecofolder

Keywords: Petri nets, Ecology

Functional Description: Ecofolder is a software for unfolding Petri nets with resets according to the Esparza/Römer/Vogler unfolding algorithm.

URL: <https://github.com/giannkas/ecofolder>

Publication: [aguirresamboni:hal-03750233](#)

Contact: Giann Karlo Aguirre Samboni

Participants: Giann Karlo Aguirre Samboni, Stefan Haar

7.2 New platforms

Participants: - .

8 New results

8.1 Unfoldings for reset Petri nets

Participants: Giann Karlo Aguirre Samboni, Thomas Chatain, Stefan Haar.

Reset Petri nets are a particular class of Petri nets where transition firings can remove all tokens from a place without checking if this place actually holds tokens or not. In [9], we look at partial order semantics of reset Petri nets. In particular, we propose a pomset bisimulation for comparing their concurrent behaviours. Building on this pomset bisimulation we then propose a generalization of the standard finite complete prefixes of unfolding for this class of Petri nets. In a different vein, we introduce in [12] the systematic use of (1-safe) reset Petri nets for the analysis of Ecosystems, and a dedicated unfolding procedure to represent the dynamics of a reset net in an ordinary Petri net.

8.2 Avoid One's Doom: Finding Cliff-Edge Configurations in Petri Nets

Participants: Giann-Karlo Aguirre Samboni, Stefan Haar, Stefan Schwoon, Nick Würdemann.

A crucial question in analyzing a concurrent system is to determine its long-run behaviour, and in particular, whether there are irreversible choices in its evolution, leading into parts of the reachability space from which there is no return to other parts. Casting this problem in the unifying framework of safe Petri nets, our previous work [3] has provided techniques for identifying attractors, i.e. terminal strongly connected components of the reachability space, whose attraction basins we wish to determine. In [16] we provide a solution for the case of safe Petri nets. Our algorithm uses net unfoldings and provides a map of all of the system's configurations (concurrent executions) that act as cliff-edges, i.e. any maximal extension for those configurations lies in some basin that is considered fatal. The computation turns out to require only a relatively small prefix of the unfolding.

8.3 Fast All-Digital Clock Frequency Adaptation Circuit for Voltage Droop Tolerance.

Participants: Matthias Függer.

In classical synchronous designs, supply voltage droops can be handled by accounting for them in clock margins. However, this results in a significant performance hit even if droops are rare. By contrast, adaptive strategies detect such potentially hazardous events and either initiate a rollback to a previous state or proactively reduce clock speed in order to prevent timing violations. The performance of such solutions critically depends on a very fast response to droops. State-of-the-art solutions incur synchronization delays in the order of several clock cycles to avoid, with sufficient probability, that the clock signal is affected by metastability. We present in [10] an all-digital circuit that can respond to droops within a fraction of a clock cycle. This is achieved by using potentially metastable measurement values to delay clock signals while they undergo synchronization, instead of after they are synchronized. The challenge is to ensure that this strategy does not lead to harmful glitches or metastable upsets within the circuit. To this end, we verify our solution by formally proving correctness. We complement our findings by simulations of a 65 nm ASIC design confirming the results of our analysis.

8.4 Analysis of Recurrent Neural Networks via Property-Directed Verification of Surrogate Models.

Participants: Serge Haddad, Igor Khmelnitsky, Benoît Barbot.

Angluin's L^* algorithm learns the minimal (complete) deterministic finite automaton (DFA) of a regular language using membership and equivalence queries. Its probabilistic approximatively correct (PAC) version substitutes an equivalence query by a large enough set of random membership queries to get a high level confidence to the answer. Thus it can be applied to any kind of (also non-regular) device and may be viewed as an algorithm for synthesizing an automaton abstracting the behavior of the device based on observations. In [11], we are interested on how Angluin's PAC learning algorithm behaves for devices which are obtained from a DFA by introducing some noise. More precisely we study whether Angluin's algorithm reduces the noise and produces a DFA closer to the original one than the noisy device. We propose several ways to introduce the noise: (1) the noisy device inverts the classification of words w.r.t. the DFA with a small probability, (2) the noisy device modifies with a small probability the letters of the word before asking its classification w.r.t. the DFA, and (3) the noisy device combines the classification of a word w.r.t. the DFA and its classification w.r.t. a counter automaton. Our experiments were performed on several hundred DFAs. Our main contributions, bluntly stated, consist in showing that: (1) Angluin's algorithm behaves well whenever the noisy device is produced by a random process, (2) but poorly with a

structured noise, and, that (3) almost surely randomness yields systems with non-recursively enumerable languages.

8.5 MobsPy: A Meta-species Language for Chemical Reaction Networks.

Participants: Fabricio Cravo, Matthias Függer.

Chemical reaction networks are widely used to model biochemical systems. However, when the complexity of these systems increases, the chemical reaction networks are prone to errors in the initial modeling and subsequent updates of the model.

In [14], we present the Meta-species-oriented Biochemical Systems Language (MobsPy), a language designed to simplify the definition of chemical reaction networks in Python. MobsPy is built around the notion of meta-species, which are sets of species that can be multiplied to create higher-dimensional orthogonal characteristics spaces and inheritance of reactions. Reactions can modify these characteristics. For reactants, queries allow to select a subset from a meta-species and use them in a reaction. For products, queries specify the dimensions in which a modification occurs. We demonstrate the simplification capabilities of the MobsPy language at the hand of a running example and a circuit from literature. The MobsPy Python package includes functions to perform both deterministic and stochastic simulations, as well as easily configurable plotting. The MobsPy package is indexed in the Python Package Index and can thus be installed via pip.

8.6 Analyzing Robustness of Angluin's L* Algorithm in Presence of Noise.

Participants: Serge Haddad, Igor Khmelnitsky, Benoît Barbot.

Angluin's L* algorithm learns the minimal (complete) deterministic finite automaton (DFA) of a regular language using membership and equivalence queries. Its probabilistic approximatively correct (PAC) version substitutes an equivalence query by a large enough set of random membership queries to get a high level confidence to the answer. Thus it can be applied to any kind of (also non-regular) device and may be viewed as an algorithm for synthesizing an automaton abstracting the behavior of the device based on observations. In [17], we are interested in how Angluin's PAC learning algorithm behaves for devices which are obtained from a DFA by introducing some noise. More precisely we study whether Angluin's algorithm reduces the noise and produces a DFA closer to the original one than the noisy device. We propose several ways to introduce the noise: (1) the noisy device inverts the classification of words w.r.t. the DFA with a small probability, (2) the noisy device modifies with a small probability the letters of the word before asking its classification w.r.t. the DFA, and (3) the noisy device combines the classification of a word w.r.t. the DFA and its classification w.r.t. a counter automaton. Our experiments were performed on several hundred DFAs. Our main contributions, bluntly stated, consist in showing that: (1) Angluin's algorithm behaves well whenever the noisy device is produced by a random process, (2) but poorly with a structured noise, and, that (3) almost surely randomness yields systems with non-recursively enumerable languages.

8.7 Timed Alignments in Process Mining

Participants: Thomas Chatain.

We study in [13] the conformance checking for timed models, that is, process models that consider both the sequence of events in a process as well as the timestamps at which each event is recorded. Time-aware process mining is a growing subfield of research, and as tools that seek to discover timing

related properties in processes develop, so does the need for conformance checking techniques that can tackle time constraints and provide insightful quality measures for time-aware process models. In particular, one of the most useful conformance artefacts is the alignment, that is, finding the minimal changes necessary to correct a new observation to conform to a process model. In this paper, we set our problem of timed alignment and solve two cases each corresponding to a different metric over time processes. For the first, we have an algorithm whose time complexity is linear both in the size of the observed trace and the process model, while for the second we have a quadratic time algorithm for linear process models.

In [19], we solve the case where the metrics used to compare timed processes allows mixed moves, i.e. an error on the timestamp of an event may or may not have propagated to its successors, and provide linear time algorithms for distance computation and alignment on models with sequential causal processes.

9 Bilateral contracts and grants with industry

-

Participants: - .

10 Partnerships and cooperations

Participants: Matthias Függer, Stefan Haar, Serge Haddad.

10.1 International research visitors

10.1.1 Visits of international scientists

- Nick Würdemann, PhD student from Oldenburg (Germany) visited the team twice, from March 7 to March 18 and from September 3 to October 18.

10.2 National initiatives

Matthias Függer is co-PI of the ANR project [Distributed Algorithms for Microbiological, started in September 2021. Systems \(Dreamy\)](#) 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).

Serge Haddad is a participant of the ANR project [Maveriq](#) and local coordinator at LMF for this project.

10.3 Regional initiatives

- Stefan Haar: DIGICOSME grant ESCAPE (PhD thesis, Giann Karlo Aguirre Samboni) on causal analysis of Ecosystems.
- Matthias Fuegger:
 - since 10/2022: Acquired AAP Proof of Concept de l’OI BioProbe grant on bacteria and microalgae modeling and optimization (PROCEED, 6 k€); Thomas Chatain and Stefan Haar participate in the project.

- since 10/2022: Deep reinforcement learning for distributed systems: 3-year PhD grant at the doctoral school ED STIC
- since 07/2022: 1-year RFSI Ile de France Postdoc grant on Distributed Algorithms in Birth–Death Systems (DAB, 55 k€)
- since 10/2021: Distributed Circuit Design for Bacterial Systems: 3-year PhD grant at the doctoral school ED STIC

11 Dissemination

11.1 Promoting scientific activities

Steering Committees

- Matthias Fuegger is a member of the steering committee for the IEEE symposium on Asynchronous Circuits and Systems (ASYNC).
- Serge Haddad is a member of the steering committee for the on Application and Theory of Petri Nets and Concurrency (see [here](#) for the 2022 edition)

Member of the organizing committees

- Matthias Fuegger has co-organized the *Workshop on Computing among Cells (CELLS 2022)* at the 36th international symposium on distributed computing (DISC 2022)

11.1.1 Scientific events: selection

Member of the conference program committees

- Matthias Fuegger was a member of the program committee for the IEEE International Symposium on Design and Diagnostics of Electronic Circuits and Systems (DDECS 2022)
- Stefan Haar was a member of the program committee for the *43rd International Conference on Application and Theory of Petri Nets and Concurrency*,
- Thomas Chatain was a PC member of the *4th International Conference on Process Mining (ICPM'22)*.

11.1.2 Journal

Member of editorial boards

- Stefan Haar is an associate editor for the journal *Discrete Event Dynamic Systems: Theory and Applications*

Reviewer - reviewing activities

- Numerous reviews for journals and conferences.

11.1.3 Invited talks

- On November 29, 2022, Serge Haddad gave an invited Talk at Inria Rennes on "Expressiveness of Deterministic Single-Clock Timed Automata"

11.1.4 Research administration

- Serge Haddad
 - has been, since 2020, a member of 'conseil scientifique et d'administration (CSA)' for the Labex CIMI, Toulouse, and
 - was from 2012 to 2022 a member of the 'conseil d'orientation scientifique (COS)' of the LIS laboratory, Marseille (UMR 7020), and
 - has been since 2020 the representative of ENS Paris-Saclay in the supervisory board ('Conseil') of the Graduate School Computer Science, Université Paris-Saclay.

2018-2022 Membre du conseil d'orientation scientifique (COS) du LIS de Marseille (UMR 7020)

- Stefan Haar is adjoint director for research of the UPSaclay graduate school Computer Science.

11.2 Teaching - Supervision - Juries

11.2.1 Teaching

- Serge Haddad is 'Responsable' of M2 FESUP Informatique at ENS Paris-Saclay.
- 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 teaches basic and advanced algorithmics (L3, M2 FESUP) and probabilistic features of computer science (M1).
- STEFAN SCHWOON :
 - *Responsable L3 Informatique*, ENS Paris-Saclay
 - Teaching in M1 MPRI : cours *Initiation à la Vérification*
 - L3 Info : *Architecture et Système*, project *Programmation orienté objet*, TD *Langages Formels*
 - ENSIIE: Concepts et Model-Checking

11.2.2 Supervision

- MATTHIAS FÜGGER :
 - * Raghda Elshehaby (PhD), 6 month co-supervision of visiting PhD student of Andreas Steininger (TU Wien);
 - * Zhuofan Xu (PhD), co-supervision with Benedikt Bollig (ENS Paris-Saclay) and Thomas Nowak (ENS Paris-Saclay)
 - * Thomas Chevet (Postdoc), co-supervision with Thomas Nowak (ENS Paris-Saclay);
 - * Fabricio Cravo (PhD), co-supervision with Janna Burman (Université Paris-Saclay) and Thomas Nowak (ENS Paris-Saclay)
 - * interns:
 - Misaal Bedi (MSc), co-supervision with Thomas Nowak (ENS Paris-Saclay) and Manish Kushwaha (INRAE), studies at IISER Pune, India.

- Gayathri Prakash (MSc), co-supervision with Thomas Nowak (ENS Paris-Saclay) and Manish Kushwaha (INRAE), studies at IIT-Madras, India.
 - Artur Piana Broilo (MSc), co-supervision with Thomas Nowak (ENS Paris-Saclay), studies at CentraleSupélec and UFRGS, Brazil
 - Anisse Belhadj (MSc), co-supervision with Thomas Nowak (ENS Paris-Saclay)
 - Letizia Diniz (MSc), co-supervision with Thomas Nowak (ENS Paris-Saclay)
- STEFAN HAAR is supervising,
 - with FRANCK POMMEREAU (Univ. Evry) the PhD thesis of GIANN KARLO AGUIRRE SAMBONI entitled *Ecosystem Causal Analysis Using Petri net Unfoldings*, started in the fall of 2020
 - and with CORINNE CURT (INRAE Recover) the PhD thesis of SOUHILA FOUNAS on multi-risque modeling.

In addition, he supervised the M1 internship of SUNHEANG TY on *Quantum Boolean Networks*.

11.2.3 Juries

- Philippe Dague was supervisor of the Ph Thesis of LULU HE on *Formal verification at design stage of diagnosis related properties for discrete event and real-time systems*, defended in May 2022 at Université Paris-Saclay.
- Thomas Chatain was an external examiner for the PhD thesis of Hanin Abdulrahman, University of Newcastle-upon-Tyne, UK, supervisor Jason Steggles, and a member of the jury of the PhD of Emily Clément, University of Rennes, France, supervisors Thierry Jéron, Nicolas Markey and David Mentré.
- Stefan Haar was
 - an examiner for the thesis of ABDUL MATJIHNORDHEEN on June 7 2022 at Rennes University, entitled *Automated Verification and Synthesis of Distributed Systems*.
 - a reviewer of the PhD thesis of SARARiva November 21st, 2022, at University of Nice, on *Factorization of Discrete Dynamical System*.

11.2.4 Internal or external Inria responsibilities

- Stefan Haar participates in the local evaluation committee for associated teams at INRIA Saclay.

12 Scientific production

12.1 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>.

12.2 Publications of the year

International journals

- [9] T. Chatain, M. Comlan, D. Delfieu, L. Jezequel and O. H. Roux. ‘Pomset bisimulation and unfolding for reset Petri nets’. In: *Information and Computation* 283 (Feb. 2022), p. 104674. DOI: [10.1016/j.ic.2020.104674](https://doi.org/10.1016/j.ic.2020.104674). URL: <https://hal.archives-ouvertes.fr/hal-03650582>.
- [10] 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* 41.8 (2022), pp. 2518–2531. DOI: [10.1109/TCAD.2021.3097599](https://doi.org/10.1109/TCAD.2021.3097599). URL: <https://hal.archives-ouvertes.fr/hal-03536906>.
- [11] I. Khmelnitsky, D. Neider, R. Roy, X. Xie, B. Barbot, B. Bollig, A. Finkel, S. Haddad, M. Leucker and L. Ye. ‘Analysis of Recurrent Neural Networks via Property-Directed Verification of Surrogate Models’. In: *International Journal on Software Tools for Technology Transfer* (2022). URL: <https://hal.archives-ouvertes.fr/hal-03794442>.

International peer-reviewed conferences

- [12] G. K. Aguirre-Samboní, C. Gauchere, S. Haar and F. Pommereau. ‘Reset Petri Net Unfolding Semantics for Ecosystem Hypergraphs’. In: International Workshop on Petri Nets and Software Engineering (PNSE 2022). Vol. 3170. CEUR Workshop Proceedings. Bergen, Norway, 20th July 2022, pp. 213–214. URL: <https://hal-univ-evry.archives-ouvertes.fr/hal-03750233>.
- [13] T. Chatain and N. Rino. ‘Timed Alignments’. In: 2022 4th International Conference on Process Mining (ICPM). Bolzano, France: IEEE, 23rd Oct. 2022, pp. 112–119. DOI: [10.1109/ICPM57379.2022.9980687](https://doi.org/10.1109/ICPM57379.2022.9980687). URL: <https://hal.science/hal-03703712>.
- [14] F. Cravo, M. Függer, T. Nowak and G. Prakash. ‘MobsPy: A Meta-species Language for Chemical Reaction Networks’. In: CMSB 2022 - International Conference on Computational Methods in Systems Biology. Vol. LNBI-13447. Computational Methods in Systems Biology. Bucharest, Romania: Springer International Publishing, 19th Aug. 2022, pp. 277–285. DOI: [10.1007/978-3-031-15034-0_14](https://doi.org/10.1007/978-3-031-15034-0_14). URL: <https://hal.science/hal-03876041>.
- [15] L. He, P. Dague and L. Ye. ‘Using Delay Blocks to Make Non-Diagnosable Discrete Event Systems Diagnosable’. In: 33rd International Workshop on Principle of Diagnosis – DX 2022. Toulouse, France, 14th Sept. 2022. URL: <https://hal.science/hal-03773712>.
- [16] G. Karlo Aguirre-Samboní, S. Haar, L. Paulevé, S. Schwoon and N. Würdemann. ‘Avoid One’s Doom: Finding Cliff-Edge Configurations in Petri Nets’. In: GandALF 2022 - Games, Automata, Logics, and Formal Verification. Madrid, Spain, 21st Sept. 2022. DOI: [10.4204/EPTCS.370.12](https://doi.org/10.4204/EPTCS.370.12). URL: <https://hal.inria.fr/hal-03773855>.
- [17] I. Khmelnitsky, S. Haddad, L. Ye, B. Barbot, B. Bollig, M. Leucker, D. Neider and R. Roy. ‘Analyzing Robustness of Angluin’s L* Algorithm in Presence of Noise’. In: GandALF 2022 - 13th International Symposium on Games, Automata, Logics and Formal Verification. Vol. 370. Madrid, Spain, 20th Sept. 2022, pp. 81–96. DOI: [10.4204/EPTCS.370.6](https://doi.org/10.4204/EPTCS.370.6). URL: <https://hal.science/hal-03794320>.

Scientific book chapters

- [18] M. Függer, C. Lenzen and U. Schmid. ‘On Specifications and Proofs of Timed Circuits’. In: *Principles of Systems Design*. Vol. 13660. Lecture Notes in Computer Science. Springer Nature Switzerland, 29th Dec. 2022, pp. 107–130. DOI: [10.1007/978-3-031-22337-2_6](https://doi.org/10.1007/978-3-031-22337-2_6). URL: <https://hal.science/hal-03935679>.

Reports & preprints

- [19] N. Rino and T. Chatain. *Timed Alignments with Mixed Moves*. 26th Oct. 2022. URL: <https://hal.inria.fr/hal-03831470>.

12.3 Cited publications

- [20] 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>.
- [21] R. Alur, K. Etessami and M. Yannakakis. ‘Realizability and Verification of MSC Graphs’. In: *Theor. Comput. Sci.* 331.1 (2005), pp. 97–114.
- [22] 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-icompi0.pdf>.
- [23] 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>.
- [24] 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>.
- [25] 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>.
- [26] 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>.
- [27] 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.
- [28] J. Esparza and K. Heljanko. *Unfoldings - A Partial-Order Approach to Model Checking*. EATCS Monographs in Theoretical Computer Science. Springer, 2008.
- [29] 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.
- [30] 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.
- [31] 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.

- [32] 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.
- [33] 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.
- [34] 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.
- [35] 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>.
- [36] 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>.