RESEARCH CENTRE

Inria Saclay Centre at Institut Polytechnique de Paris

IN PARTNERSHIP WITH:

CNRS, Université Paris-Saclay

2023 ACTIVITY REPORT

Team MEXICO

Modeling and Exploitation of Interaction and Concurrency

Inria teams are typically groups of researchers working on the definition of a common project, and objectives, with the goal to arrive at the creation of a project-team. Such project-teams may include other partners (universities or research institutions)

IN COLLABORATION WITH: Laboratoire de Méthodes Formelles

DOMAIN

Algorithmics, Programming, Software and Architecture

THEME

Proofs and Verification

Contents

Team MEXICO

Creation of the Team: 2023 June 30

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]
- Matthias Fuegger [CNRS, Researcher, until Jun 2023]
- Stefan Haar [INRIA, Senior Researcher, from Jul 2023, HDR]

Faculty Members

- Thomas Chatain [ENS PARIS-SACLAY, Associate Professor, until Jun 2023]
- Serge Haddad [ENS PARIS-SACLAY, Professor, until Jun 2023, HDR]
- Stefan Schwoon [ENS PARIS-SACLAY, Associate Professor, until Jun 2023, HDR]

PhD Students

- Giann Karlo Aguirre Samboni [INRIA, Thesis defence Dec 14, 2023]
- Fabricio Cravo [CNRS]
- Mélanie PIETRI [CNRS]
- Zhuofan Xu [CNRS]

External Collaborator

- Benoît Barbot [UNIV PARIS EST, until Jun 2023]

2 Overall objectives

Please note that this document reports on part of 2023 only; in fact, after more than twelve years of existence, the MEXICO team has been terPminated on June 30, 2023.

Introduction. The general objective of the MEXICO team is the analysis and control of complex, concurrent and composite systems, using intensively our expertise in formal models for dynamical systems.

We have chosen to modify the detailed presentation of objectives and domains, comparted to the way they were presented in previous years, to account for the evolution of our work over the past years.

2.1 Overall objectives, Evolution and Perspectives

The goals of the MEXICO team have been structured around three main transversal (and fundamental) *objectives*, and four application domains. The main objectives have been

1. Model and exploit *concurrency*, *asynchrony* and *distribution*,
2. address the challenges posed by the *interaction* of diverse and semi-transparent components, and
3. manage *quantitative* aspects of behaviour;

these objectives intersected and transversed the application domains of

1. telecommunications,
2. biological regulation networks,
3. metabolic networks and

4. transportation systems.

Globally, the fundamental objectives still hold unchanged, and have manifested themselves in our activities and results. It should be noted that the intention behind the objective 'interaction' had been mainly to to address combinations of web services or similar, *software*-related entities from telecommunications; by contrast, the current realization of the interaction objective has been, over the past years, the dynamic networks and distributed algorithms, with the focus on *hardware* and *'wetware'*.

This phenomenon is not limited to one objective,but indicative of a more general, fundamental shift in our application domains:

- The telecommunications domain had been, until about the 2010s, a rather favorable ground for our activities of axis 1 , plus our work on partial order semantics. After that, the focus in the application area changed, and we were not able to secure industrial partnerships in telecommunication.

- The field of transportation systems had still been active in the previous reporting period, with collaborative projects on multimodel transport networks on the one hand, and autonomous vehicles at the other. This was the reason that this application domain was part of our intended targets at the time of the previous evaluation, but the submissions of follow-up proposals were unsuccesful, and the partnerships dissolved; were the MEXICO team to be renewed, this domain would no longer be part of its objectives.

- The field of biological networks, on the other hand, is still very active, while undergoing several modifications. First of all, the work on *metabolic pathways* in MEXICO ended with the end of the temporary stay of Philippe Dague in the team. Then, the subject of *cellular reprogramming* had been carried by the ANR-FNR project ALGORECELL, with partners at Institut Pasteur and University of Luxemburg ; but here again, a follow-up project submission was unsuccessful, and the cooperation with these partners has continued without support. The research on the scientific issues raised by cellular reprogramming is continuing and will be pursued under different auspices, see below.

- By contrast, the methodology of modeling biological networks by discrete Petri nets and extracting behavioural information via unfoldings of these models, is being successfully transferred, adapted and extended in the modelling and analysis of *ecosystems.* The project *EcoSystem Causal Analysis with Petri nEt unfoldings (ESCAPE)*, with Université Evry and INRAE Montpellier, is terminating end of 2023, with the expected completion and defence of Giann Karlo Aguirre Samboní's PhD thesis. Another project, SMART, jointly funded by INRAE and INRIA, is under way in cooperation with the RECOVER team at INRAE Aix-en-Provence, on the subject of harnessing environmental *multi-risks*.

- The fields of asynchronous circuits and distributed algorithms for Microbiological Systems is currently in a very active phase, and its amount and intensity of activity have greatly increased, not the least due to the ANR project DREAMY one of whose principal investigators is Matthias Függer (the other being Thomas Nowak, also at LMF). Matthias is also preparing a new biological lab to be hosted at ENS Paris-Saclay to be opened early in 2024.

Overall, and adding the continuing activity on *process mining*, we thus have a fundamental change in objectives, reflected by the five axes detailed below. These axes will continue in different forms, owed to the termination of the team. 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, and the study of process alignments.

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 [8] 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 $\mathscr{N}$ and an observed alarm pattern $\mathscr{A}$, 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* $\mathscr{A}$. The Petri-net based approach generalizes to dynamically evolving topologies, in dynamical systems modeled by graph grammars, see [21].

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 [28, 31], elements of a methodology for distributed computation of unfoldings between several supervisors, underwritten by algebraic properties of the category of Petri nets have been developed. Generalizations, in particular to Graph Grammars, are still do be done.

Computing diagnosis in a distributed way is only one aspect of a much vaster topic, that of *distributed diagnosis* (see [26, 29]). 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 [20, 23]). 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 [32, 33]. Automatizing such procedures for the supervision and management of distributed and locally monitored asynchronous systems is a long-term goal to which *MExICo* hopes to contribute.

Hybrid Systems

Participants: Philippe Dague, Lina Ye, Serge Haddad.

Hybrid systems constitute a model for cyber-physical systems which integrates continuous-time dynamics (modes) governed by differential equations, and discrete transitions which switch instantaneously from one mode to another. Thanks to their ease of programming, hybrid systems have been integrated to power electronics systems, and more generally in cyber-physical systems. In order to guarantee that such systems meet their specifications, classical methods consist in finitely abstracting the systems by discretization of the (infinite) state space, and deriving automatically the appropriate mode control from the specification using standard graph techniques.

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

Contextual Nets

Participants: Stefan Schwoon.

Assuring the correctness of concurrent systems is notoriously difficult due to the many unforeseeable ways in which the components may interact and the resulting state-space explosion. A well-established approach to alleviate this problem is to model concurrent systems as Petri nets and analyse their unfoldings, essentially an acyclic version of the Petri net whose simpler structure permits easier analysis [27].

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 [34] 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 [19]; 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 [22].
- 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 [25].
- 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 [35] .

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 [10]); (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 [30] [24]. 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 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 differenciate ? Do mutations occur ? etc), and progress there hinges on our capacity to analyse, predict, monitor and control complex and variegated processes. We have applied Petri net unfolding techniques for the efficient computation of attractors in a regulatory network; that is, to identify strongly connected reachability components that correspond to stable evolutions, e.g. of a cell that differentiates into a specific functionality (or mutation). This constitutes the starting point of a broader research with Petri net unfolding techniques in regulation. In fact, the use of ordinary Petri nets for capturing regulatory network (RN) dynamics overcomes the limitations of traditional RN models : those impose e.g. Monotonicity properties in the influence that one factor had upon another, i.e. always increasing or always decreasing, and were thus unable to cover all actual behaviours. Rather, we follow the more refined model of boolean networks of automata, where the local states of the different factors jointly detemine which state transitions are possible. For these connectors, ordinary PNs constitute a first approximation, improving greatly over the literature but leaving room for improvement in terms of introducing more refined logical connectors. Future work thus involves transcending this class of PN models. Via unfoldings, one has access – provided efficient techniques are available – to all behaviours of the model, rather than over-or under-approximations as previously. This opens the way to efficiently searching in particular for determinants of the cell fate : which attractors are reachable from a given stage, and what are the factors that decide in favor of one or the other attractor, etc. Our current research focusses cellular reprogramming on the one hand, and distributed algorithms in wild or synthetic biological systems on the other. The latter is a distributed algorithms' view on microbiological systems, both with the goal to model and analyze existing microbiological systems as distributed systems, and to design and implement distributed algorithms in synthesized microbiological systems. Envisioned major long-term goals are drug production and medical treatment via synthesized bacterial colonies. We are approaching our goal of a distributed algorithm's view of microbiological systems from several directions: (i) Timing plays a crucial role in microbiological systems. Similar to modern VLSI circuits, dominating loading effects and noise render classical delay models unfeasible. In previous work we showed limitations of current delay models and presented a class of new delay models, so called involution channels. In [26] we showed that involution channels are still in accordance with Newtonian physics, even in presence of noise. (ii) In [7] we analyzed metastability in circuits by a three-valued Kleene logic, presented a general technique to build circuits that can tolerate a certain degree of metastability at its inputs, and showed the presence of a computational hierarchy. Again, we expect metastability to play a crucial role in microbiological systems, as similar to modern VLSI circuits, loading effects are pronounced. (iii) We studied agreement problems in highly dynamic networks without stability guarantees [28], [27]. We expect such networks to occur in bacterial cultures where bacteria communicate by producing and sensing small signal molecules like AHL. Both works also have theoretically relevant implications: The work in [27] presents the first approximate agreement protocol in a multidimensional space with time complexity independent of the dimension, working also in presence of Byzantine faults. In [28] we proved a tight lower bound on convergence rates and time complexity of asymptotic and approximate agreement in dynamic and classical static fault models. (iv) We are currently working with Manish Kushwaha (INRA), and Thomas Nowak (LRI) on biological infection models for E. coli colonies and M13 phages.

In the context of the ESCAPE project (PhD thesis of G.K. Aguirre Samboni, started in October 2020) we are now extending our research on causal analysis of complex biological networks to the domain of *ecosystems*, in cooperation with INRAE researcher Cédric Gaucherel. The cooperation with INRAE has been intensifiying 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.

4.2 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

Completion of the ECOFOLDER tool for the causal analysis of ecosystems using unfoldings.

Successful completion of the cooperation on unfolding prefixes for high-level Petri nets [17], successfully presented at the 2023 Petri Nets conference.

7 New results

A preliminary comment is in order: several new results have been published in the second half of 2023, and cannot be cited here for this reason; we comment only on those available in the publication list below.

7.1 PALS: Distributed Gradient Clocking on Chip

Participants: Matthias Függer.

Consider an arbitrary network of communicating modules on a chip, each requiring a local signal telling it when to execute a computational step. There are three common solutions to generating such a local clock signal: 1) by deriving it from a single, central clock source; 2) by local, free-running oscillators; or 3) by handshaking between neighboring modules. Conceptually, each of these solutions is the result of a perceived dichotomy in which (sub)systems are either clocked or asynchronous. We present in [13] a solution and its implementation that lies between these extremes. Based on a distributed gradient clock synchronization (GCS) algorithm, we show a novel design providing modules with local clocks, the frequency bounds of which are almost as good as those of free-running oscillators, yet neighboring modules are guaranteed to have a phase offset substantially smaller than one clock cycle. Concretely, parameters obtained from a 15-nm application specific integrated circuit (ASIC) simulation running at 2 GHz yield mathematical worst-case bounds of 20 ps on the phase offset for a 32×32 node grid network.

7.2 Continuity of Thresholded Mode-Switched ODEs and Digital Circuit Delay Models

Participants: Matthias Függer.

Thresholded mode-switched ODEs are restricted dynamical systems that switch ODEs depending on digital input signals only, and produce a digital output signal by thresholding some internal signal. Such systems arise in recent digital circuit delay models, where the analog signals within a gate are governed by ODEs that change depending on the digital inputs. In [14], we prove the continuity of the mapping from digital input signals to digital output signals for a large class of thresholded mode-switched ODEs. This continuity property is known to be instrumental for ensuring the faithfulness of the model w.r.t. propagating short pulses. We apply our result to several instances of such digital delay models, thereby proving them to be faithful.

7.3 On the Susceptibility of QDI Circuits to Transient Faults

Participants: Matthias Függer.

By design, quasi delay-insensitive (QDI) circuits exhibit higher resilience against timing variations as compared to their synchronous counterparts. Since computation in QDI circuits is event-based rather than clock-triggered, spurious events due to transient faults such as radiation-induced glitches, a priori are of higher concern in QDI circuits. In this work we propose a formal framework with the goal to gain a deeper understanding on how susceptible QDI circuits are to transient faults. We introduce in [16] a worst-case model for transients in circuits. We then prove an equivalence of faults within this framework and use this result to provably exhaustively check a widely used QDI circuit, a linear Muller pipeline, for its susceptibility to produce non-stable output signals.

7.4 Taking Complete Finite Prefixes To High Level, Symbolically

Participants: Thomas Chatain, Stefan Haar.

Unfoldings are a well known partial-order semantics of P/T Petri nets that can be applied to various model checking or verification problems. For high-level Petri nets, the so-called symbolic unfolding generalizes this notion. A complete finite prefix of the unfolding of a P/T Petri net contains all information to verify, e.g., reachability of markings. In [17], we unite these two concepts and define complete finite prefixes of the symbolic unfolding of high-level Petri nets. For a class of safe high-level Petri nets, we generalize the well-known algorithm by Esparza et al. for constructing small such prefixes. Additionally, we identify a more general class of nets with infinitely many reachable markings, for which an approach with an adapted cutoff criterion extends the complete prefix methodology, in the sense that the original algorithm cannot be applied to the P/T net represented by a high-level net.

7.5 Introducing Divergence for Infinite Probabilistic Modelseural Networks via Property-Directed Verification of Surrogate Models.

Participants: Serge Haddad, Lina Ye, Benoît Barbot.

Computing the reachability probability in infinite state probabilistic models has been the topic of numerous works. In [15], we introduce a new property called divergence that when satisfied allows to compute reachability probabilities up to an arbitrary precision. One of the main interest of divergence is that our algorithm does not require the reachability problem to be decidable. Then we study the decidability of divergence for probabilistic versions of pushdown automata and Petri nets where the weights associated with transitions may also depend on the current state. This should be contrasted with most of the existing works that assume weights independent of the state. Such an extended framework is motivated by the modeling of real case studies. Moreover, we exhibit some divergent subclasses of channel systems and pushdown automata, particularly suited for specifying open distributed systems and networks prone to performance collapsing in order to compute the probabilities related to service requirements.

8 Partnerships and cooperations

Participants: Serge Haddad, .

8.1 National initiatives

Dreamy (354 kEur): A key advantage of biological computing devices is their ability to sense, compute, and especially to respond to their biological environment, e.g., bacteria can be programmed to act as autonomous robots within the human body. Local presence of certain molecules in the environment allows sensing of neighboring cell types and acting accordingly, e.g., by activating an immune response. Current designs of synthetic circuits in bacteria, however, face severe resource limitations: each genetic part added to the cell imposes an additional burden, becoming progressively toxic for the cell. The most common design techniques for biological logic gates rely on gene regulation via DNA-binding proteins, nucleic acid (DNA/RNA) interactions, or more recently the CRISPR machinery. Each comes with its own constraints: like limited availability of orthogonal signals for use within the cell (DNA-binding), small dynamic range (RNA-based), or reduced growth rates (the CRISPR machinery). This has led to recent efforts to distribute circuits among several cells to reduce the resource load per cell, taking the formative steps towards distributed bacterial circuits. The DREAMY research project seeks to develop innovative solutions to the problem of building distributed circuits in bacteria from an algorithmic, theoretical perspective that contributes to real-world implementable solutions.

Involved groups: LMF (FR), LISN (FR), DISCO (FR), ALGO group (University of Geneva, CH), Micalis (INRAE, FR), L2S (FR)

ANR MAVeriQ (ANR-20-CE25-0012), 2021-2025.
Partners:

- IRIF (Université de Paris - CNRS)
- INRIA Rennes
- LACL (Université Paris Est Créteil)
- VERIMAG (Université Grenoble Alpes - CNRS)

Participant: Serge Haddad.
MAVEriQ stands for "Methods of Analysis for Verification of Quantitative properties". Its goal is to promote unified methods for the quantitative verification of timed, stochastic and/or hybrid systems.

Other funding

Funding for Giann Karlo Aguirre Samboni's PhD thesis comes from the DIGICOSME grant ESCAPE, DIGICOSME RD 242-ESCAPE-15203; participants: Giann Karlo Aguirre Samboni, Franck Pommereau, Stefan Haar.

NERF project on *Formal representation of Nucleotide Excision Repair pathway*, jointly with CEA team Genome transcriptional regulation; 13KEur, funded by the *interdisciplinary object (OI) LivingMachines@Work*, Université Paris-Saclay; participants in MEXICO: Thomas Chatain, Serge Haddad and Stefan Haar.

PROCEED [2023]: AAP Preuve de Concept of OI BioProbe (ENS Paris-Saclay).

9 Dissemination

Participants: All Members.

9.1 Conference Committees, Responsibilities in the Scientific Community

THOMAS CHATAIN

- Member of the jury for *SIF-Gilles Kahn award* for outstainding PhD theses
- Head of *Interactions* pole at LMF lab
- PC member of the *International Conference on Process Mining (ICPM)* since 2019

MATTHIAS FÜGGER:

- Topic chair for Emerging technologies, AI application for HW design & test, new computing paradigms at IEEE DDECS 2024
- Co-founded Workshop on Computing among Cells (CELLS 2019) at DISC'19

STEFAN HAAR

- Adjoint Director for Research of the Université Paris-Saclay's *Graduate School Computer Science*
- Associate editor for *Journal of Discrete Events Dynamic Systems: Theory and Application* (JDEDS).
- Workshops co-chair for ETAPS 2023
- Member of the of the program committee for the *International Conference on Application and Theory of Petri Nets and Concurrency* 2023.

SERGE HADDAD

- Member of the scientific and administrative council (CSA) of Labex CIMI of Toulouse and a member of the scientific orientation council (COS) of LIS of Marseille (UMR 7020).
- Representative of ENS Paris-Saclay to the supervisory board ('conseil') of the Graduate School Computer Science, Université Paris-Saclay
- Member of the steering committee of the International Conference on Application and Theory of Petri Nets and Concurrency.

STEFAN SCHWOON

- Participation in the organisation of *QONFEST '21*, in particular as webmaster

9.2 Teaching - Supervision - Juries

9.2.1 Teaching

- SERGE HADDAD is a full-time professor, and both THOMAS CHATAIN and STEFAN SCHWOON are full-time associate professors, all in the Computer Science Department at ENS Paris-Saclay. In addition, THOMAS CHATAIN is the head of the *M1* and STEFAN SCHWOON of the *L3* program in Computer Science at at ENS Paris-Saclay.
- MATTHIAS FÜGGER
  - Fall 2023: Algorithmics and Bioinformatics (Függer, Haar, Nowak), 4h per week, half semester
  - Fall 2023: Initiation to Research: Synthetic Biology (Függer, Nowak), 2h per week, semester
- STEFAN HAAR has taught *Analysis of dynamics in biological networks* in the *Bioinformatics* program of *Université Paris-Saclay*, at the M2 level, for 24 h of teaching, 4h per week, during a half semester; and *Algorithmics and Bio-Informatics* at the M1 level of the MPRI (Parisian Master of research in Computer Science) program at ENS Paris-Saclay, 4h per week, half semester (for the first time in Fall 2023.)

9.2.2 Supervision

PhD Supervision:

- MATTHIAS FÜGGER:
  - Melanie Pietri (PhD), co-supervision with Bruno Le Pioufle (ENS Paris-Saclay), Sakina Chantoiseau-Bensalem (ENS Paris-Saclay), and Thomas Nowak (ENS Paris-Saclay) since 2023
  - Zhuofan Xu (PhD), co-supervision with Benedikt Bollig (ENS Paris-Saclay) and Thomas Nowak (ENS Paris-Saclay) since 2022
  - Fabricio Cravo (PhD), co-supervision with Janna Burman (Université Paris-Saclay) and Thomas Nowak (ENS Paris-Saclay) since 2021
  - Raghda Elshehaby (PhD), 6 month co-supervision of visiting PhD student of Andreas Steininger (TU Wien) 2022 - 2023
- STEFAN HAAR:
  - Giann Karlo Aguirre Samboni; co-supervisor: Franck Pommereau (Université d'Evry), 2020-2023 (defence on Dec 14, 2023)
  - Souhila FOUNAS (at INRAE Aix-en-Provence); main supervisor: Corinne Curt (RECOVER team, INRAE), since Dec. 2022

9.2.3 Juries

PhD committees:

THOMAS CHATAIN

- Thomas Chatain was a reviewer for the PhD thesis of Federica Adobbati (director Luca Brenardinello) at University of Milano Bicocca (Italy) in March 2023.

STEFAN HAAR

- *Garant* and jury member for the *HDR* of Paolo Ballarini, Université Paris-Saclay, defended 2023
- President of the PhD Committee for the PhD thesis of Rūta Binkytė-Sadauskienė, École Polytechnique, 2023

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.

[2] T. Chatain, M. Boltenhagen and J. Carmona. 'Anti-Alignments – Measuring The Precision of Process Models and Event Logs'. In: *Information Systems* (May 2021). DOI: 10.1016/j.is.2020.101708. URL: https://inria.hal.science/hal-02383546.

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

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

[5] D.-J. Cho, M. Függer, C. Hopper, M. Kushwaha, T. Nowak and Q. Soubeyran. 'Distributed computation with continual population growth'. In: *Distributed Computing* (Oct. 2021). DOI: 10.1007/s00446-021-00404-8. URL: https://inria.hal.science/hal-03519002.

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

[7] 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. URL: https://hal.science/hal-03408731.

[8] 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.

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

[10] 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.

[11] 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.

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

10.2 Publications of the year

International journals

[13] J. Bund, M. Függer and M. Medina. 'PALS: Distributed Gradient Clocking on Chip'. In: *IEEE Transactions on Very Large Scale Integration (VLSI) Systems* 31.11 (Nov. 2023), pp. 1740–1753. DOI: 10.1109/TVLSI.2023.3311178. URL: https://hal.science/hal-04266029.

International peer-reviewed conferences

[14] A. Ferdowsi, M. Függer, T. Nowak and U. Schmid. 'Continuity of Thresholded Mode-Switched ODEs and Digital Circuit Delay Models'. In: 26th ACM International Conference on Hybrid Systems: Computation and Control (HSCC). San Antonio, United States, 9th May 2023. URL: https://hal.science/hal-03973076.

[15] A. Finkel, S. Haddad and L. Ye. 'Introducing Divergence for Infinite Probabilistic Models'. In: *Reachability Problems*. RP 2023 - 17th International Conference on Reachability Problems. Vol. LNCS-14235. Lecture Notes in Computer Science. Nice, France: Springer Nature Switzerland, 5th Oct. 2023, pp. 127–140. DOI: 10.1007/978-3-031-45286-4_10. URL: https://hal.science/hal-04271224.

[16] R. E. Shehaby, M. Függer and A. Steininger. 'On the Susceptibility of QDI Circuits to Transient Faults'. In: *Lecture Notes in Computer Science*. FORMATS 2023 - International Conference on Formal Modeling and Analysis of Timed Systems. Vol. LNCS-14138. Formal Modeling and Analysis of Timed Systems. Antwerp, Belgium: Springer Nature Switzerland, 29th Aug. 2023, pp. 69–85. DOI: 10.1007/978-3-031-42626-1_5. URL: https://hal.science/hal-04266002.

[17] N. Würdemann, T. Chatain and S. Haar. 'Taking Complete Finite Prefixes To High Level, Symbolically (Full Version)'. In: *LNCS*. Petri Nets 2023 - 44TH International Conference on Applications and Theory of Petri Nets and Concurrency. Caparica (Lisbonne), Portugal: Springer, 25th Mar. 2023. URL: https://inria.hal.science/hal-04029490.

Reports & preprints

[18] N. Würdemann, T. Chatain, S. Haar and L. Panneke. *Taking Complete Finite Prefixes To High Level, Symbolically*. 2023. DOI: 10.48550/arXiv.2311.11443. URL: https://inria.hal.science/hal-04297773.

10.3 Cited publications

[19] 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.

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

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

[22] 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.

[23] P. Bhateja, P. Gastin, M. Mukund and K. Narayan Kumar. 'Local testing of message sequence charts is difficult'. In: *Proceedings of the 16th International Symposium on Fundamentals of Computation Theory (FCT'07)*. Ed. by E. Csuhaj-Varjú and Z. Ésik. Vol. 4639. Lecture Notes in Computer Science. Budapest, Hungary: Springer, Aug. 2007, pp. 76–87. DOI: 10.1007/978-3-540-74240-1_8. URL: http://www.lsv.ens-cachan.fr/Publis/PAPERS/PDF/BGMN-fct07.pdf.

[24] 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.

[25] 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.

[26] 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.

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

[28] 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.

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

[30] 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.

[31] 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.

[32] 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.

[33] 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.

[34] 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. URL: https://hal.archives-ouvertes.fr/hal-00926202.

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