IN PARTNERSHIP WITH:
CNRS

Ecole normale supérieure de Cachan

Activity Report 2018

Project-Team MEXICO

Modeling and Exploitation of Interaction and Concurrency

IN COLLABORATION WITH: Laboratoire specification et vérification (LSV)

RESEARCH CENTER
Saclay - Île-de-France

THEME
Proofs and Verification

Table of contents

Project-Team MEXICO

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

Keywords:

Computer Science and Digital Science:

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
B1.1.10. - Systems and synthetic biology
B6.3.1. - Web
B6.3.2. - Network protocols
B6.3.3. - Network Management
B7.1. - Traffic management
B7.2.1. - Smart vehicles

1. Team, Visitors, External Collaborators

Research Scientists

Stefan Haar [Team leader, Inria, Senior Researcher, HDR]
Laurent Fribourg [CNRS, Researcher, HDR]
Matthias Függer [CNRS, Researcher]

Faculty Members

Thomas Chatain [Ecole Normale Supérieure Cachan, Associate Professor]
Philippe Dague [Univ Paris-Sud, Professor, from Oct 2018]
Serge Haddad [Ecole Normale Supérieure Cachan, Professor, HDR]
Claudine Picaronny [Ecole Normale Supérieure Cachan, Professor, until Jul 2018]
Stefan Schwoon [Ecole Normale Supérieure Cachan, Associate Professor]

Post-Doctoral Fellows
Da Jung Cho [CNRS, from May 2018]
Daniele Zonetti [CNRS, from May 2018]

PhD Students
Mathilde Boltenhagen [CNRS, from Nov 2018]
Yann Duplouy [Institut de recherche technologique System X, until Nov 2018]
Jawher Jerray [Univ Paris-Nord, from Oct 2018]
Igor Khmelnitsky [Ecole Normale Supérieure Cachan, from Nov 2018]
Juraj Kolcˇák [Inria]
Hugues Mandon [Inria]
Adnane Saoud [CNRS]

Administrative Assistants
Adeline Lochet [Inria, from Jun 2018 until Nov 2018]
Emmanuelle Perrot [Inria]

External Collaborator
Benoît Barbot [Univ Paris-Est Marne La Vallée]

2. Overall Objectives

2.1. Scientific Objectives

2.1.1. Introduction.

In the increasingly networked world, reliability of applications becomes ever more critical as the number of users of, e.g., communication systems, web services, transportation etc., grows steadily. Management of networked systems, in a very general sense of the term, therefore is a crucial task, but also a difficult one.

MExICo strives to take advantage of distribution by orchestrating cooperation between different agents that observe local subsystems, and interact in a localized fashion.

The need for applying formal methods in the analysis and management of complex systems has long been recognized. It is with much less unanimity that the scientific community embraces methods based on asynchronous and distributed models. Centralized and sequential modeling still prevails.

However, we observe that crucial applications have increasing numbers of users, that networks providing services grow fast both in the number of participants and the physical size and degree of spatial distribution. Moreover, traditional *isolated* and *proprietary* software products for local systems are no longer typical for emerging applications.

In contrast to traditional centralized and sequential machinery for which purely functional specifications are efficient, we have to account for applications being provided from diverse and non-coordinated sources. Their distribution (e.g. over the Web) must change the way we verify and manage them. In particular, one cannot ignore the impact of quantitative features such as delays or failure likelihoods on the functionalities of composite services in distributed systems.

We thus identify three main characteristics of complex distributed systems that constitute research challenges:

- *Concurrency* of behavior;
- *Interaction* of diverse and semi-transparent components; and
- management of *Quantitative* aspects of behavior.

2.1.2. Concurrency

The increasing size and the networked nature of communication systems, controls, distributed services, etc. confront us with an ever higher degree of parallelism between local processes. This field of application for our work includes telecommunication systems and composite web services. The challenge is to provide sound theoretical foundations and efficient algorithms for management of such systems, ranging from controller synthesis and fault diagnosis to integration and adaptation. While these tasks have received considerable attention in the *sequential* setting, managing *non-sequential* behavior requires profound modifications for existing approaches, and often the development of new approaches altogether. We see concurrency in distributed systems as an opportunity rather than a nuisance. Our goal is to *exploit* asynchronicity and distribution as an advantage. Clever use of adequate models, in particular *partial order semantics* (ranging from Mazurkiewicz traces to event structures to MSCs) actually helps in practice. In fact, the partial order vision allows us to make causal precedence relations explicit, and to perform diagnosis and test for the dependency between events. This is a conceptual advantage that interleaving-based approaches cannot match. The two key features of our work will be *(i)* the exploitation of concurrency by using asynchronous models with partial order semantics, and *(ii)* distribution of the agents performing management tasks.

2.1.3. Interaction

Systems and services exhibit non-trivial *interaction* between specialized and heterogeneous components. A coordinated interplay of several components is required; this is challenging since each of them has only a limited, partial view of the system's configuration. We refer to this problem as *distributed synthesis* or *distributed control*. An aggravating factor is that the structure of a component might be semi-transparent, which requires a form of *grey box management*.

2.1.4. Quantitative Features

Besides the logical functionalities of programs, the *quantitative* aspects of component behavior and interaction play an increasingly important role.

- *Real-time* properties cannot be neglected even if time is not an explicit functional issue, since transmission delays, parallelism, etc, can lead to time-outs striking, and thus change even the logical course of processes. Again, this phenomenon arises in telecommunications and web services, but also in transport systems.
- In the same contexts, *probabilities* need to be taken into account, for many diverse reasons such as unpredictable functionalities, or because the outcome of a computation may be governed by race conditions.
- Last but not least, constraints on *cost* cannot be ignored, be it in terms of money or any other limited resource, such as memory space or available CPU time.

2.1.5. Evolution and Perspectives

Since the creation of *MExICo*, the weight of *quantitative* aspects in all parts of our activities has grown, be it in terms of the models considered (weighted automata and logics), be it in transforming verification or diagnosis verdict into probabilistic statements (probabilistic diagnosis, statistical model checking), or within the recently started SystemX cooperation on supervision in multi-modal transport systems. This trend is certain to continue over the next couple of years, along with the growing importance of diagnosis and control issues.

In another development, the theory and use of partial order semantics has gained momentum in the past four years, and we intend to further strengthen our efforts and contacts in this domain to further develop and apply partial-order based deduction methods.

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

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

1. **(Re-)programming in discrete concurrent models.** Cellular regulatory networks exhibit highly complex concurrent behaviours that is influenced by a high number of perturbations such as mutations. We are in particular investigating discrete models, both in the form of boolean networks and of Petri nets, to harness this complexity, and to obtain viable methods for two interconnected and central challenges:
   - find *attractors*, i.e. long-run stable states or sets of states, that indicate possible phenotypes of the organism under study, and
   - determine *reprogramming* strategies that apply perturbations in such a way as to steer the cell's long-run behaviour into some desired phenotype, or away from an undesired one.
2. **Distributed Algorithms in wild or synthetic biological systems.** Since the arrival of Matthias Függer in the team, we also work, on the multi-cell level, with a distributed algorithms' view on microbiological systems, both with the goal to model and analyze existing microbiological systems as distributed systems, and to design and implement distributed algorithms in synthesized microbiological systems. Major long-term goals are drug production and medical treatment via synthesized bacterial colonies.

3. Research Program

3.1. Concurrency

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

Concurrency: Property of systems allowing some interacting processes to be executed in parallel.

Diagnosis: The process of deducing from a partial observation of a system aspects of the internal states or events of that system; in particular, *fault diagnosis* aims at determining whether or not some non-observable fault event has occurred.

Conformance Testing: Feeding dedicated input into an implemented system IS and deducing, from the resulting output of I, whether I respects a formal specification S.

3.1.1. Introduction

It is well known that, whatever the intended form of analysis or control, a *global* view of the system state leads to overwhelming numbers of states and transitions, thus slowing down algorithms that need to explore the state space. Worse yet, it often blurs the mechanics that are at work rather than exhibiting them. Conversely, respecting concurrency relations avoids exhaustive enumeration of interleavings. It allows us to focus on 'essential' properties of non-sequential processes, which are expressible with causal precedence relations. These precedence relations are usually called causal (partial) orders. Concurrency is the explicit absence of such a precedence between actions that do not have to wait for one another. Both causal orders and concurrency are in fact essential elements of a specification. This is especially true when the specification is constructed in a distributed and modular way. Making these ordering relations explicit requires to leave the framework of state/interleaving based semantics. Therefore, we need to develop new dedicated algorithms for tasks such as conformance testing, fault diagnosis, or control for distributed discrete systems. Existing solutions for these problems often rely on centralized sequential models which do not scale up well.

3.1.2. Diagnosis

Participants: Stefan Haar, Serge Haddad, Stefan Schwoon.

Fault Diagnosis for discrete event systems is a crucial task in automatic control. Our focus is on *event oriented* (as opposed to *state oriented*) model-based diagnosis, asking e.g. the following questions:

given a - potentially large - *alarm pattern* formed of observations,

- what are the possible *fault scenarios* in the system that *explain* the pattern ?
- Based on the observations, can we deduce whether or not a certain - invisible - fault has actually occurred ?

Model-based diagnosis starts from a discrete event model of the observed system - or rather, its relevant aspects, such as possible fault propagations, abstracting away other dimensions. From this model, an extraction or unfolding process, guided by the observation, produces recursively the explanation candidates.

In asynchronous partial-order based diagnosis with Petri nets [49], [50], [51], 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 [38]

3.1.2.1. Observability and Diagnosability

Diagnosis algorithms have to operate in contexts with low observability, i.e., in systems where many events are invisible to the supervisor. Checking *observability* and *diagnosability* for the supervised systems is therefore a crucial and non-trivial task in its own right. Analysis of the relational structure of occurrence nets allows us to check whether the system exhibits sufficient visibility to allow diagnosis. Developing efficient methods for both verification of *diagnosability checking* under concurrency, and the *diagnosis* itself for distributed, composite and asynchronous systems, is an important field for *MExICo*.

3.1.2.2. Distribution

Distributed computation of unfoldings allows one to factor the unfolding of the global system into smaller *local* unfoldings, by local supervisors associated with sub-networks and communicating among each other. In [50], [40], 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 [47], [53]). 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 [37], [43]). 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 [55], [56]. Automatizing such procedures for the supervision and management of distributed and locally monitored asynchronous systems is a long-term goal to which *MExICo* hopes to contribute.

3.1.3. Hybrid Systems

Participants: Laurent Fribourg, 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. These methods face the well-known problem of "curse of dimensionality", and cannot generally treat systems of dimension exceeding 5 or 6. Thanks to the introduction of original compositional techniques [25], [30], [13] as well as finer estimations of integration errors [3], we are now able to control several case studies of greater dimension. Actually, in the real world, many parameters of hybrid models are not known precisely, and require adjustements to experimental data. We plan to elaborate methods based on parameter estimation and machine learning techniques in order to define formal stability criteria and well-posed learning problems in the framework of hybrid systems with nonlinear dynamics.

3.1.4. Contextual Nets

Participant: Stefan Schwoon.

Assuring the correctness of concurrent systems is notoriously difficult due to the many unforeseeable ways in which the components may interact and the resulting state-space explosion. A well-established approach to alleviate this problem is to model concurrent systems as Petri nets and analyse their unfoldings, essentially an acyclic version of the Petri net whose simpler structure permits easier analysis [48].

However, Petri nets are inadequate to model concurrent read accesses to the same resource. Such situations often arise naturally, for instance in concurrent databases or in asynchronous circuits. The encoding tricks typically used to model these cases in Petri nets make the unfolding technique inefficient. Contextual nets, which explicitly do model concurrent read accesses, address this problem. Their accurate representation of concurrency makes contextual unfoldings up to exponentially smaller in certain situations. An abstract algorithm for contextual unfoldings was first given in [39]. 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 [57] is currently at an abstract level. In the short term, we want to put this work into practice, requiring some theoretical groundwork, as well as programming and experimentation.

Another well-known approach to verifying concurrent systems is *partial-order reduction*, exemplified by the tool SPIN. Although it is known that both partial-order reduction and unfoldings have their respective strengths and weaknesses, we are not aware of any conclusive comparison between the two techniques. Spin comes with a high-level modeling language having an explicit notion of processes, communication channels, and variables. Indeed, the reduction techniques implemented in Spin exploit the specific properties of these features. On the other side, while there exist highly efficient tools for unfoldings, Petri nets are a relatively general low-level formalism, so these techniques do not exploit properties of higher language features. Our work on contextual unfoldings and CMPs represents a first step to make unfoldings exploit richer models. In the long run, we wish raise the unfolding technique to a suitable high-level modelling language and develop appropriate tool support.

3.2. Management of Quantitative Behavior

Participants: Thomas Chatain, Stefan Haar, Serge Haddad.

3.2.1. Introduction

Besides the logical functionalities of programs, the *quantitative* aspects of component behavior and interaction play an increasingly important role.

- *Real-time* properties cannot be neglected even if time is not an explicit functional issue, since transmission delays, parallelism, etc, can lead to time-outs striking, and thus change even the logical course of processes. Again, this phenomenon arises in telecommunications and web services, but also in transport systems.
- In the same contexts, *probabilities* need to be taken into account, for many diverse reasons such as unpredictable functionalities, or because the outcome of a computation may be governed by race conditions.
- Last but not least, constraints on *cost* cannot be ignored, be it in terms of money or any other limited resource, such as memory space or available CPU time.

Traditional mainframe systems were proprietary and (essentially) localized; therefore, impact of delays, unforeseen failures, etc. could be considered under the control of the system manager. It was therefore natural, in verification and control of systems, to focus on *functional* behavior entirely.

With the increase in size of computing system and the growing degree of compositionality and distribution, quantitative factors enter the stage:

- calling remote services and transmitting data over the web creates *delays*;
- remote or non-proprietary components are not "deterministic", in the sense that their behavior is uncertain.

Time and *probability* are thus parameters that management of distributed systems must be able to handle; along with both, the *cost* of operations is often subject to restrictions, or its minimization is at least desired. The mathematical treatment of these features in distributed systems is an important challenge, which *MExICo* is addressing; the following describes our activities concerning probabilistic and timed systems. Note that cost optimization is not a current activity but enters the picture in several intended activities.

3.2.2. Probabilistic distributed Systems

Participants: Stefan Haar, Serge Haddad.

3.2.2.1. Non-sequential probabilistic processes

Practical fault diagnosis requires to select explanations of *maximal likelihood*. For partial-order based diagnosis, this leads therefore to the question what the probability of a given partially ordered execution is. In Benveniste et al. [42], [35], we presented a model of stochastic processes, whose trajectories are partially ordered, based on local branching in Petri net unfoldings; an alternative and complementary model based on Markov fields is developed in [52], which takes a different view on the semantics and overcomes the first model's restrictions on applicability.

Both approaches abstract away from real time progress and randomize choices in *logical* time. On the other hand, the relative speed - and thus, indirectly, the real-time behavior of the system's local processes - are crucial factors determining the outcome of probabilistic choices, even if non-determinism is absent from the system.

In another line of research [44] we have studied the likelihood of occurrence of non-sequential runs under random durations in a stochastic Petri net setting. It remains to better understand the properties of the probability measures thus obtained, to relate them with the models in logical time, and exploit them e.g. in *diagnosis*.

3.2.2.2. Distributed Markov Decision Processes

Participant: Serge Haddad.

Distributed systems featuring non-deterministic and probabilistic aspects are usually hard to analyze and, more specifically, to optimize. Furthermore, high complexity theoretical lower bounds have been established for models like partially observed Markovian decision processes and distributed partially observed Markovian decision processes. We believe that these negative results are consequences of the choice of the models rather than the intrinsic complexity of problems to be solved. Thus we plan to introduce new models in which the associated optimization problems can be solved in a more efficient way. More precisely, we start by studying connection protocols weighted by costs and we look for online and offline strategies for optimizing the mean cost to achieve the protocol. We have been cooperating on this subject with the SUMO team at Inria Rennes; in the joint work [36]; 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.

3.2.3. Large scale probabilistic systems

Addressing large-scale probabilistic systems requires to face state explosion, due to both the discrete part and the probabilistic part of the model. In order to deal with such systems, different approaches have been proposed:

- Restricting the synchronization between the components as in queuing networks allows to express the steady-state distribution of the model by an analytical formula called a product-form [41].

- 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 [46].
- 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 [58] .

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 [2]); (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.

3.2.4. Real time distributed systems

Nowadays, software systems largely depend on complex timing constraints and usually consist of many interacting local components. Among them, railway crossings, traffic control units, mobile phones, computer servers, and many more safety-critical systems are subject to particular quality standards. It is therefore becoming increasingly important to look at networks of timed systems, which allow real-time systems to operate in a distributed manner.

Timed automata are a well-studied formalism to describe reactive systems that come with timing constraints. For modeling distributed real-time systems, networks of timed automata have been considered, where the local clocks of the processes usually evolve at the same rate [54] [45]. It is, however, not always adequate to assume that distributed components of a system obey a global time. Actually, there is generally no reason to assume that different timed systems in the networks refer to the same time or evolve at the same rate. Any component is rather determined by local influences such as temperature and workload.

3.2.4.1. Implementation of Real-Time Concurrent Systems

Participants: Thomas Chatain, Stefan Haar, Serge Haddad.

This was one of the tasks of the ANR ImpRo.

Formal models for real-time systems, like timed automata and time Petri nets, have been extensively studied and have proved their interest for the verification of real-time systems. On the other hand, the question of using these models as specifications for designing real-time systems raises some difficulties. One of those comes from the fact that the real-time constraints introduce some artifacts and because of them some syntactically correct models have a formal semantics that is clearly unrealistic. One famous situation is the case of Zeno executions, where the formal semantics allows the system to do infinitely many actions in finite time. But there are other problems, and some of them are related to the distributed nature of the system. These are the ones we address here.

One approach to implementability problems is to formalize either syntactical or behavioral requirements about what should be considered as a reasonable model, and reject other models. Another approach is to adapt the formal semantics such that only realistic behaviors are considered.

These techniques are preliminaries for dealing with the problem of implementability of models. Indeed implementing a model may be possible at the cost of some transformation, which make it suitable for the target device. By the way these transformations may be of interest for the designer who can now use high-level features in a model of a system or protocol, and rely on the transformation to make it implementable.

We aim at formalizing and automating translations that preserve both the timed semantics and the concurrent semantics. This effort is crucial for extending concurrency-oriented methods for logical time, in particular for exploiting partial order properties. In fact, validation and management - in a broad sense - of distributed systems is not realistic *in general* without understanding and control of their real-time dependent features; the link between real-time and logical-time behaviors is thus crucial for many aspects of *MExICo*'s work.

4. Application Domains

4.1. Telecommunications

Participants: Stefan Haar, Serge Haddad.

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 Systems

Participants: Thomas Chatain, Matthias Függer, 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 Da-Jung Cho, Manish Kushwaha (INRA), and Thomas Nowak (LRI) on biological infection models for *E. coli* colonies and M13 phages.

4.3. Autonomous Vehicles

Participant: Serge Haddad.

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.

5. Highlights of the Year

5.1. Highlights of the Year

5.1.1. Reaching agreement in unstable times

Reaching approximate agreement in a distributed system among a set of local input values is a problem that often is repeatedly solved in artificial and natural distributed systems. Time efficient algorithms for this problem are thus of great theoretical and practical relevance. In [28] we studied the performance of such algorithms in dynamic networks. We showed lower time complexity bounds, demonstrating that already relatively simple broadcast and averaging algorithms achieve optimal time complexity. The results also imply new tight lower time complexity bounds for approximate agreement in classic distributed computing models with stable network architectures; solving a previously open problem.

5.1.2. New Semantics and State Spaces for Biological networks (and beyond)

We have gained major new insights into the dynamics of biological networks by

- obtaining [34], on the one hand, bi-directional translations between Contextual nets and BNs and correspondences between results on synchronism sensitivities. Taking advantage of CPN semantics enabling more behaviour than the generalized asynchronous updating mode, we propose an encoding of BNs that ensures correct abstraction of any multivalued refinement; and

- [20], [32] investigating update modes for discrete networks. It is commonly expected that Boolean networks produce an over-approximation of behaviours (reachable configurations), and that subsequent refinements would only prune some impossible transitions. However, we show that even generalized asynchronous updating of Boolean networks, which subsumes the usual updating modes including synchronous and fully asynchronous, does not capture all transitions doable in a multivalued or timed refinement. We introduce a new semantics for interpreting BNs which meets with a correct abstraction of any multivalued refinements, with any update mode. This semantics subsumes all the usual updating modes, while enabling new behaviours achievable by more concrete models. Moreover, it appears that classical dynamical analyses of reachability and attractors have a simpler computational complexity: – reachability can be assessed in a polynomial number of iterations (instead of being PSPACE-complete with update modes); – attractors are hypercubes, and deciding the existence of attractors with a given upper-bounded dimension is in NP (instead of PSPACE-complete with update modes). The computation of iterations is in NP in the very general case, and is linear when local functions are monotonic, or with some usual representations of functions of BNs (binary decision diagrams, Petri nets, automata networks, etc.). In brief, the most permissive semantics of BNs enables a correct abstract reasoning on dynamics of BNs, with a greater tractability than previously introduced update modes. These works open new perspectives in concurrent semantics, and at the same time will allow to capture hitherto inaccessible phenotypes and pathways in biological networks.

5.1.3. Awards

-

6. New Software and Platforms

6.1. 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++

- Participants: Benoît Barbot, Hilal Djafri, Marie Duflot-Kremer, Paolo Ballarini and Serge Haddad
- Contact: Benoît Barbot
- URL: http://www.lsv.ens-cachan.fr/~barbot/cosmos/

6.2. 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.

- Participants: Alban Linard, Fabrice Kordon, Laure Petrucci and Serge Haddad
- Partners: LIP6 - LSV - LIPN (Laboratoire d'Informatique de l'Université Paris Nord)
- Contact: Serge Haddad
- URL: http://www.cosyverif.org/

6.3. 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.

- Participant: Stefan Schwoon
- Contact: Stefan Schwoon
- URL: http://www.lsv.ens-cachan.fr/~schwoon/tools/mole/

7. New Results

7.1. Contract Based Design of Symbolic Controllers for Interconnected Multiperiodic Sampled-Data Systems

This paper deals with the synthesis of symbolic controllers for interconnected sampled-data systems where each component has its own sampling period. A compositional approach based on continuous-time assume-guarantee contracts is used. We provide sufficient conditions guaranteeing for a sampled-data system, satisfaction of an assume-guarantee contract and completeness of trajectories. Then, compositional results can be used to reason about interconnection of mul-tiperiodic sampled-data systems. We then show how discrete abstractions and symbolic control techniques can be applied to enforce the satisfaction of contracts and ensure completeness of trajectories. Finally, theoretical results are applied to a vehicle platooning problem on a circular road, which show the effectiveness of our approach.

7.2. Boolean Networks: Beyond Generalized Asynchronicity

Boolean networks are commonly used in systems biology to model dynamics of biochemical networks by abstracting away many (and often unknown) parameters related to speed and species activity thresholds. It is then expected that Boolean networks produce an over-approximation of behaviours (reachable configurations), and that subsequent refinements would only prune some impossible transitions. However, we show that even generalized asynchronous updating of Boolean networks, which subsumes the usual updating modes including synchronous and fully asynchronous, does not capture all transitions doable in a multi-valued or timed refinement. We define a structural model transformation which takes a Boolean network as input and outputs a new Boolean network whose asynchronous updating simulates both synchronous and asynchronous updating of the original network, and exhibits even more behaviours than the generalized asynchronous updating. We argue that these new behaviours should not be ignored when analyzing Boolean networks, unless some knowledge about the characteristics of the system explicitly allows one to restrict its behaviour.

7.3. Most Permissive Semantics of Boolean Networks

The usual update modes of Boolean networks (BNs), including synchronous and (generalized) asynchronous, fail to capture behaviours introduced by multivalued refinements. Thus, update modes do not allow a correct abstract reasoning on dynamics of biological systems, as they may lead to reject valid BN models. We introduce a new semantics for interpreting BNs which meets with a correct abstraction of any multivalued refinements, with any update mode. This semantics subsumes all the usual updating modes, while enabling new behaviours achievable by more concrete models. Moreover, it appears that classical dynamical analyses of reachability and attractors have a simpler computational complexity: – reachability can be assessed in a polynomial number of iterations (instead of being PSPACE-complete with update modes); – attractors are hypercubes, and deciding the existence of attractors with a given upper-bounded dimension is in NP (instead of PSPACE-complete with update modes). The computation of iterations is in NP in the very general case, and is linear when local functions are monotonic, or with some usual representations of functions of BNs (binary decision diagrams, Petri nets, automata networks, etc.). In brief, the most permissive semantics of BNs enables a correct abstract reasoning on dynamics of BNs, with a greater tractability than previously introduced update modes. This technical report lists the main definitions and properties of the most permissive semantics of BNs, and draw some remaining open questions.

7.4. Concurrency in Boolean networks

Boolean networks (BNs) are widely used to model the qualitative dynamics of biological systems. Besides the logical rules determining the evolution of each component with respect to the state of its regulators, the scheduling of components updates can have a dramatic impact on the predicted behaviours. In this paper, we explore the use of Contextual Petri Nets (CPNs) to study dynamics of BNs with a concurrency theory perspective. After showing bi-directional translations between CPNs and BNs and analogies between results on synchronism sensitivies, we illustrate that usual updating modes for BNs can miss plausible behaviours, i.e., incorrectly conclude on the absence/impossibility of reaching specific configurations. Taking advantage of CPN semantics enabling more behaviour than the generalized asynchronous updating mode, we propose an encoding of BNs ensuring a correct abstraction of any multivalued refinement, as one may expect to achieve when modelling biological systems with no assumption on its time features.

7.5. On the Composition of Discrete and Continuous-time Assume-Guarantee Contracts for Invariance

Many techniques for verifying invariance properties are limited to systems of moderate size. In this paper, we propose an approach based on assume-guarantee contracts and compositional reasoning for verifying invariance properties of a broad class of discrete-time and continuous-time systems consisting of interconnected components. The notion of assume-guarantee contracts makes it possible to divide responsibilities among the system components: a contract specifies an invariance property that a component must fulfill under some assumptions on the behavior of its environment (i.e. of the other components). We define weak and strong semantics of assume-guarantee contracts for both discrete-time and continuous-time systems. We then establish a certain number of results for compositional reasoning, which allow us to show that a global invariance property of the whole system is satisfied when all components satisfy their own contract. Interestingly, we show that the weak satisfaction of the contract is sufficient to deal with cascade compositions, while strong satisfaction is needed to reason about feedback composition. Specific results for systems described by differential inclusions are then developed. Throughout the paper, the main results are illustrated using simple examples.

7.6. Compositional synthesis of state-dependent switching control

We present a correct-by-design method of state-dependent control synthesis for sampled switching systems. Given a target region R of the state space, our method builds a capture set S and a control that steers any element of S into R. The method works by iterated backward reachability from R. The method is also used to synthesize a recurrence control that makes any state of R return to R infinitely often. We explain how the

synthesis method can be performed in a compositional manner, and apply it to the synthesis of a compositional control of a concrete floor-heating system with 11 rooms and up to $2^1 1 = 2048$ toswitching modes.

7.7. An Improved Algorithm for the Co3ntrol Synthesis of Nonlinear Sampled Switched Systems

A novel algorithm for the control synthesis for nonlinear switched systems is presented in this paper. Based on an existing procedure of state-space bisection and made available for nonlinear systems with the help of guaranteed integration, the algorithm has been improved to be able to consider longer patterns of modes with a better pruning approach. Moreover, the use of guaranteed integration also permits to take bounded perturbations and varying parameters into account. It is particularly interesting for safety critical applications, such as in aeronautical, military or medical fields. The whole approach is entirely guaranteed and the induced controllers are correct-by-design. Some experimentations are performed to show the important gain of the new algorithm.

7.8. Control Synthesis for Stochastic Switched Systems using the Tamed Euler Method

In this paper, we explain how, under the one-sided Lipschitz (OSL) hypothesis, one can find an error bound for a variant of the Euler-Maruyama approximation method for stochastic switched systems. We then explain how this bound can be used to control stochastic switched switched system in order to stabilize them in a given region. The method is illustrated on several examples of the literature.

7.9. The Complexity of Diagnosability and Opacity Verification for Petri Nets

Diagnosability and opacity are two well-studied problems in discrete-event systems. We revisit these two problems with respect to expressiveness and complexity issues. We first relate different notions of diagnosability and opacity. We consider in particular fairness issues and extend the definition of Germanos et al. [ACM TECS, 2015] of weakly fair diagnosability for safe Petri nets to general Petri nets and to opacity questions. Second, we provide a global picture of complexity results for the verification of diagnosability and opacity. We show that diagnosability is NL-complete for finite state systems, PSPACE-complete for safe Petri nets (even with fairness), and EXPSPACE-complete for general Petri nets without fairness, while non diagnosability is inter-reducible with reachability when fault events are not weakly fair. Opacity is ESPACE-complete for safe Petri nets (even with fairness) and undecidable for general Petri nets already without fairness.

7.10. Integrating Simulink Models into the Model Checker Cosmos

We present an implementation for Simulink model executions in the statistical model-checker Cosmos. We take profit of this implementation for an hybrid modeling combining Petri nets and Simulink models.,Nous présentons une implémentation pour l'exécution de modèles Simulink dans le model-checker Cosmos. Cette implémentation est ensuite utilisée pour la simulation de modèles hybrides, combinant des réseaux de Petri et des modèles Simulink.

7.11. Bounds Computation for Symmetric Nets

Monotonicity in Markov chains is the starting point for quantitative abstraction of complex probabilistic systems leading to (upper or lower) bounds for probabilities and mean values relevant to their analysis. While numerous case studies exist in the literature, there is no generic model for which monotonicity is directly derived from its structure. Here we propose such a model and formalize it as a subclass of Stochastic Symmetric (Petri) Nets (SSNs) called Stochastic Monotonic SNs (SMSNs). On this subclass the monotonicity is proven by coupling arguments that can be applied on an abstract description of the state (symbolic marking). Our class includes both process synchronizations and resource sharings and can be extended to model open or cyclic closed systems. Automatic methods for transforming a non monotonic system into a monotonic one matching the MSN pattern, or for transforming a monotonic system with large state space into one with reduced state space are presented. We illustrate the interest of the proposed method by expressing standard monotonic models and modelling a flexible manufacturing system case study.

7.12. Distributed computation of vector clocks in Petri nets unfolding for test selection

Petri net unfoldings with time stamps allow to build distributed testers for distributed systems. However, the construction of the annotated unfolding of a distributed system currently remains a centralized task. In the aforemention paper, we extend a distributed unfolding technique in order to annotate the resulting unfolding with time stamps. This allows for distributed construction of distributed testers for distributed systems.

7.13. Hyper Partial Order Logic

We define HyPOL, a local hyper logic for partial order models, expressing properties of sets of runs. These properties depict shapes of causal dependencies in sets of partially ordered executions,with similarity relations defined as isomorphisms of past observations. Unsurprisingly, since comparison of projections are included, satisfiability of this logic is undecidable. We then addressmodel checking of HyPOL and show that, already for safe Petri nets, the problem is undecidable. Fortunately, sensible restrictions of observations and nets allow us to bring back model checking ofHyPOL to a decidable problem, namely model checking of MSO on graphs of bounded treewidth.

7.14. Integrating Simulink Models into the Model Checker Cosmos

We present an implementation for Simulink model executions in the statistical model-checker Cosmos. We take profit of this implementation for hybrid modeling and simulations combining Petri nets and Simulink models.

7.15. Site-Directed Deletion

We introduce a new bio-inspired operation called a site-directed deletion motivated from site-directed mutagenesis performed by enzymatic activity of DNA polymerase: Given two strings x and y, a site-directed deletion partially deletes a substring of x guided by the string y that specifies which part of a substring can be deleted. We study a few decision problems with respect to the new operation and examine the closure properties of the (iterated) site-directed deletion operations. We, then, define a site-directed deletion-closed (and-free) language L and investigate its decidability properties when L is regular or context-free.

7.16. Site-Directed Insertion: Decision Problems, Maximality and Minimality

Site-directed insertion is an overlapping insertion operation that can be viewed as analogous to the overlap assembly or chop operations that concatenate strings by overlapping a suffix and a prefix of the argument strings. We consider decision problems and language equations involving site-directed insertion. By relying on the tools provided by semantic shuffle on trajectories we show that one variable equations involving site-directed insertion and regular constants can be solved. We consider also maximal and minimal variants of the site-directed insertion operation.

7.17. A Faithful Binary Circuit Model with Adversarial Noise

Accurate delay models are important for static and dynamic timing analysis of digital circuits, and mandatory for formal verification. However, Függer et al. [IEEE TC 2016] proved that pure and inertial delays, which are employed for dynamic timing analysis in state-of-the-art tools like ModelSim, NC-Sim and VCS, do not yield faithful digital circuit models. Involution delays, which are based on delay functions that are mathematical involutions depending on the previous-output-to-input time offset, were introduced by Függer et al. [DATE'15] as a faithful alternative (that can easily be used with existing tools). Although involution delays were shown to predict real signal traces reasonably accurately, any model with a deterministic delay function is naturally limited in its modeling power. In this paper, we thus extend the involution model, by adding non-deterministic delay variations (random or even adversarial), and prove analytically that faithfulness is not impaired by this

generalization. Albeit the amount of non-determinism must be considerably restricted to ensure this property, the result is surprising: the involution model differs from non-faithful models mainly in handling fast glitch trains, where small delay shifts have large effects. This originally suggested that adding even small variations should break the faithfulness of the model, which turned out not to be the case. Moreover, the results of our simulations also confirm that this generalized involution model has larger modeling power and, hence, applicability.

7.18. Tight Bounds for Asymptotic and Approximate Consensus

We study the performance of asymptotic and approximate consensus algorithms under harsh environmental conditions. The asymptotic consensus problem requires a set of agents to repeatedly set their outputs such that the outputs converge to a common value within the convex hull of initial values. This problem, and the related approximate consensus problem, are fundamental building blocks in distributed systems where exact consensus among agents is not required or possible, e.g., man-made distributed control systems , and have applications in the analysis of natural distributed systems, such as flocking and opinion dynamics. We prove tight lower bounds on the contraction rates of asymptotic consensus algorithms in dynamic networks, from which we deduce bounds on the time complexity of approximate consensus algorithms. In particular, the obtained bounds show optimality of asymptotic and approximate consensus algorithms presented in [Charron-Bost et al., ICALP'16] for certain dynamic networks, including the weakest dynamic network model in which asymptotic and approximate consensus are solvable. As a corollary we also obtain asymptotically tight bounds for asymptotic consensus in the classical asynchronous model with crashes. Central to our lower bound proofs is an extended notion of valency, the set of reachable limits of an asymptotic consensus algorithm starting from a given configuration. We further relate topological properties of valencies to the solvability of exact consensus , shedding some light on the relation of these three fundamental problems in dynamic networks.

7.19. Pomsets and Unfolding of Reset Petri Nets

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 this paper we look at partial order semantics of such 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 to the class of safe reset Petri nets.

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

Naive handling of supply voltage droops in synchronous circuits results in conservative bounds on clock speeds, resulting in poor performance even if droops are rare. 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. However, state-of-the-art solutions incur synchronization delay to avoid that the clock signal is affected by metastability. Addressing the challenges discussed by Keith Bowman in his ASYNC 2017 keynote talk, we present an all-digital circuit that can respond to droops within a fraction of a clock cycle. This is achieved by delaying clock signals based on measurement values while they undergo synchronization simultaneously. We verify our solution by formally proving correctness, complemented by VHDL and Spice simulations of a 65 nm ASIC design confirming the theoretically obtained results.

7.21. Fast Multidimensional Asymptotic and Approximate Consensus

We study the problems of asymptotic and approximate consensus in which agents have to get their values arbitrarily close to each others' inside the convex hull of initial values, either without or with an explicit decision by the agents. In particular, we are concerned with the case of multidimensional data, i.e., the agents' values are d-dimensional vectors. We introduce two new algorithms for dynamic networks, subsuming

classical failure models like asynchronous message passing systems with Byzantine agents. The algorithms are the first to have a contraction rate and time complexity independent of the dimension d. In particular, we improve the time complexity from the previously fastest approximate consensus algorithm in asynchronous message passing systems with Byzantine faults by Mendes et al. [Distrib. Comput. 28].

7.22. Parameter Space Abstraction and Unfolding Semantics of Discrete Regulatory Networks

The modelling of discrete regulatory networks combines a graph specifying the pairwise influences between the variables of the system, and a parametrisation from which can be derived a discrete transition system. Given the influence graph only, the exploration of admissible parametrisations and the behaviours they enable is computationally demanding due to the combinatorial explosions of both parametrisation and reachable state space. This article introduces an abstraction of the parametrisation space and its refinement to account for the existence of given transitions, and for constraints on the sign and observability of influences. The abstraction uses a convex sub-lattice containing the concrete parametrisation space specified by its infimum and supremum parametrisations. It is shown that the computed abstractions are optimal, i.e., no smaller convex sublattice exists. Although the abstraction may introduce over-approximation, it has been proven to be conservative with respect to reachability of states. Then, an unfolding semantics for Parametric Regulatory Networks is defined, taking advantage of concurrency between transitions to provide a compact representation of reachable transitions. A prototype implementation is provided: it has been applied to several examples of Boolean and multi-valued networks, showing its tractability for networks with numerous components.

7.23. Interval Iteration Algorithm for MDPs and IMDPs

Markov Decision Processes (MDP) are a widely used model including both non-deterministic and probabilistic choices. Minimal and maximal probabilities to reach a target set of states, with respect to a policy resolving non-determinism, may be computed by several methods including value iteration. This algorithm, easy to implement and efficient in terms of space complexity, iteratively computes the probabilities of paths of increasing length. However, it raises three issues: (1) defining a stopping criterion ensuring a bound on the approximation, (2) analysing the rate of convergence, and (3) specifying an additional procedure to obtain the exact values once a sufficient number of iterations has been performed. The first two issues are still open and, for the third one, an upper bound on the number of iterations has been proposed. Based on a graph analysis and transformation of MDPs, we address these problems. First we introduce an interval iteration algorithm, for which the stopping criterion is straightforward. Then we exhibit its convergence rate. Finally we significantly improve the upper bound on the number of iterations required to get the exact values. We extend our approach to also deal with Interval Markov Decision Processes (IMDP) that can be seen as symbolic representations of MDPs.

7.24. Diagnosability of Repairable Faults

The diagnosis problem for discrete event systems consists in deciding whether some fault event occurred or not in the system, given partial observations on the run of that system. Diagnosability checks whether a correct diagnosis can be issued in bounded time after a fault, for all faulty runs of that system. This problem appeared two decades ago and numerous facets of it have been explored, mostly for permanent faults. It is known for example that diagnosability of a system can be checked in polynomial time, while the construction of a diagnoser is exponential. The present paper examines the case of transient faults, that can appear and be repaired. Diagnosability in this setting means that the occurrence of a fault should always be detected in bounded time, but also before the fault is repaired. Checking this notion of diagnosability is proved to be PSPACE-complete. It is also shown that faults can be reliably counted provided the system is diagnosable for faults and for repairs.

7.25. Metastability-Containing Circuits

In digital circuits, metastability can cause deteriorated signals that neither are logical 0 nor logical 1, breaking the abstraction of Boolean logic. Synchronizers, the only traditional countermeasure, exponentially decrease the odds of maintained metastability over time. We propose a fundamentally different approach: It is possible to deterministically contain metastability by fine-grained logical masking so that it cannot infect the entire circuit. At the heart of our approach lies a time-and value-discrete model for metastability in synchronous clocked digital circuits, in which metastability is propagated in a worst-case fashion. The proposed model permits positive results and passes the test of reproducing Marino's impossibility results. We fully classify which functions can be computed by circuits with standard registers. Regarding masking registers, we show that more functions become computable with each clock cycle, and that masking registers permit exponentially smaller circuits for some tasks. Demonstrating the applicability of our approach, we present the first fault-tolerant distributed clock synchronization algorithm that deterministically guarantees correct behavior in the presence of metastability. As a consequence, clock domains can be synchronized without using synchronizers, enabling metastability-free communication between them.

8. Bilateral Contracts and Grants with Industry

8.1. Bilateral Contracts with Industry

Our cooperation with industry took place in the context of a multi-lateral SystemX project, see below.

9. Partnerships and Cooperations

9.1. Regional Initiatives

- Serge Haddad and Yann Duplouy have been participating in the *Simulation pour la sécurité du véhicule autonome* (SVA) project at SystemX, in cooperation with Renault, on the application of formal methods to the development of embedded systems for autonomous vehicles.
- Matthias Függer co-leads the Digicosme working group HicDiesMeus on "Highly Constrained Discrete Agents for Modeling Natural Systems" (parsys.lri.fr/HicDiesMeus).
- Matthias Függer participates in the Farman project Dicimus in collaboration with Thomas Nowak (LRI). The project is on modeling of bacterial interactions using techniques from distributed computing theory and VLSI design.

9.2. National Initiatives

- Thomas Chatain, Stefan Haar, Serge Haddad and Stefan Schwoon are participating in the ANR Project ALGORECELL.
- Matthias Függer participates in the ANR project FREDDA on verification and synthesis of distributed algorithms.
- Laurent Fribourg participates in Digicosme Emergence Project "CODECSY" in collaboration with Antoine Girard (CentraleSupelec).

9.3. European Initiatives

Serge Haddad is a member of the European project ERC EQualIS "Enhancing the Quality of Interacting Systems" headed by Patricia Bouyer.

9.4. International Initiatives

9.4.1. Inria Associate Teams Not Involved in an Inria International Labs

9.4.1.1. LifeForm

Title: Life Sciences need formal Methods !

International Partner (Institution - Laboratory - Researcher):

Newcastle University (United Kingdom) - School of Computing Science - Victor Khomenko

Start year: 2016

See also: http://projects.lsv.ens-cachan.fr/LifeForm/

This project extends an existing cooperation between the MEXICO team and Newcastle University on partial-order based formal methods for concurrent systems. We enlarge the partnership to bioinformatics and synthetic biology. The proposal addresses addresses challenges concerning formal specification, verification, monitoring and control of synthetic biological systems, with use cases conducted in the Center for Synthetic Biology and the Bioeconomy (CSBB) in Newcastle. A main challenge is to create a solid modelling framework based on Petri-net type models that allow for causality analysis and rapid state space exploration for verification, monitoring and control purposes; a potential extension to be investigated concerns the study of attractors and cell reprogramming in Systems Biology.

9.4.2. Inria International Partners

9.4.2.1. Informal International Partners

Josep Carmona (UPC Barcelona) visited us in April and July 2018. He collaborated with Thomas Chatain on process mining.

9.5. International Research Visitors

9.5.1. Visits to International Teams

9.5.1.1. Research Stays Abroad

- Juraj Kolcˇák has started, in August 2018, a 6-month research visit in the MMM group / NII Tokyo (Japan), funded by the ERATO project, to work with the PI, Prof. Ichiro Hasuo. Stefan Haar has visited that group from Oct 29 to Friday Nov 2, preceded by a visit to Prof. Tatsuya Akutsu's group at Kyoto University (Uji campus) on Oct 26.

10. Dissemination

10.1. Promoting Scientific Activities

10.1.1. Scientific Events Organisation

10.1.1.1. General Chair, Scientific Chair

Matthias Függer was

- general co-chair of the IEEE International Symposium on Asynchronous Circuits and Systems (ASYNC) 2018 (www.async2018.wien)
- general co-chair of the Workshop of Emergent Algorithms and Network Dynamics (WENDY) 2018 (wendy.paris)

Serge Haddad is a member of the steering commitee of the Petri Nets conference.

10.1.2. Scientific Events Selection

10.1.2.1. Chair of Conference Program Committees

Thomas Chatain was co-chair of the program committe of ACSD 2018 interes.institute/acsd2018/.

10.1.2.2. Member of the Conference Program Committees

- Matthias Függer was a PC member of
  - the *21st IEEE International Symposium on Design and Diagnostics of Electronic Circuits and Systems(DDECS 2018)*,
- Stefan Haar was a PC member of
  - the *18th International Conference on Applications of Concurrency to Systems Design(ACSD 2018)*,
  - the workshop *Algorithms and Theories for the Analysis of Event Data 2018(ATAED 2018)*, and
  - the *International Workshop on Petri Nets and Modeling (PeMod '18)*.
- Laurent Fribourg was a PC member of
  - *Model-Based Design of Cyber Physical Systems* (CyPhy'18), October 4-5, 2018, Torino, Italy,
  - *12th International Conference on Reachability Problems* (RP'18), September 24-26 2018, Marseille, France,
  - *8th International Conference on New Computational Methods for Inverse Problems* (NCMIP'18), Ecole normale supérieure Paris-Saclay, France.
- Serge Haddad was a PC member of
  - *12th International Workshop on Verification and Evaluation of Computer and Communication Systems (VECOS 2018)*, Grenoble, France, September 2018
  - *5th International Symposium on Formal Approaches to Parallel and Distributed Systems (4PAD 2018)*, Orléans, France, July 2018

10.1.2.3. Reviewer

- Matthias Függer reviewed for Automatica, ASYNC'18, DDECS'18, Philisophical Transactions, DISC'18, PODC'18, SIROCCO'18, STACS'18.
- Stefan Schwoon reviewed for MFCS and FSTTCS.
- Stefan Haar reviewed for FOSSACS 2019.

10.1.3. Journals

10.1.3.1. Member of the Editorial Boards

- Matthias Függer is guest editor for the special issue *Selected Papers from the 24th IEEE International Symposium on Asynchronous Circuits and Systems - ASYNC 2018*
- Stefan Haar is an associate editor for *Journal of Discrete Event Dynamic Systems*

10.1.3.2. Reviewer - Reviewing Activities

- Thomas Chatain reviewed for *Journal of Discrete Event Dynamic Systems*, *Transactions of the Society for Modeling and Simulation International*.
- Stefan Schwoon reviewed for *Journal of Discrete Event Dynamic Systems*, *Acta Informatica*, *ACM Transactions on Programming Languages and Systems*.
- Stefan Haar reviewed for *Scientific Annals of Computer Science* and *IEEE Transactions on Automatic Control*.

10.1.4. Invited Talks

- Serge Haddad gave an invited talk on “Time and Stochastic Petri Nets” at the tutorials of Petri Nets 2018, the 25th June 2018, Bratislava, Slovakia
- Matthias Függer gave an invited talk at ACSD’18 on "Challenges of circuit design: Circuits as robust distributed algorithms"

10.1.5. Scientific Expertise

- Stefan Schwoon acted as reviewer for the ERC 2018 Starting Grant call.
- Serge Haddad was expert for the allocation of the grants ”Prime d’Investissement Recherche de l’Université” of Sorbonne Université

10.2. Teaching - Supervision - Juries

10.2.1. Teaching

Note: we only list the teaching activities of researchers here, not those of our assistant and full professors.

Licence: Stefan Haar taught one half of the L3 course on formal languages (18 h EQ TD) at ENS Paris-Saclay.

Master: Matthias Függer and Stefan Haar each taught a module of 10 h EQTD in the *Jaques Herbrand* master MI course *Introduction à la recherche*.

Laurent Fribourg taught one half of M2 course on “Hybrid Automata” at MPRI (Master Parisien de Recherche en Informatique).

10.2.2. Supervision

PhD:

- Thomas Chatain is the supervisor of the PhD thesis of Mathilde Boltenhagen.
- Stefan Haar is the supervisor of the PhD theses of
  - **Juraj Kolcˇák** Unfoldings and Abstract Interpretation for Parametric Biological Regulatory Networks, started in March 2017, and of
  - **Hugues Mandon** on *Computational models and algorithms for the prediction of cell reprogramming strategies*, started on Oct. 1st,

  both at ENS Paris-Saclay.
- Laurent Fribourg is the supervisor of the PhD theses of
  - **Adnane Saoud***Compositional controller synthesis for cyber-physical systems*, started in October 2016, co-supervised by Antoine Girard (CentraleSupelec), funded by Digicosme projet Emergence Codecsys
  - **Jawher Jerray***Formal analysis of real-time systems*, started in October 2018, co-supervised by Etienne André (Paris 13), funded by University Paris 13, ED Galilée.
- Serge Haddad is the supervisor of the PhD thesis of Igor Khmelnitsky on Machine Learning and Verification of Infinite-state Systems co-supervised with Alain Finkel.

10.2.3. Juries

- Thomas Chatain reviewed the PhD thesis of Thi Thanh Huyen Nguyen, directed by Laure Petrucci and defended at Université Paris 13 in December 2018.
- Stefan Schwoon reviewed the PhD thesis of Adrien Pommellet, directed by Tayssir Touili and defended at Université Paris 13 in July 2018. He also acted as examinator for the PhD thesis of Huu Vu Nguyen.

10.3. Popularization

Laurent Fribourg was interviewed by *L’Édition de l’ Université Paris-Saclay* in: “La cyberphysique prépare l’usine de demain”, May 2018.

10.3.1. Internal or external Inria responsibilities

- Laurent Fribourg is Head of Institut Farman (FR 3311 CNRS & ENS Paris-Saclay).
- Serge Haddad is Head of the Computer Science Department of ENS Paris-Saclay.
- Stefan Haar is the president of Inria's COST-GTRI.

11. Bibliography

Major publications by the team in recent years

[1] S. HAAR, S. HADDAD, T. MELLITI, S. SCHWOON. *Optimal Constructions for Active Diagnosis*, in "Proceedings of the 33rd Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS'13)", Guwahati, India, A. SETH, N. VISHNOI (editors), Leibniz International Proceedings in Informatics, Leibniz-Zentrum für Informatik, December 2013, vol. 24, pp. 527-539, http://www.lsv.ens-cachan.fr/Publis/PAPERS/PDF/HHMS13-fsttcs.pdf

[2] S. HADDAD, J. MAIRESSE, H.-T. NGUYEN. *Synthesis and Analysis of Product-form Petri Nets*, in "Fundamenta Informaticae", 2013, vol. 122, n° 1-2, pp. 147-172, https://hal.archives-ouvertes.fr/hal-00925774

[3] A. LE COËNT, J. ALEXANDRE DIT SANDRETTO, A. CHAPOUTOT, L. FRIBOURG. *An Improved Algorithm for the Control Synthesis of Nonlinear Sampled Switched Systems*, in "Formal Methods in System Design", November 2017, vol. 53, n° 3, pp. 363-383 [*DOI :* 10.1007/s10703-017-0305-8], https://hal.archives-ouvertes.fr/hal-01399337

Publications of the year

Doctoral Dissertations and Habilitation Theses

[4] E. LEFAUCHEUX. *Controlling information in Probabilistic Systems*, Université Rennes 1, September 2018, https://hal.inria.fr/tel-01946840

Articles in International Peer-Reviewed Journals

[5] B. BÉRARD, S. HAAR, S. SCHMITZ, S. SCHWOON. *The Complexity of Diagnosability and Opacity Verification for Petri Nets*, in "Fundamenta Informaticae", July 2018, vol. 161, n° 4, pp. 317–349 [*DOI :* 10.3233/FI-2018-1706], https://hal.inria.fr/hal-01852119

[6] É. FABRE, L. HÉLOUËT, E. LEFAUCHEUX, H. MARCHAND. *Diagnosability of Repairable Faults*, in "Discrete Event Dynamic Systems", 2018, vol. 28, n° 2, pp. 183-213 [*DOI :* 10.1007/s10626-017-0255-8], https://hal.inria.fr/hal-01646911

[7] S. FRIEDRICHS, M. FÜGGER, C. LENZEN. *Metastability-Containing Circuits*, in "IEEE Transactions on Computers", March 2018, pp. 1167 - 1183 [*DOI :* 10.1109/TC.2018.2808185], https://hal.inria.fr/hal-01936292

[8] P. GASTIN, B. MONMEGE. *A Unifying Survey on Weighted Logics and Weighted Automata: Core Weighted Logic: Minimal and Versatile Specification of Quantitative Properties*, in "Soft Computing", February 2018, vol. 22, n° 4, pp. 1047-1065 [*DOI :* 10.1007/s00500-015-1952-6], https://hal.archives-ouvertes.fr/hal-01130216

[9] S. HADDAD. *Memoryless Determinacy of Finite Parity Games: Another Simple Proof*, in "Information Processing Letters", 2018, n^{o} 132, 3 p. , https://hal.inria.fr/hal-01541508

[10] S. HADDAD, B. MONMEGE. *Interval Iteration Algorithm for MDPs and IMDPs*, in "Theoretical Computer Science", July 2018, vol. 735, pp. 111 - 131 [*DOI :* 10.1016/J.TCS.2016.12.003], https://hal.archives-ouvertes.fr/hal-01809094

[11] J. KOLČÁK, D. ŠAFRÁNEK, S. HAAR, L. PAULEVÉ. *Parameter Space Abstraction and Unfolding Semantics of Discrete Regulatory Networks*, in "Theoretical Computer Science", 2018 [*DOI :* 10.1016/J.TCS.2018.03.009], https://hal.archives-ouvertes.fr/hal-01734805

[12] F. KORDON, H. GARAVEL, L. HILLAH, E. PAVIOT-ADET, L. JEZEQUEL, F. HULIN-HUBARD, E. AMPARORE, M. BECCUTI, B. BERTHOMIEU, H. EVRARD, P. G. JENSEN, D. LE BOTLAN, T. LIEBKE, J. MEIJER, J. SRBA, Y. THIERRY-MIEG, J. VAN DE POL, K. WOLF. *MCC'2017-The Seventh Model Checking Contest*, in "LNCS Transactions on Petri Nets and Other Models of Concurrency ", 2018, vol. 11090, pp. 181-209 [*DOI :* 10.1007/978-3-662-58381-4_9], https://hal.inria.fr/hal-01917492

[13] A. LE COËNT, L. FRIBOURG, N. MARKEY, F. DE VUYST, L. CHAMOIN. *Compositional synthesis of state-dependent switching control*, in "Theoretical Computer Science", November 2018, vol. 750, pp. 53-68 [*DOI :* 10.1016/J.TCS.2018.01.021], https://hal-utc.archives-ouvertes.fr/hal-01958874

[14] A. LE COËNT, L. FRIBOURG, N. MARKEY, F. DE VUYST, L. CHAMOIN. *Compositional synthesis of state-dependent switching control*, in "Theoretical Computer Science", November 2018, vol. 750, pp. 53-68 [*DOI :* 10.1016/J.TCS.2018.01.021], https://hal.archives-ouvertes.fr/hal-01860379

International Conferences with Proceedings

[15] B. BARBOT, B. BÉRARD, Y. DUPLOUY, S. HADDAD. *Integrating Simulink Models into the Model Checker Cosmos*, in "39th International Conference on Applications and Theory of Petri Nets and Concurrency", Bratislava, Slovakia, Lecture Notes in Computer Sciences, Springer, June 2018, vol. 10877, pp. 363-373 [*DOI :* 10.1007/978-3-319-91268-4_19], https://hal.archives-ouvertes.fr/hal-01916467

[16] B. BÉRARD, S. HAAR, L. HELOUET. *Hyper Partial Order Logic*, in "FSTTCS 2018 - Foundations of Software Technology and Theoretical Computer Science", Ahmedabad, India, December 2018, https://hal.inria.fr/hal-01884390

[17] D.-J. CHO, Y.-S. HAN, H. KIM, K. SALOMAA. *Site-Directed Deletion*, in "22nd International Conference on Developments in Language Theory (DLT 2018)", Tokyo, Japan, September 2018, https://hal.inria.fr/hal-01937635

[18] D.-J. CHO, Y.-S. HAN, K. SALOMAA, T. J. SMITH. *Site-Directed Insertion: Decision Problems, Maximality and Minimality*, in "20th International Conference on Descriptional Complexity of Formal Systems", Halifax, Canada, July 2018, https://hal.inria.fr/hal-01937654

[19] T. CHATAIN, R. GROSU. *Message from the ACSD 2018 Program Chairs*, in "2018 18th International Conference on Application of Concurrency to System Design (ACSD)", Bratislava, France, IEEE, June 2018 [*DOI :* 10.1109/ACSD.2018.00005], https://hal.inria.fr/hal-01936815

[20] T. CHATAIN, S. HAAR, L. PAULEVÉ. *Boolean Networks: Beyond Generalized Asynchronicity*, in "AUTOMATA 2018 - 24th IFIP WG 1.5 International Workshop on Cellular Automata and Discrete Complex Systems", Ghent, Belgium, J. M. BAETENS, M. KUTRIB (editors), Lecture Notes in Computer Science, Springer, June 2018, vol. 10875, pp. 29-42 [*DOI :* 10.1007/978-3-319-92675-9_3], https://hal.archives-ouvertes.fr/hal-01768359

[21] M. FÜGGER, A. KINALI, C. LENZEN, B. WIEDERHAKE. *Fast All-Digital Clock Frequency Adaptation Circuit for Voltage Droop Tolerance*, in "24th IEEE International Symposium on Asynchronous Circuits and Systems (ASYNC)", Wien, Austria, May 2018, https://hal.inria.fr/hal-01936403

[22] S. HAAR, J. KOLČÁK, L. PAULEVÉ. *Combining Refinement of Parametric Models with Goal-Oriented Reduction of Dynamics*, in "VMCAI 2019 - 20th International Conference on Verification, Model Checking, and Abstract Interpretation", Lisbon, Portugal, January 2019, https://arxiv.org/abs/1811.12377 , https://hal.archives-ouvertes.fr/hal-01940174

[23] L. JEZEQUEL, T. CHATAIN, M. COMLAN, D. DELFIEU, O. H. ROUX. *Pomsets and Unfolding of Reset Petri Nets*, in "12th International Conference on Language and Automata Theory and Applications (LATA 2018)", Ramat Gan, Israel, S. T. KLEIN, C. MARTÍN-VIDE, D. SHAPIRA (editors), Language and Automata Theory and Applications, Springer, April 2018, https://hal.archives-ouvertes.fr/hal-01766530

[24] L. JEZEQUEL, A. MADALINSKI, S. SCHWOON. *Distributed computation of vector clocks in Petri nets unfolding for test selection*, in "Workshop on Discrete Event Systems (WODES)", Sorrento, Italy, May 2018, https://hal.inria.fr/hal-01735406

[25] A. SAOUD, A. GIRARD, L. FRIBOURG. *Contract Based Design of Symbolic Controllers for Interconnected Multiperiodic Sampled-Data Systems*, in "57th IEEE Conference on Decision and Control (CDC 2018)", Miami Beach, FL, United States, December 2018, https://hal.archives-ouvertes.fr/hal-01857389

Conferences without Proceedings

[26] M. FÜGGER, J. MAIER, R. NAJVIRT, T. NOWAK, U. SCHMID. *A Faithful Binary Circuit Model with Adversarial Noise*, in "DATE 2018 - Design, Automation and Test in Europe Conference and Exhibition", Dresden, Germany, March 2018 [*DOI :* 10.23919/DATE.2018.8342219], https://hal.archives-ouvertes.fr/hal-01738254

[27] M. FÜGGER, T. NOWAK. *Fast Multidimensional Asymptotic and Approximate Consensus*, in "International Symposium on DIStributed Computing (DISC) 2018", New Orleans, United States, October 2018 [*DOI :* 10.4230/LIPIcs.DISC.2018.27], https://hal.archives-ouvertes.fr/hal-01936316

[28] M. FÜGGER, T. NOWAK, M. SCHWARZ. *Tight Bounds for Asymptotic and Approximate Consensus*, in "ACM Symposium on Principles of Distributed Computing (PODC'18)", Egham, United Kingdom, July 2018 [*DOI :* 10.1145/3212734.3212762], https://hal.archives-ouvertes.fr/hal-01799039

[29] A. LE COËNT, L. FRIBOURG, J. VACHER. *Control Synthesis for Stochastic Switched Systems using the Tamed Euler Method*, in "6th IFAC Conference on Analysis and Design of Hybrid Systems ADHS 2018 ADHS'18,", Oxford, United Kingdom, July 2018, vol. 51, n^{o} 16, https://hal.archives-ouvertes.fr/hal-01959845

[30] A. SAOUD, A. GIRARD, L. FRIBOURG. *On the Composition of Discrete and Continuous-time Assume-Guarantee Contracts for Invariance*, in "European Control Conference (ECC 2018)", Limassol, Cyprus, 2018, https://hal.archives-ouvertes.fr/hal-01712710

Research Reports

[31] B. BARBOT, B. BÉRARD, Y. DUPLOUY, S. HADDAD. *Integrating Simulink Models into the Model Checker Cosmos*, LSV, ENS Cachan, CNRS, Inria, Université Paris-Saclay, Cachan (France) ; LIP6, Sorbonne Université, CNRS, UMR 7606 ; LACL, Université Paris-Est, March 2018, https://hal.archives-ouvertes.fr/hal-01725835

[32] T. CHATAIN, S. HAAR, L. PAULEVÉ. *Most Permissive Semantics of Boolean Networks*, LRI, Univ. Paris-Sud, CNRS, Inria, Université Paris-Saclay ; LSV, ENS Cachan, CNRS, Inria, Université Paris-Saclay, Cachan (France), 2018, https://arxiv.org/abs/1808.10240 , https://hal.archives-ouvertes.fr/hal-01864693

Other Publications

[33] B. BARBOT, M. BECCUTI, G. FRANCESCHINIS, S. HADDAD, C. PICARONNY. *Bounds Computation for Symmetric Nets*, March 2018, working paper or preprint, https://hal.inria.fr/hal-01726011

[34] T. CHATAIN, S. HAAR, J. KOLČÁK, A. THAKKAR, L. PAULEVÉ. *Concurrency in Boolean networks*, October 2018, working paper or preprint, https://hal.inria.fr/hal-01893106

References in notes

[35] S. ABBES, A. BENVENISTE, S. HAAR. *A Petri net model for distributed estimation*, in "Proc. MTNS 2004, Sixteenth International Symposium on Mathematical Theory of Networks and Systems, Louvain (Belgium), ISBN 90-5682-517-8", 2004

[36] S. AKSHAY, N. BERTRAND, S. HADDAD, L. HELOUET. *The steady-state control problem for Markov decision processes*, in "Qest 2013", Buenos Aires, Argentina, K. R. JOSHI, M. SIEGLE, M. STOELINGA, P. R. D'ARGENIO (editors), Springer, September 2013, vol. 8054, pp. 290-304, https://hal.inria.fr/hal-00879355

[37] R. ALUR, K. ETESSAMI, M. YANNAKAKIS. *Realizability and Verification of MSC Graphs*, in "Theor. Comput. Sci.", 2005, vol. 331, n^{o} 1, pp. 97–114

[38] P. BALDAN, TH. CHATAIN, S. HAAR, B. KÖNIG. *Unfolding-based Diagnosis of Systems with an Evolving Topology*, in "Information and Computation", October 2010, vol. 208, n^{o} 10, pp. 1169-1192, http://www.lsv.ens-cachan.fr/Publis/PAPERS/PDF/BCHK-icomp10.pdf

[39] P. BALDAN, A. CORRADINI, B. KÖNIG, S. SCHWOON. *McMillan's complete prefix for contextual nets*, in "Transactions on Petri Nets and Other Models of Concurrency", November 2008, vol. 1, pp. 199–220, Volume 5100 of Lecture Notes in Computer Science

[40] P. BALDAN, S. HAAR, B. KOENIG. *Distributed Unfolding of Petri Nets*, in "Proc.FOSSACS 2006", LNCS, Springer, 2006, vol. 3921, pp. 126-141, Extended version: Technical Report CS-2006-1. Department of Computer Science, University Ca' Foscari of Venice

[41] F. BASKETT, K. M. CHANDY, R. R. MUNTZ, F. G. PALACIOS. *Open, Closed, and Mixed Networks of Queues with Different Classes of Customers*, in "J. ACM", April 1975, vol. 22, pp. 248–260, http://doi.acm.org/10.1145/321879.321887

[42] A. BENVENISTE, É. FABRE, S. HAAR. *Markov Nets: Probabilistic Models for distributed and concurrent Systems*, in "IEEE Transactions on Automatic Control", 2003, vol. 48 (11), pp. 1936-1950, Extended version: IRISA Research Report 1538

[43] P. BHATEJA, P. GASTIN, M. MUKUND, K. NARAYAN KUMAR. *Local testing of message sequence charts is difficult*, in "Proceedings of the 16th International Symposium on Fundamentals of Computation Theory (FCT'07)", Budapest, Hungary, E. CSUHAJ-VARJÚ, Z. ÉSIK (editors), Lecture Notes in Computer Science, Springer, August 2007, vol. 4639, pp. 76-87 [*DOI :* 10.1007/978-3-540-74240-1_8], http://www.lsv.ens-cachan.fr/Publis/PAPERS/PDF/BGMN-fct07.pdf

[44] A. BOUILLARD, S. HAAR, S. ROSARIO. *Critical paths in the Partial Order Unfolding of a Stochastic Petri Net*, in "Proceedings of the 7th International Conference on Formal Modelling and Analysis of Timed Systems (FORMATS'09)", Budapest, Hungary, J. OUAKNINE, F. VAANDRAGER (editors), Lecture Notes in Computer Science, Springer, September 2009, vol. 5813, pp. 43-57 [*DOI :* 10.1007/978-3-642-04368-0_6], http://www.lsv.ens-cachan.fr/Publis/PAPERS/PDF/BHR-formats09.pdf

[45] P. BOUYER, S. HADDAD, P.-A. REYNIER. *Timed Unfoldings for Networks of Timed Automata*, in "Proceedings of the 4th International Symposium on Automated Technology for Verification and Analysis (ATVA'06)", Beijing, ROC, S. GRAF, W. ZHANG (editors), Lecture Notes in Computer Science, Springer, October 2006, vol. 4218, pp. 292-306, http://www.lsv.ens-cachan.fr/Publis/PAPERS/PDF/BHR-atva06.pdf

[46] G. CHIOLA, C. DUTHEILLET, G. FRANCESCHINIS, S. HADDAD. *Stochastic Well-Formed Colored Nets and Symmetric Modeling Applications*, in "IEEE Transactions on Computers", November 1993, vol. 42, n^{o} 11, pp. 1343-1360, http://www.lsv.ens-cachan.fr/Publis/PAPERS/PS/CDFH-toc93.ps

[47] R. DEBOUK, D. TENEKETZIS. *Coordinated decentralized protocols for failure diagnosis of discrete-event systems*, in "Journal of Discrete Event Dynamical Systems: Theory and Application", 2000, vol. 10, pp. 33–86

[48] J. ESPARZA, K. HELJANKO. *Unfoldings - A Partial-Order Approach to Model Checking*, EATCS Monographs in Theoretical Computer Science, Springer, 2008

[49] É. FABRE, A. BENVENISTE, C. JARD, S. HAAR. *Diagnosis of Asynchronous Discrete Event Systems, a Net Unfolding Approach*, in "IEEE Trans. Aut. Control", 2003, vol. 48 (5), pp. 714-727

[50] É. FABRE, A. BENVENISTE, C. JARD, S. HAAR. *Distributed monitoring of concurrent and asynchronous systems*, in "Discrete Event Dynamic Systems: theory and application", 2005, vol. 15 (1), pp. 33-84, Preliminary version: Proc. CONCUR 2003, LNCS 2761, pp.1–28, Springer

[51] S. HAAR, A. BENVENISTE, É. FABRE, C. JARD. *Partial Order Diagnosability Of Discrete Event Systems Using Petri Net Unfoldings*, in "42nd IEEE Conference on Decision and Control (CDC)", 2003

[52] S. HAAR. *Probabilistic Cluster Unfoldings*, in "Fundamenta Informaticae", 2003, vol. 53 (3-4), pp. 281-314

[53] S. LAFORTUNE, Y. WANG, T.-S. YOO. *Diagnostic Decentralisé Des Systèmes A Evénements Discrets*, in "Journal Europeen des Systèmes Automatisés (RS-JESA)", August 2005, vol. 99, n^{o} 99, pp. 95–110

[54] K. G. LARSEN, P. PETTERSSON, W. YI. *Compositional and symbolic model-checking of real-time systems*, in "Proc. of RTSS 1995", IEEE Computer Society, 1995, pp. 76-89

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

[56] L. RICKER, K. RUDIE. *Knowledge Is a Terrible Thing to Waste: Using Inference in Discrete-Event Control Problems*, in "IEEE Transactions on Automatic Control", MarchSeptember 2007, vol. 52, n^{o} 3, pp. 428–441

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

[58] H. L. S. YOUNES, R. G. SIMMONS. *Statistical probabilistic model checking with a focus on time-bounded properties*, in "Inf. Comput.", September 2006, vol. 204, pp. 1368–1409 [*DOI :* 10.1016/J.IC.2006.05.002], http://dl.acm.org/citation.cfm?id=1182767.1182770