

RESEARCH CENTRE

Rennes - Bretagne Atlantique

IN PARTNERSHIP WITH:

CNRS, Université Rennes 1

2020

ACTIVITY REPORT

Project-Team

SUMO

SUper**VISION** of large **MO**dular and **DI**stributed systems

IN COLLABORATION WITH: Institut de recherche en informatique et
systèmes aléatoires (IRISA)

DOMAIN

**Algorithmics, Programming, Software
and Architecture**

THEME

Proofs and Verification

Contents

Project-Team SUMO	1
1 Team members, visitors, external collaborators	2
2 Overall objectives	3
2.1 Context	3
2.2 Necessity of quantitative models	3
2.3 Specificities of distributed systems	3
2.4 New issues raised by large systems	3
3 Research program	4
3.1 Introduction	4
3.2 Axis 1: Quantitative models	4
3.3 Axis 2: Large systems	5
3.4 Axis 3: Population models	5
3.5 Axis 4: Data-driven models	5
3.6 Transversal concern: missing models	6
4 Application domains	7
4.1 Smart transportation systems	7
4.2 Management of telecommunication networks and of data centers	7
4.3 Collaborative workflows	7
5 New software and platforms	8
5.1 New software	8
5.1.1 SIMSTORS	8
5.1.2 MOCHY	8
6 New results	9
6.1 New results on Axis 1: Quantitative models	9
6.1.1 Verification of Real-Time Models	9
6.1.2 Verification of Stochastic Models	9
6.1.3 Energy Games	10
6.2 New results on Axis 2: Large Systems Models	10
6.2.1 Supervisory Control	10
6.2.2 Multi-agent systems	11
6.3 New results on Axis 3: Population Models	11
6.4 New results on Axis 4: Data-driven Models	12
6.4.1 Crowdsourcing	12
6.4.2 Guarded Attribute Grammars	13
6.5 New results on Transversal Concern: Missing Models	13
7 Bilateral contracts and grants with industry	14
7.1 Bilateral contracts with industry	14
8 Partnerships and cooperations	14
8.1 International initiatives	14
8.1.1 Inria International Labs	14
8.1.2 Inria Associate Team not involved in an ILL	15
8.1.3 Inria International Partners	16
8.2 International research visitors	16
8.2.1 Visits to international teams	16
8.3 National initiatives	16

9 Dissemination	18
9.1 Promoting scientific activities	18
9.1.1 Scientific events: organisation	18
9.1.2 Scientific events: selection	18
9.1.3 Journal	18
9.1.4 Invited talks	19
9.1.5 Leadership within the scientific community	19
9.1.6 Scientific expertise	19
9.1.7 Research administration	19
9.2 Teaching - Supervision - Juries	19
9.2.1 Teaching	19
9.2.2 Supervision	20
9.2.3 Juries	21
9.3 Popularization	21
9.3.1 Education	21
10 Scientific production	21
10.1 Major publications	21
10.2 Publications of the year	22
10.3 Cited publications	24

Project-Team SUMO

Creation of the Team: 2013 January 01, updated into Project-Team: 2015 January 01

Keywords

Computer sciences and digital sciences

- A1.2.2. – Supervision
- A1.3. – Distributed Systems
- A1.5. – Complex systems
- A2.3. – Embedded and cyber-physical systems
- A2.4. – Formal method for verification, reliability, certification
- A2.4.2. – Model-checking
- A4.5. – Formal methods for security
- A6.4.3. – Observability and Controlability
- A6.4.6. – Optimal control
- A7.1.1. – Distributed algorithms
- A7.2. – Logic in Computer Science
- A8.2. – Optimization
- A8.6. – Information theory
- A8.11. – Game Theory

Other research topics and application domains

- B5.2.2. – Railway
- B6.2. – Network technologies
- B6.3.3. – Network Management
- B7.1. – Traffic management
- B8.5.2. – Crowd sourcing

1 Team members, visitors, external collaborators

Research Scientists

- Nathalie Bertrand [Team leader, Inria, Senior Researcher, HDR]
- Éric Badouel [Inria, Researcher, HDR]
- Éric Fabre [Inria, Senior Researcher, HDR]
- Blaise Genest [CNRS, Senior Researcher, HDR]
- Loïc Hélouët [Inria, Researcher, HDR]
- Thierry Jérón [Inria, Senior Researcher, HDR]
- Hervé Marchand [Inria, Researcher, HDR]
- Nicolas Markey [CNRS, Senior Researcher, HDR]
- Ocan Sankur [CNRS, Researcher]

Faculty Member

- Hugo Bazille [Univ de Rennes I, ATER, until Aug 2020]

Post-Doctoral Fellow

- Adrian Puerto Aubel [Inria, until Jun 2020]

PhD Students

- Sihem Cherrared [Inria, until Jan 2020]
- Emily Clement [Mitsubishi Electric, CIFRE]
- Arij Elmajed [Nokia, CIFRE, until Jan 2020]
- Leo Henry [Univ de Rennes I]
- Anirban Majumdar [CNRS]
- Abdul Majith Noordheen [Inria]
- Victor Roussanaly [Univ de Rennes I, until Oct 2020]
- Suman Sadhukhan [Inria]
- Bastien Thomas [Univ de Rennes I]

Technical Staff

- Antoine Thebault [Inria, Engineer, from Oct 2020]

Administrative Assistant

- Laurence Dinh [Inria]

External Collaborator

- Reiya Noguchi [Mitsubishi Electric]

2 Overall objectives

2.1 Context

Most software-driven systems we commonly use in our daily life are huge hierarchical assemblings of components. This observation runs from the micro-scale (multi-core chips) to the macro-scale (data centers), and from hardware systems (telecommunication networks) to software systems (choreographies of web services). The main characteristics of these pervasive applications are size, complexity, heterogeneity, and modularity (or concurrency). Besides, several such systems are actively used before they are fully mastered, or they have grown so much that they now raise new problems that are hardly manageable by human operators. While these systems and applications are becoming more essential, or even critical, the need for their *reliability*, *efficiency* and *manageability* becomes a central concern in computer science. The main objective of SUMO is to develop theoretical tools to address such challenges, according to the following axes.

2.2 Necessity of quantitative models

Several disciplines in computer science have of course addressed some of the issues raised by large systems. For example, formal methods (essentially for verification purposes), discrete-event systems (diagnosis, control, planning, and their distributed versions), but also concurrency theory (modelling and analysis of large concurrent systems). Practical needs have oriented these methods towards the introduction of quantitative aspects, such as time, probabilities, costs, and their combinations. This approach drastically changes the nature of questions that are raised. For example, verification questions become the reachability of a state in a limited time, the average sojourn duration in a state, the probability that a run of the system satisfies some property, the existence of control strategies with a given winning probability, etc. In this setting, exact computations are not always appropriate as they may end up with unaffordable complexities, or even with undecidability. Approximation strategies then offer a promising way around, and are certainly also a key to handling large systems. Approaches based on discrete-event systems follow the same trend towards quantitative models. For diagnosis aspects, one is interested in the most likely explanations to observed malfunctions, in the identification of the most informative tests to perform, or in the optimal placement of sensors. For control problems, one is of course interested in optimal control, in minimizing communications, in the robustness of the proposed controllers, in the online optimization of QoS (Quality of Service) indicators, etc.

2.3 Specificities of distributed systems

While the above questions have already received partial answers, they remain largely unexplored in a distributed setting. We focus on structured systems, typically a network of dynamic systems with known interaction topology, the latter being either static or dynamic. Interactions can be synchronous or asynchronous. The state-space explosion raised by such systems has been addressed through two techniques. The first one consists in adopting true-concurrency models, which take advantage of the parallelism to reduce the size of the trajectory sets. The second one looks for modular or distributed "supervision" methods, taking the shape of a network of local supervisors, one per component. While these approaches are relatively well understood, their mixing with quantitative models remains a challenge (as an example, there exists no proper setting assembling concurrency theory with stochastic systems). This field is largely open both for modeling, analysis and verification purposes, and for distributed supervision techniques. The difficulties combine with the emergence of data-driven distributed systems (as web services or data centric systems), where the data exchanged by the various components influence both the behaviors of these components and the quantitative aspects of their reactions (e.g. QoS). Such systems call for symbolic or parametric approaches for which a theory is still missing.

2.4 New issues raised by large systems

Some existing distributed systems like telecommunication networks, data centers, or large-scale web applications have reached sizes and complexities that reveal new management problems. One can

no longer assume that the model of the managed systems is static and fully known at any time and any scale. To scale up the management methods to such applications, one needs to be able to design reliable abstractions of parts of the systems, or to dynamically build a part of their model, following the needs of the management functions to realize. Besides, one does not wish to define management objectives at the scale of each single component, but rather to pilot these systems through high-level policies (maximizing throughput, minimizing energy consumption, etc.) These distributed systems and management problems have connections with other approaches for the management of large structured stochastic systems, such as Bayesian networks (BN) and their variants. The similarity can actually be made more formal: inference techniques for BN rely on the concept of conditional independence, which has a counterpart for networks of *dynamic* systems and is at the core of techniques like distributed diagnosis, distributed optimal planning, or the synthesis of distributed controllers. The potential of this connection is largely unexplored, but it suggests that one could derive from it good approximate management methods for large distributed dynamic systems.

3 Research program

3.1 Introduction

Since its creation in 2015, SUMO has successfully developed formal methods for large quantitative systems, in particular addressing verification, synthesis and control problems. Our current motivation is to expand this by putting emphasis on new concerns, such as algorithm efficiency, imprecision handling, and the more challenging objective of addressing incomplete or missing models. In the following we list a selection of detailed research goals, structured into four axes according to model classes: quantitative models, large systems, population models, and data-driven models. Some correspond to the pursuit of previously obtained results, others are more prospective.

3.2 Axis 1: Quantitative models

The analysis and control of quantitative models will remain at the heart of a large part of our research activities. In particular, we have two starting collaborative projects focusing on **timed models**, namely our ANR project TickTac and our collaboration with MERCE. The main expected outcome of TickTac is an open-source tool implementing the latest algorithms and allowing for quick prototyping of new algorithms. Several other topics will be explored in these collaborations, including robustness issues, game-theoretic problems, as well as the development of efficient algorithms, *e.g.* based on CEGAR approach or specifically designed for subclasses of automata (*e.g.* automata with few clocks and/or having a specific structure, as in [42]). Inspired by our collaboration with Alstom, we also aim at developing symbolic techniques for analysing non-linear timed models.

Stochastic models are another important focus for our research. On the one hand, we want to pursue our work on the optimization of non-standard properties for Markov decision processes, beyond the traditional verification questions, and explore *e.g.* long-run probabilities, and quantiles. Also, we aim at lifting our work on decisiveness from purely stochastic [40, 41] to non-deterministic and stochastic models in order to provide approximation schemes for the probability of (repeated) reachability properties in infinite-state Markov decision processes. On the other hand, in order to effectively handle large stochastic systems, we will pursue our work on approximation techniques. We aim at deriving simpler models, enjoying or preserving specific properties, and at determining the appropriate level of abstraction for a given system. One needs of course to quantify the approximation degrees (distances), and to preserve essential features of the original systems (explainability). This is a connection point between formal methods and the booming learning methods.

Regarding **diagnosis/opacity** issues, we will explore further the quantitative aspects. For diagnosis, the theory needs extensions to the case of incomplete or erroneous models, and to reconfigurable systems, in order to develop its applicability (see Sec. 3.6). There is also a need for non-binary causality analysis (*e.g.* performance degradations in complex systems). For opacity, we aim at quantifying the effort attackers must produce *vs* how much of a secret they can guess. We also plan to synthesize robust controllers resisting to sensor failures/attacks.

3.3 Axis 2: Large systems

Part of the background of SUMO is on the analysis and management of concurrent and modular/distributed systems, which we view as two main approaches to address state explosion problems. We will pursue the study of these models (including their quantitative features): verification of timed concurrent systems, robust distributed control of modular systems, resilient control to coalitions of attackers, distributed diagnosis, modular opacity analysis, distributed optimal planning, etc. Nevertheless, we have identified two new lines of effort, inspired by our application domains.

Reconfigurable systems. This is mostly motivated by applications at the convergence of virtualization techs with networking (Orange and Nokia PhDs). Software defined networks, either in the core (SDN/NFV) or at the edge (IoT) involve distributed systems that change structure constantly, to adapt to traffic, failures, maintenance, upgrades, etc. Traditional verification, control, diagnosis approaches (to mention only those) assume static and known models that can be handled as a whole. This is clearly insufficient here: one needs to adapt existing results to models that (sometimes automatically) change structure, incorporate new components/users or lose some, etc. At the same time, the programming paradigms for such systems (chaos monkey) incorporate resilience mechanisms, that should be considered by our models.

Hierarchical systems. Our experience with the regulation of subway lines (Alstom) revealed that large scale complex systems are usually described at a single level of granularity. Determining the appropriate granularity is a problem in itself. The control of such systems, with humans in the loop, can not be expressed at this single level, as tasks become too complex and require extremely skilled staff. It is rather desirable to describe models simultaneously at different levels of granularity, and to perform control at the appropriate level: humans in charge of managing the system by high level objectives, and computers in charge of implementing the appropriate micro-control sequences to achieve these tasks.

3.4 Axis 3: Population models

We want to step up our effort in parameterized verification of systems consisting of many identical components, so-called population models. In a nutshell our objectives summarize as "from Boolean to quantitative".

Inspired by our experience on the analysis of populations of yeasts, we aim at developing the quantitative analysis and control of population models, *e.g.* using Markov decision processes together with quantitative properties, and focusing on generating strategies with fast convergence.

As for broadcast networks, the challenge is to model the mobility of nodes (representing mobile ad hoc networks) in a faithful way. The obtained model should reflect on the one hand, the placement of nodes at a given time instant, and on the other hand, the physical movement of nodes over time. In this context, we will also use game theory techniques which allows one to study cooperative and conflictual behaviors of the nodes in the network, and to synthesize correct-by-design systems in adversarial environments.

As a new application area, we target randomized distributed algorithms. Our goal is to provide probabilistic variants of threshold automata [47] to represent fault-tolerant randomized distributed algorithms, designed for instance to solve the consensus problem. Most importantly, we then aim at developing new parameterized verification techniques, that will enable the automated verification of the correctness of such algorithms, as well as the assessment of their performances (in particular the expected time to termination).

In this axis, we will investigate whether fluid model checking and mean-field approximation techniques apply to our problems. More generally, we aim at a fruitful cross-fertilizing of these approaches with parameterized model-checking algorithms.

3.5 Axis 4: Data-driven models

In this axis, we will consider data-centric models, and in particular their application to crowd-sourcing. Many data-centric models such as Business Artifacts [48] orchestrate simple calls and answers to tasks performed by a single user. In a crowd-sourcing context, tasks are realized by pools of users, which may result in imprecise, uncertain and (partially) incompatible information. We thus need mechanisms to reconcile and fuse the various contributions in order to produce reliable information. Another aspect to

consider concerns answers of higher-order: how to allow users to return intentional answers, under the form of a sub-workflow (coordinated set of tasks) which execution will provide the intended value. In the framework of the ANR Headwork we will build on formalisms such as GAG (guarded attribute grammars) or variants of business artifacts to propose formalisms adapted to crowd-sourcing applications, and tools to analyze them. To address imprecision, we will study techniques to handle fuzziness in user answers, will explore means to set incentives (rewards) dynamically, and to set competence requirements to guide the execution of a complex workflow, in order to achieve an objective with a desired level of quality.

In collaboration with Open Agora, CESPAs and University of Yaoundé (Cameroun) we intend to implement in the GAG formalism some elements of argumentation theory (argumentation schemes, speech acts and dialogic games) in order to build a tool for the conduct of a critical discussion and the collaborative construction of expertise. The tool would incorporate point of view extraction (using clustering mechanisms), amendment management and consensus building mechanisms.

3.6 Transversal concern: missing models

We are concerned with one important lesson derived from our involvement in several application domains. Most of our background gets in force as soon as a perfect model of the system under study is available. Then verification, control, diagnosis, test, etc. can mobilize a solid background, or suggest new algorithmic problems to address. In numerous situations, however, assuming that a model is available is simply unrealistic. This is a major bottleneck for the impact of our research. We therefore intend to address this difficulty, in particular for the following domains.

- Model building for diagnosis. As a matter of fact, diagnosis theory hardly touches the ground to the extent that complete models of normal behavior are rarely available, and the identification of the appropriate abstraction level is unclear. Knowledge of faults and their effects is even less accessible. Also, the actual implemented systems may differ significantly from behaviors described in the norms. One therefore needs a theory for incomplete and erroneous models. Besides, one is often less bothered by partial observations than drowned by avalanches of alerts when malfunctions occur. Learning may come to the rescue, all the more that software systems may be deployed in sandpits and damaged for experimentation, thus allowing the collection of masses of labeled data. Competition on that theme clearly comes from Machine Learning techniques.
- Verification of large scale software. For some verification problems like the one we address in the IPL HAC-Specis, one does not have access to a formal model of the distributed program under study, but only to executions in a simulator. Formal verification poses new problems due to the difficulties to capture global states, to master state space explosion by gathering and exploiting concurrency information.
- Learning of stochastic models. Applications in bioinformatics often lead to large scale models, involving numerous chains of interactions between chemical species and/or cells. Fine grain models can be very precise, but very inefficient for inference or verification. Defining the appropriate levels of description/abstraction, given the available data and the verification goals, remains an open problem. This cannot be considered as a simple data fitting problem, as elements of biological knowledge must be combined with the data in order to preserve explainability of the phenomena.
- Testing and learning timed models: during conformance testing of a black-box implementation against its formal specification, one wants to detect non-conformances but may also want to learn the implementation model. Even though mixing testing and learning is not new, this is more recent and challenging for continuous-time models.
- Process mining. We intend to extend our work on process discovery using Petri net synthesis [39] by using negative information (*e.g.* execution traces identified as outliers) and quantitative information (probabilistic or fuzzy sets of execution traces) in order to infer more robust and precise models.

4 Application domains

4.1 Smart transportation systems

The smart-city trend aims at optimizing all functions of future cities with the help of digital technologies. We focus on the segment of urban trains, which will evolve from static and scheduled offers to reactive and eventually on-demand transportation offers. We address two challenges in this field. The first one concerns the optimal design of robust subway lines. The idea is to be able to evaluate, at design time, the performance of time tables and of different regulation policies. In particular, we focus on robustness issues: how can small perturbations and incidents be accommodated by the system, how fast will return to normality occur, when does the system become unstable? The second challenge concerns the design of new robust regulation strategies to optimize delays, recovery times, and energy consumption at the scale of a full subway line. These problems involve large-scale discrete-event systems, with temporal and stochastic features, and translate into robustness assessment, stability analysis and joint numerical/combinatorial optimization problems on the trajectories of these systems.

4.2 Management of telecommunication networks and of data centers

Telecommunication-network management is a rich provider of research topics for the team, and some members of SUMO have a long background of contacts and transfer with industry in this domain. Networks are typical examples of large distributed dynamic systems, and their management raises numerous problems ranging from diagnosis (or root-cause analysis), to optimization, reconfiguration, provisioning, planning, verification, etc. They also bring new challenges to the community, for example on the modeling side: building or learning a network model is a complex task, specifically because these models should reflect features like the layering, the multi-resolution view of components, the description of both functions, protocols and configuration, and they should also reflect dynamically-changing architectures. Besides modeling, management algorithms are also challenged by features like the size of systems, the need to work on abstractions, on partial models, on open systems, etc. The networking technology is now evolving toward software-defined networks, virtualized-network functions, multi-tenant systems, etc., which reinforces the need for more automation in the management of such systems.

Data centers are another example of large-scale modular dynamic and reconfigurable systems: they are composed of thousands of servers, on which virtual machines are activated, migrated, resized, etc. Their management covers issues like troubleshooting, reconfiguration, optimal control, in a setting where failures are frequent and mitigated by the performance of the management plane. We have a solid background in the coordination of the various autonomic managers that supervise the different functions/layers of such systems (hardware, middleware, web services, ...) Virtualization technologies now reach the domain of networking, and telecommunication operators/vendors evolve towards providers of distributed open clouds. This convergence of IT and networking strongly calls for new management paradigms, which is an opportunity for the team.

4.3 Collaborative workflows

A current trend is to involve end-users in collection and analysis of data. Examples of this trend are contributive science, crisis-management systems, and crowd-sourcing applications. All these applications are data-centric and user-driven. They are often distributed and involve complex, and sometimes dynamic workflows. In many cases, there are strong interactions between data and control flows: indeed, decisions taken regarding the next tasks to be launched highly depend on collected data. For instance, in an epidemic-surveillance system, the aggregation of various reported disease cases may trigger alerts. Another example is crowd-sourcing applications where user skills are used to complete tasks that are better performed by humans than computers. In return, this requires addressing imprecise and sometimes unreliable answers. We address several issues related to complex workflows and data. We study declarative and dynamic models that can handle workflows, data, uncertainty, and competence management.

Once these models are mature enough, we plan to build prototypes to experiment them on real use cases from contributive science, health-management systems, and crowd-sourcing applications. We also plan to define abstraction schemes allowing formal reasoning on these systems.

5 New software and platforms

5.1 New software

5.1.1 SIMSTORS

Name: Simulator for stochastic regulated systems

Keywords: Simulation, Public transport, Stochastic models, Distributed systems

Functional Description: SIMSTORS is a software for the simulation of stochastic concurrent timed systems. The heart of the software is a variant of stochastic and timed Petri nets, whose execution is controlled by a regulation policy (a controller), or a predetermined theoretical schedule. The role of the regulation policy is to control the system to realize objectives or a schedule when it exists with the best possible precision. SIMSTORS is well adapted to represent systems with randomness, parallelism, tasks scheduling, and resources. From 2015 to 2018, it was used for the P22 collaboration with Asstom Transport, to model metro traffic and evaluate performance of regulation solutions. In 2020, it was at the heart of a collaboration on multi-modal networks with Alstom transport Madrid. This software allows for step by step simulation, but also for efficient performance analysis of systems such as production cells or train systems. The initial implementation was released in 2015, and the software is protected by the APP.

Since then, SIMSTORS has been extended along two main axes: on one hand, SIMSTORS models were extended to handle situations where shared resources can be occupied by more than one object (this is of paramount importance to represent conveyors, roads occupied by cars, or train tracks with smoothed scheduling allowing shared sections among trains) with priorities, constraint on their ordering and individual characteristics. This allows for instance to model vehicles with different speeds on a road, while handling safety distance constraints. On the other hand, SIMSTORS models were extended to allow control of stochastic nets based on decision rules that follow optimization schemes. In 2020, it was used to define efficient traffic management techniques with planning during a collaboration with Roma 3 University.

Release Contributions: modeling of continuous vehicles movements

URL: <http://www.irisa.fr/sumo/Software/SIMSTORS/>

Authors: Abd El Karim Kecir, Loïc Hérouët

Contact: Loïc Hérouët

5.1.2 MOCHY

Name: MOdels for Concurrent and HYbrid systems

Keywords: Public transport, Hybrid models, Simulation, Performance analysis

Functional Description: Allows for the modeling of hybrid systems, schedule and regulation algorithms to optimize Key Performance Indicators. The tool allows for the fast simulation of these regulated models, to measure performance indicators.

Release Contributions: Co-simulation of time Petri nets and timetables (model for regulated urban train systems with a hold-on policy). Performance analysis for simple Key Performance Indicators.

URL: <https://adt-mochy.gitlabpages.inria.fr/mochy/>

Authors: Loïc Hérouët, Antoine Thebault, Didier Vojtisek

Contact: Loïc Hérouët

6 New results

6.1 New results on Axis 1: Quantitative models

6.1.1 Verification of Real-Time Models

Participants : Emily Clement, Blaise Genest, Loïc Hélouët, Thierry Jéron, Nicolas Markey, Reiya Noguchi, Ocan Sankur

Incremental methods for checking real-time consistency. Requirement engineering is a key phase in the development process. Ensuring that requirements are consistent is essential so that they do not conflict and admit implementations. In [28], we consider the formal verification of rt-consistency, which imposes that the inevitability of definitive errors of a requirement should be anticipated, and that of partial consistency, which was recently introduced as a more effective check. We generalize and formalize both notions for discrete-time timed automata, develop three incremental algorithms, and present experimental results.

Timed Negotiations. Negotiations were introduced in [44] as a model for concurrent systems with multiparty decisions. What is very appealing with negotiations is that it is one of the very few non-trivial concurrent models where several interesting problems, such as soundness, i.e. absence of deadlocks, can be solved in PTIME [43]. In this paper, we introduce the model of timed negotiations and consider the problem of computing the minimum and the maximum execution times of a negotiation. The latter can be solved using the algorithm of [45] computing costs in negotiations, but surprisingly minimum execution time cannot.

This year, we have proposed new algorithms [18] to compute both minimum and maximum execution time, that work in much more general classes of negotiations than [45], that only considered sound and deterministic negotiations. Further, we uncover the precise complexities of these questions, ranging from PTIME to Δ_2^P -complete. In particular, we show that computing the minimum execution time is more complex than computing the maximum execution time in most classes of negotiations we consider.

Robust strategies for reachability in timed automata. Reachability can be decided in polynomial space in timed automata. However, this result assumes arbitrary precision in the dates at which transitions have to be taken in order to reach the target state. In [26], we introduce and study *permissive strategies* in the setting of timed systems: instead of suggesting exact dates at which to take transitions, permissive strategies propose intervals of dates; the exact date at which the transitions are taken are chosen by an opponent. We develop an algorithm for computing *maximally permissive* strategies in acyclic timed automata and acyclic timed games.

6.1.2 Verification of Stochastic Models

Participants : Nathalie Bertrand, Hugo Bazille, Éric Fabre, Blaise Genest

Diagnosis and Degradation Control for Probabilistic Systems. Systems prone to faults are often equipped with a controller whose aim consists in restricting the behaviour of the system in order to perform a diagnosis. Such a task is called active diagnosis. However to avoid that the controller degrades the system in view of diagnosis, a second objective in terms of quality of service is usually assigned to the controller. In the framework of stochastic systems, a possible specification, called safe active diagnosis requires that the probability of correctness of the infinite (random) run is non null. In [13], we introduce and study two alternative specifications that are in many contexts more realistic. The notion of (γ, ν) -fault freeness associates with each run a value depending on the discounted length of its correct prefix where the discounting factor is γ . The controller has to ensure that the average of this value is above the threshold ν . The notion of α -resiliency requires that asymptotically, at every time step, a proportion greater than α of correct runs remain correct. From a semantic point of view, we determine the equivalences and (non) implications between the three notions of degradations both for finite and infinite systems.

From an algorithmic point of view, we establish the border between decidability and undecidability of the diagnosability problems. Furthermore in the positive case, we exhibit their precise complexity and propose a synthesis of the controller which may require an infinite memory.

Approximate Verification of Dynamic Bayesian Networks. We are interested in studying the evolution of large homogeneous populations of cells, where each cell is assumed to be composed of a group of biological players (species) whose dynamics is governed by a complex biological pathway, identical for all cells. Modeling the inherent variability of the species concentrations in different cells is crucial to understand the dynamics of the population. In [16], we focus on handling this variability by modeling each species by a random variable that evolves over time. This appealing approach runs into the curse of dimensionality since exactly representing a joint probability distribution involving a large set of random variables quickly becomes intractable as the number of variables grows. To make this approach amenable to biopathways, we explore different techniques to (i) approximate the exact joint distribution at a given time point, and (ii) track its evolution as time elapses.

6.1.3 Energy Games

Participants : Loïc Hélouët, Nicolas Markey

Games with reachability objectives under energy constraints. In [35], we study games with reachability objectives under energy constraints: such games are played on weighted graphs, and the aim is to reach a target state while always keeping the sum of the weights within a given interval. We prove that under strict energy constraints (either only lower-bound constraint or interval constraint), those games are LOGSPACE-equivalent to energy games with the same energy constraints but without reachability objective (i.e., for infinite runs). We then consider two relaxations of the upper-bound constraints (while keeping the lower-bound constraint strict): in the first one, called weak upper bound, the upper bound is absorbing, i.e., when the upper bound is reached, the extra energy is not stored; in the second one, we allow for temporary violations of the upper bound, imposing limits on the number or on the amount of violations.

We prove that when considering weak upper bound, reachability objectives require memory, but can still be solved in polynomial-time for one-player arenas; we prove that they are in coNP in the two-player setting. Allowing for bounded violations makes the problem PSPACE-complete for one-player arenas and EXPTIME-complete for two players. We then address the problem of existence of bounds for a given arena. We show that with reachability objectives, existence can be a simpler problem than the game itself, and conversely that with infinite games, existence can be harder.

6.2 New results on Axis 2: Large Systems Models

6.2.1 Supervisory Control

Participants : Loïc Hélouët, Hervé Marchand

Synthesis of Supervisors Robust Against Sensor Deception Attacks. A cyber-physical systems is usually composed of various physical and software components that interact with each other in ways that change with context. For such large system, one challenge is to ensure security against malicious users. As a security problem, we consider in [37] feedback control systems where sensor readings may be compromised by a malicious attacker intending on causing damage to the system. We study this problem at the supervisory layer of the control system, using discrete event systems techniques. We assume that the attacker can edit the outputs from the sensors of the system before they reach the supervisory controller. In this context, we formulate the problem of synthesizing a supervisor that is robust against the class of edit attacks on the sensor readings and present a solution methodology for this problem. This methodology blends techniques from games on automata with imperfect information with results from supervisory control theory of partially-observed discrete event systems. Necessary and sufficient conditions are provided for the investigated problem.

Enforcing Opacity in Modular Systems. Following preceding works regarding the enforcement of confidential informations for Cyber-physical systems, we considered in [30], the opacity control problem with coalition between attackers. In discrete-event systems, the opacity of a secret ensures that some behaviors or states cannot be inferred with certainty from partial observation of the system. Enforcing opacity in a discrete-event system, encoded by a finite labelled transition system (LTS), is a way to avoid information leakage. Checking opacity is decidable but costly (EXPTIME in the worst cases). We addressed opacity for modular systems in which every module, represented by an LTS, has to protect its own secret (a set of secret states S) w.r.t. a local attacker. Once the system is composed, we assume a coalition between the attackers that share their local view (called the global attacker). Assuming the global attacker can observe all interactions between modules, we provide a reduced-complexity opacity verification technique and an algorithm for constructing local controllers that enforces opacity for each secret separately.

6.2.2 Multi-agent systems

Participants : Arthur Queffelec, Ocan Sankur

Multi-agent path planning problems. In [14], we study a variant of the multi-agent path finding (MAPF) problem in which the group of agents are required to stay connected with a supervising base station throughout the execution. In addition, we consider the problem of covering an area with the same connectivity constraint. We show that both problems are PSPACE-complete on directed and undirected topological graphs while checking the existence of a bounded plan is NP-complete when the bound is given in unary (and PSPACE-hard when the encoding is in binary). Moreover, we identify a realistic class of topological graphs on which the decision problem falls in NLOGSPACE although the bounded versions remain NP-complete for unary encoding.

6.3 New results on Axis 3: Population Models

Participants : Nathalie Bertrand, Blaise Genest, Anirban Majumdar, Nicolas Markey, Suman Sadhukhan, Ocan Sankur

Dynamic network congestion games. In [23], we study congestion games which are a classical type of games studied in game theory, in which n players choose a resource, and their individual cost increases with the number of other players choosing the same resource. In network congestion games (NCGs), the resources correspond to simple paths in a graph, e.g. representing routing options from a source to a target. We introduce a variant of NCGs, referred to as dynamic NCGs: in this setting, players take transitions synchronously, they select their next transitions dynamically, and they are charged a cost that depends on the number of players simultaneously using the same transition. We study, from a complexity perspective, standard concepts of game theory in dynamic NCGs: social optima, Nash equilibria, and subgame perfect equilibria. Our contributions are the following: the existence of a strategy profile with social cost bounded by a constant is in PSPACE and NP-hard. (Pure) Nash equilibria always exist in dynamic NCGs; the existence of a Nash equilibrium with bounded cost can be decided in EXPSPACE, and computing a witnessing strategy profile can be done in doubly-exponential time. The existence of a subgame perfect equilibrium with bounded cost can be decided in 2EXPSPACE, and a witnessing strategy profile can be computed in triply-exponential time.

Model checking randomized distributed algorithms. Randomization is a powerful paradigm to solve hard problems, especially in distributed computing. Proving the correctness, and assessing the performances, of randomized distributed algorithms, is a very challenging research objective, that the verification community has started to address. In [12], we review existing model checking approaches to the verification of randomized distributed algorithms and identify further research directions.

Concurrent Games with Arbitrarily Many Players. Traditional concurrent games on graphs involve a fixed number of players, who take decisions simultaneously, determining the next state of the game. We recently introduced a parameterized variant of concurrent games on graphs, where the parameter is precisely the number of players. Parameterized concurrent games are described by finite graphs, in which the transitions bear finite-word languages to describe the possible move combinations that lead from one vertex to another. In the invited contribution [20], we report on results on two problems for such concurrent games with arbitrary many players. To start with, we studied the problem of determining whether the first player, say Eve, has a strategy to ensure a reachability objective against any strategy profile of her opponents as a coalition. In particular Eve's strategy should be independent of the number of opponents she actually has. We establish the precise complexities of the problem for reachability objectives. Second, we considered a synthesis problem, where one aims at designing a strategy for each of the (arbitrarily many) players so as to achieve a common objective. For safety objectives, we show that this kind of distributed synthesis problem is decidable [21].

Expressiveness of population protocols. We considered the number of states necessary for (leaderless) population protocols to be as expressive as unrestricted population protocols. It is well known since 2004 that the expressive power of unrestricted population protocols is exactly (quantifier-free) Presburger arithmetic with remainder predicates. In [24], we prove that it is sufficient to have a number of states polynomial in the size of a quantifier-free presburger arithmetic formula with remainder predicates. There is no difference between protocols with and without leaders. This result is surprising as it is known that for fast protocols, there is an exponential gap between the speed of leader and leaderless protocols.

6.4 New results on Axis 4: Data-driven Models

6.4.1 Crowdsourcing

Participants : Loïc Hélouët, Rituraj Singh

Realization of complex workflows in a crowdsourcing environment. Crowdsourcing consists in hiring workers on internet to perform large amounts of simple, independent and replicated work units, before assembling the returned results. A challenge to solve intricate problems is to define orchestrations of tasks, and allow higher-order answers where workers can suggest a process to obtain data rather than a plain answer. Another challenge is to guarantee that an orchestration with correct input data terminates, and produces correct output data.

In [25], we have proposed complex workflows, a data-centric model for crowdsourcing based on orchestration of concurrent tasks and higher order schemes. We will consider termination (whether some/all runs of a complex workflow terminate) and correctness (whether some/all runs of a workflow terminate with data satisfying FO requirements). We show that existential termination/correctness are undecidable in general excepted for specifications with bounded recursion. However, universal termination/correctness are decidable when constraints on inputs are specified in a decidable fragment of FO, and are at least in co-2EXPTIME .

Quality improvement in Crowdsourcing. In [29], we have addressed quality issues in crowdsourcing. Crowdsourcing is a way to solve problems that need human contribution. Crowdsourcing platforms distribute replicated tasks to workers, pay them for their contribution, and aggregate answers to produce a reliable conclusion. A fundamental problem is to infer a correct answer from the set of returned results. Another challenge is to obtain a reliable answer at a reasonable cost: unlimited budget allows hiring experts or large pools of workers for each task but a limited budget forces to use resources at best. This paper considers crowdsourcing of simple boolean tasks. We first define a probabilistic inference technique, that considers difficulty of tasks and expertise of workers when aggregating answers. We then propose CrowdInc, a greedy algorithm that reduce the cost needed to reach a consensual answer. CrowdInc distributes resources dynamically to tasks according to their difficulty. We show on several benchmarks that CrowdInc achieves good accuracy, reduces costs, and we compare its performance to existing solutions.

In [36], we have extended this approach to quality and cost improvement for complex tasks realized on a crowdsourcing platform. We have designed new algorithms to replicate, distribute tasks and assemble the returned results during the realization of a complex workflow orchestrating a data production process. The algorithm is dynamic, and optimizes cost and confidence in aggregated data to decide whether tasks should be replicated. We have shown that our algorithm exhibits better performance than static allocation of tasks to crowdworkers, both in terms of costs and accuracy.

6.4.2 Guarded Attribute Grammars

Participants : Adrian Puerto Aubel, Éric Badouel

In [10], we address the problem of component reuse in the context of service-oriented programming and more specifically for the design of user-centric distributed collaborative systems modelled by Guarded Attribute Grammars. Following the contract-based specification of components we developed an approach to an interface theory for the components of a collaborative system in three stages: we define a composition of interfaces that specifies how the component behaves with respect to its environment, we introduce an implementation order on interfaces and finally a residual operation on interfaces characterizing the systems that, when composed with a given component, can complement it in order to realize a global specification.

6.5 New results on Transversal Concern: Missing Models

Participants : Sihem Cherrared, Éric Fabre, Blaise Genest, Léo Henry, Thierry Jérón, Nicolas Markey

Learning discrete-time Markov chains. When models are missing, learning models from observations of a system is a powerful tool. In [3], we consider learning Discrete Time Markov Chains (DTMC), with different AI methods such as frequency estimation or Laplace smoothing. While models learnt with such methods converge asymptotically towards the exact system, a more practical question in the realm of trusted machine learning is how accurate a model learnt with a limited time budget is. Existing approaches provide bounds on how close the model is to the original system, in terms of bounds on local (transition) probabilities, which has unclear implication on the global behavior. In this work, we provide global bounds on the error made by such a learning process, in terms of global behaviors formalized using temporal logic. More precisely, we propose a learning process ensuring a bound on the error in the probabilities of these properties. While such learning process cannot exist for the full LTL logic, we provide one ensuring a bound that is uniform over all the formulas of CTL. Further, given one time-to-failure property, we provide an improved learning algorithm. Interestingly, frequency estimation is sufficient for the latter, while Laplace smoothing is needed to ensure non-trivial uniform bounds for the full CTL logic.

Learning timed automata. Active learning of timed languages is concerned with the inference of timed automata by observing some of the timed words in their languages. The learner can query for the membership of words in the language, or propose a candidate model and ask if it is equivalent to the target. The major difficulty of this framework is the inference of clock resets, which are central to the dynamics of timed automata but not directly observable.

Interesting first steps have already been made by restricting to the subclass of event-recording automata [46], where clock resets are tied to observations. In order to advance towards learning of general timed automata, in [27], we generalize this method to a new class, called reset-free event-recording automata, where some transitions may reset no clocks.

Central to our contribution is the notion of invalidity, and the algorithm and data structures to deal with it, allowing on-the-fly detection and pruning of reset hypotheses that contradict observations. This notion is a key to any efficient active-learning procedure for generic timed automata.

Learning models for telecommunication management. The fault diagnosis literature (about discrete event systems) is abundant for situations where system models *including faults and their consequences* are given as a starting point. When addressing real life applications such as fault diagnosis in telecommunication networks, one is immediately faced with the lack of models. Already for the normal behaviour

of these systems, not to mention faults and their effects. We have addressed this challenge from two angles. In a collaboration with Orange Labs, we have considered a self-modeling approach, that assemble generic component models in order to fit a network architecture, capturing the physical, virtualization, functional and service layers, at appropriate granularities. This allows a quick adaptation to changing configurations, an important feature of programmable networks. Models obtained by this approach can be validated by fault injections on platform mockups. They are then translated into an appropriate formalism (constraint networks or Bayesian networks) to serve as the basis for diagnosis engines. This work was at the heart of Sihem Cherrared's thesis, defended in June 2020 [32]. In another collaboration with Nokia Bell Labs [31], we explored a different approach aiming at detecting and characterizing soft performance degradations (instead of abrupt failures). This is done by detecting changes in the joint distribution of numerous performance indicators. We explored several machine learning approaches, using data collected on a platform running the real production software, with simulated traffic. Again, we used (soft) fault injections to characterize network behaviours under resource stress.

7 Bilateral contracts and grants with industry

7.1 Bilateral contracts with industry

Nokia Bell Labs - ADR SAPIENS. Several researchers of SUMO are involved in the joint research lab of Nokia Bell Labs France and Inria. We participate in the common research team SAPIENS (Smart Automated and Programmable Infrastructures for End-to-end Networks and Services), previously named "Softwarization of Everything." This team involves several other Inria teams: Convecs, Diverse and Spades. SUMO focuses on the management of reconfigurable systems, both at the edge (IoT based applications) and in the core (e.g. virtualized IMS systems). In particular, we study control and diagnosis issues for such systems. A PhD student is involved in the project: Abdul Majith (started in January 2019) on Controller Synthesis of Adaptive Systems, supervised by Hervé Marchand, Ocan Sankur and Dinh Thai Bui (Nokia Bell Labs).

Orange Labs. SUMO takes part in IOLab, the common lab of Orange Labs and Inria, dedicated to the design and management of Software Defined Networks. Our activities concern the diagnosis of malfunctions in virtualized multi-tenant networks.

This collaboration supported the Cifre PhD grant of Sihem Cherrared, supervised by Éric Fabre, Gregor Goessler (Inria Spades, Grenoble) and Sofiane Imadali (Orange Labs).

Mitsubishi Electric Research Center Europe (MERCE). Several researchers of SUMO are involved in a collaboration on the verification of real-time systems with the "Information and Network Systems (INS)" Team led by David Mentré of the "Communication & Information Systems (CIS)" Division of MERCE Rennes. The members of the team at MERCE work on different aspects of formal verification. Currently the SUMO team and MERCE jointly supervise a Cifre PhD student (Emily Clement) funded by MERCE since fall 2018; the thesis is about robustness of reachability in timed automata. Moreover Reiya Noguchi, a young engineer, member of MERCE, on leave of a Japanese operational division of Mitsubishi is also hosted and co-supervised by the SUMO team since the beginning of 2019, one day per week; we collaborate with him on the consistency of timed requirements.

8 Partnerships and cooperations

8.1 International initiatives

8.1.1 Inria International Labs

FUCHSIA

Title: Fuchsia

Duration: 2019 - 2021

Coordinator: Éric Badouel

Partners:

- ENSP, U. Yaoundé (Cameroon);
- OpenAgora;
- Epicentre.

Inria contact: Éric Badouel

Summary: Fuchsia is an associate team between Inria SUMO team in Rennes, ENSP (Ecole nationale Supérieure Polytechnique) in Yaoundé, Open Agora and Epicentre. The scientific objective of the team is to deploy flexible, adaptive and user-centric workflow systems on the Internet to enable groups to make smart decisions and coordinate their actions. The proposed solutions, based on the model of Guarded Attribute Grammars, should provide support for information gathering, deliberation and decision-making. Various applications will be considered: urban crowdsourcing, choreography of services, and crisis management systems in collaboration with Epicentre.

8.1.2 Inria Associate Team not involved in an IIL

EQUAVE

Title: *Efficient Quantitative Verification*

Duration: 2018 - 2020

Coordinator: Loïc Hélouët

Partners:

- Dpt of Computer Science and Engineering, Indian Institute of Technology Bombay (India)
- Dpt of Computer Science and Engineering, Indian Institute of Technology Dharwad (India)

Inria contact: Loïc Hélouët

Summary: Formal verification has been addressed for a long time. A lot of effort has been devoted to Boolean verification, i.e., formal analysis of systems that check whether a given property is true or false.

In many settings, a boolean verdict is not sufficient. The notions of interest are for instance the amount of confidential information leaked by a system, the proportion of some protein after a duration in some experiment in a biological system, whether a distributed protocol satisfies some property only for a bounded number of participants... This calls for quantitative verification, in which algorithms compute a value such as the probability for a property to hold, the mean cost of runs satisfying it, the time needed to achieve a complex workflow...

A second limitation of formal verification is the efficiency of algorithms. Even for simple questions, verification is rapidly PSPACE-complete. However, some classes of models allow polynomial time verification. The key techniques to master complexity are to use concurrency, approximation, etc

The objective of this project is to study efficient techniques for quantitative verification, and develop efficient algorithms for models such as stochastic games, timed and concurrent systems,

QuaSL

Title: Quantitative extensions of Strategy Logic

Duration: 2020 - 2022

Coordinator: Nicolas Markey

Partners:

- Team Astrea, University of Naples "Federico II" (Italy)

Inria contact: Nicolas Markey

Summary: Model checking aims at verifying that the executions of a computer system (usually modelled as a labelled transition system) satisfy a given property. Those properties are most often expressed using temporal logics, which provide a powerful way of specifying constraints on the occurrence of events for an execution to be valid. When reasoning about systems made of several components, we usually do not want to consider all executions: instead, we want to only consider those executions that can be triggered by some of the components as a reaction to the behaviours of other components. In an analogy with game theory (where players are components, executions are plays and valid behaviours correspond to winning conditions), temporal logics have been extended to reason about strategies; Strategy Logic can for instance express rich properties including antagonism and cooperation between groups of players. Our objectives in this project is to augment Strategy Logic with quantitative aspects: in that setting, properties are not true or false, but they take values reflecting the quality or efficiency of strategies and their associated executions. Checking such quantitative properties usually has very high complexity, if doable at all. Our recent works led to positive results, which we will extend in this associate team.

8.1.3 Inria International Partners

Informal international partners The team regularly collaborates with the following researchers:

- S. Akshay (IIT Bombay, India) on timed concurrent models;
- Andrea D’Ariano (University Roma Tre, Italy), on train regulation;
- Christel Baier (Technical University of Dresden, Germany) on verification and control of stochastic systems;
- Thomas Brihaye (Université de Mons, Belgium) on the verification of stochastic timed systems;
- Alessandro Giua and Michele Pinna (University Cagliari, Italy) on diagnosis and unfolding techniques for concurrent systems;
- Igor Konnov (Informal Systems, Austria), Marijana Lazić (Technical University Munich, Germany) and Josef Widder (Informal Systems, Austria) on the automated verification of randomized distributed algorithms;
- Stéphane Lafortune (University of Michigan, USA) on the control of cyber-physical systems;
- Kim G. Larsen (University Aalborg, Denmark) on quantitative timed games, and on topics related to urban train systems modeling;
- Laurie Ricker (Mount Allison University, Canada) on the enforcement of opacity;

8.2 International research visitors

8.2.1 Visits to international teams

Research stays abroad Léo Henry was awarded a grant from Rennes Métropole for a 2-month stay at the VUB (Brussels), in the team of Ann Nowé. The stay was planned to last from mid-february to mid-april, but it was stopped after one month because of the covid-19 pandemics.

8.3 National initiatives

ANR TickTac: Efficient Techniques for Verification and Synthesis of Real-Time Systems (2019-2023)

- [link to website](#)
- Led by Ocan Sankur (SUMO);
- SUMO participants: Emily Clement, Léo Henry, Thierry Jérón, Nicolas Markey, Victor Roussanaly, Ocan Sankur

- Partners: LSV/LMF (Paris-Saclay), ISIR (Paris), LaBRI (Bordeaux), LRDE (Paris), LIF (Marseille)

The aim of TickTac is to develop novel algorithms for the verification and synthesis of real-time systems using the timed automata formalism. One of the project's objectives is to develop an open-source and configurable model checker which will allow the community to compare algorithms. The algorithms and the tool will be used on a motion planning case study for robotics.

ANR HeadWork: Human-Centric Data-oriented WORKflows (2016-2022)

- [link to website](#)
- Led by David Gross-Amblard (Université Rennes 1);
- SUMO participants : Éric Badouel, Loïc Hérouët, Adrian Puerto Aubel, Rituraj Singh;
- Partners: Irisa team Druid (Rennes), Inria Project-Teams Valda (Paris), SUMO (Rennes) and Links (Lille), MNHN, Foule Factory.

The objective of this project is to develop techniques to facilitate development, deployment, and monitoring of crowd-based participative applications. This requires handling complex workflows with multiple participants, uncertainty in data collections, incentives, skills of contributors, ... To overcome these challenges, Headwork will define rich workflows with multiple participants, data and knowledge models to capture various kind of crowd applications with complex data acquisition tasks and human specificities. We will also address methods for deploying, verifying, optimizing, but also monitoring and adapting crowd-based workflow executions at run time.

IPL HAC-SPECIS: High-performance Application and Computers, Studying Performance and Correctness In Simulation (2016-2020)

- [link to website](#)
- Led by Arnaud Legrand (Inria Grenoble Rhône-Alpes)
- SUMO participant: Thierry Jéron.
- Partners: Inria project-teams Avalon (Lyon), POLARIS (Grenoble), HiePACS, STORM (Bordeaux), MExiCo (Saclay), MYRIADS, SUMO (Rennes), VeriDis (Nancy).

The Inria Project Lab HAC-SPECIS (High-performance Application and Computers, Studying Performance and Correctness In Simulation), is a transversal project internal to Inria. The goal of the HAC SPECIS project is to answer the methodological needs raised by the recent evolution of HPC architectures by allowing application and runtime developers to study such systems both from the correctness and performance point of view. Inside this project, we collaborate with Martin Quinson (Myriads team) on the dynamic formal verification of high performance runtimes and applications. The PhD of The Anh Pham, completed in December 2019, was granted by this project.

This year was the last year of the project, the closing meeting took place virtually in November 2020, in particular with a presentation of the achievements we got about verification of HPC using dynamic partial order methods. Since December 2019, Ehsan Azimi has been hired for a 2-year engineer position (ADT). His role is both to reinforce the basis of the verification engine in SimGrid, and to implement on top of it the results of The Anh Pham's PhD thesis and experiment them on HPC benchmarks.

National informal collaborations The team collaborates with the following researchers:

- Béatrice Bérard (LIP6, Paris 6) on problems of opacity and diagnosis, and on problems related to logics and partial orders for security;
- Patricia Bouyer (LMF, ENS Paris-Saclay) on the analysis of probabilistic timed systems and quantitative aspects of verification;
- Serge Haddad (Inria team MExiCo, LMF, ENS Paris-Saclay) on opacity and diagnosis;

- Didier Lime and Olivier H. Roux (LS2N, Université de Nantes) on stochastic and timed Petri nets;
- François Laroussinie (IRIF, U. Paris) on logics for multi-agent systems.

The team also has ongoing interactions with engineers at Alstom transports, originating from former industrial collaborations (in particular the [P22 project](#)).

9 Dissemination

9.1 Promoting scientific activities

9.1.1 Scientific events: organisation

General Chair, Scientific Chair.

- Nathalie Bertrand and Nicolas Markey are members of steering committee of the school MOVEP.
- Blaise Genest is in the Steering committee of FMAI: Formal Methods in AI.
- Hervé Marchand is a member of the IFAC Technical Committees (TC 1.3 on Discrete Event and Hybrid Systems). He is the president of the steering committee of MSR (modélisation de systèmes réactifs).

Member of the Organizing Committees.

- Éric Badouel was a member of the organization committee of CARI'20.

9.1.2 Scientific events: selection

Chair of Conference Program Committees.

- Nathalie Bertrand was PC co-chair of the 18th International Conference on Formal Modeling and Analysis of Timed Systems ([FORMATS 2020](#)).

Member of the Conference Program Committees.

- Éric Badouel was a member of the program committees of CARI'20 and VECOS'20;
- Nathalie Bertrand was a PC member of FORTE'20 and Highlights'20;
- Blaise Genest was in the PC of FMAI 2020, which has been postponed to 2021 (was supposed to be held in London in April 2020);
- Thierry Jérón served in the Program Committees of ICTSS'20 and SAC-SVT'20;
- Nicolas Markey was a PC member of LATA 2020-2021 and RV 2020;
- Ocan Sankur was in the Program Committees of AAAI 2020 and FORMATS 2020.
- Loïc Hélouët was member of the Program Committees of ACSD 2020 and Petri Nets 2020.

Reviewer. All members of the team regularly write reviews for the main conferences in our areas of expertise (LICS, ICALP, CAV, Concur, STACS, FoSSaCS, RV, WoDES, ...)

9.1.3 Journal

Member of the Editorial Boards.

- Éric Badouel is co-editor-in-Chief of ARIMA Journal.
- Hervé Marchand is associate editor of the journal *Discrete Event Dynamical Systems - Theory and applications* since january 2019.

Reviewer - Reviewing Activities. All members of the team regularly write reviews for the main journals in our areas of expertise (TCS, LMCS, Acta Informatica, DEDS, IEEE TAC, FMSD, SOSYM, Fundamenta Informaticae...)

9.1.4 Invited talks

- Nathalie Bertrand was an invited speaker at DisCoTec'20, June 2020, on Modeling and verifying randomized fault-tolerant distributed algorithms.
- Nathalie Bertrand was an invited speaker at MFCS'20, August 2020, on Concurrent games with arbitrary many players.
- Loïc Hélouët was invited to give a lecture on timed negotiations at the Chennai Mathematical Institute in February.

9.1.5 Leadership within the scientific community

- Nathalie Bertrand is co-head of the French working group on verification ("GT-Verif")

9.1.6 Scientific expertise

- Nathalie Bertrand served on the HCÉRES visit committee for Verimag in January 2020;
- Ocan Sankur reviewed a project for NWO, the Dutch research council;
- Éric Fabre served as an expert for the Ministry of Research (MESRI), in the CIR program (tax reduction for research activities of private companies and startups).

9.1.7 Research administration

- Éric Badouel is the co-director (with Amel Abda, Tunis) of LIRIMA, the Inria International Lab for Africa. He is scientific officer for the African and Middle-East region at Inria DRI (International Partnership Department). He is member of the executive board of GIS SARIMA.
- Thierry Jérón is a member of the Comité d'orientation scientifique (COS) of Inria Rennes. Since 2016 he is "réfèrent chercheur" for Inria Rennes;
- Hervé Marchand is an elected member of the "Comité de Centre" at Inria Rennes since June 2019;
- Loïc Hélouët is Member of the National Health and Security Committee of Inria. He is responsible for the Young researchers Mission at Inria Rennes. He is elected (suppletive) member of the "Comité de Centre" at Inria Rennes.

9.2 Teaching - Supervision - Juries

9.2.1 Teaching

- Licence: Nathalie Bertrand, Advanced Algorithms (ALGO2), 20h, L3, Univ Rennes 1, France;
- Licence: Loïc Hélouët, JAVA and algorithms, L2, 40h, INSA de Rennes, France.
- Master: Éric Badouel, Logic and argumentation, 32h, Univ Yaoundé I, Cameroon.
- Master: Nathalie Bertrand, Language Theory; Algorithms, 20h, Agrégation, ENS Rennes, France.
- Master: Éric Fabre, Models and Algorithms for Distributed Systems (MADS), 10h, M2, Univ Rennes 1, France;
- Master: Éric Fabre, Information Theory, 15h, M1, ENS Rennes, France.
- Master: Loïc Hélouët, Algorithms, 4h, Agrégation, ENS Rennes, France;

- Master: Loïc Hélouët, Algorithms and proof, 12h, Agrégation, ENS Rennes, France;
- Master: Nicolas Markey, Verification of Complex Systems (CSV), 10h, M2, Univ Rennes 1, France;
- Master: Nicolas Markey, Algorithms, 12h, Agrégation, ENS Rennes, France;
- Master: Nicolas Markey, Computability and Complexity, 18h, Agrégation, ENS Rennes, France;
- Master: Ocan Sankur, Verification of Complex Systems (CSV), 10h, M2, Univ Rennes 1, France;
- Master: Ocan Sankur, *Travaux pratiques*, Analyse et Conception Formelle (ACF), 26h, M1, Univ Rennes 1, France;
- Master: Ocan Sankur, Logic, 12h, Agrégation, ENS Rennes, France;

9.2.2 Supervision

PhD Students.

- PhD: Sihem Cherrared, Diagnosis of multi-tenant programmable networks, supervised by Éric Fabre, Gregor Goessler (Inria, Spades) and Sofiane Imadali (Orange), Université Rennes 1, June 2020.
- PhD: Victor Roussanaly, Efficient verification of real-time systems, supervised by Nicolas Markey and Ocan Sankur, Université Rennes 1, November 2020.
- PhD in progress: Emily Clement, Verification and synthesis of control systems: efficiency and robustness, started Dec. 2018, supervised by Thierry Jéron, Nicolas Markey, and David Mentré (Mitsubishi Electric)
- PhD in progress: Rodrigue Djeumen Djatcha, Collaborative Model for Urban Crowdsourcing, started in September 2017, University of Douala, Cameroon, supervised by Éric Badouel.
- PhD in progress: Léo Henry, Optimal test-case generation with game theory, started Oct. 2018, supervised by Thierry Jéron and Nicolas Markey.
- PhD in progress: Abdul Majith, Control of Adaptive Systems, started in Jan. 2019, supervised by Hervé Marchand, Ocan Sankur, and Dinh Thai-Bui (Nokia Bell Labs).
- PhD in progress: Anirban Majumdar, Games for distributed networks: models and algorithms, ENS Paris Saclay, started Sept 2018, supervised by Nathalie Bertrand and Patricia Bouyer (LSV/LMF, ENS Paris-Saclay).
- PhD in progress: Arthur Queffelec, Tradeoff between Robustness and Optimality in Strategic Reasoning, started Nov. 2018, supervised by Ocan Sankur and François Schwarzentruber (Logica, IriSa).
- PhD in progress: Suman Sadhukhan, Modelling and parameterized verification of mobile networks, started Oct. 2018, supervised by Nathalie Bertrand, Nicolas Markey and Ocan Sankur.
- PhD in progress: Rituraj Singh, Data-centric Workflows for Crowdsourcing Applications, started Feb. 2018, supervised by Loïc Hélouët.
- PhD in progress: Bastien Thomas, Automated verification of randomized distributed algorithms, started in Oct. 2019, supervised by Nathalie Bertrand and Josef Widder (Informal Systems, Austria).

Master Students.

- Léo Henry, Thierry Jéron and Nicolas Markey supervise (2h/week during 6 months, since October 2020) Pierre Bourse, a M1 student. The topic is machine learning extended to timed automata.

Other Internships.

- Mihir Vahanwala, IIT Bombay, internship (L3). Held virtually in May and June 2020, worked on the Skolem problem, considering decidability of a robust version.
- Ritika, IIT Bombay, internship (L3). Held virtually in May and June 2020, worked on polynomial solutions for min/max duration of timed negotiations.

9.2.3 Juries

PhD Defenses.

- Nathalie Bertrand took part in the following PhD committees (all held online): Mathieu Lehaut (LIP6, Paris; December 2020), Dimitri Antakly (LS2N, Nantes; July 2020), Stephan Plassart (LIG & Inria, Grenoble; June 2020), Marion Hallet (Université Libre de Bruxelles, Belgium; May 2020);
- Blaise Genest was reviewer for the thesis of Sukanya Basu (IIT Bombay, August 2020) and of Marie Fortin (LSV, ENS Paris-Saclay, November 2020);
- Nicolas Markey was a reviewer for the PhD thesis of Ian Cassar (University of Malta and Reykjavik University, nov. 2020).

Other Juries.

- Nathalie Bertrand was a committee member for two *maître de conférences* positions, at Université de Paris/IRIF, and at Université Paris-Est Créteil/LACL;
- Thierry Jérón was in the hiring committee for CRCN and ISFP positions at Inria Rennes Bretagne Atlantique;
- Nicolas Markey was in the assessment committee for an assistant-professor position at Aalborg University (Denmark);
- Ocan Sankur was in the hiring committee for a *maître de conférences* position at Université Aix-Marseille in 2020.

9.3 Popularization

9.3.1 Education

Several members of the team took a training course on the "[chiche project](#)", whose aim is to organize visits of researchers in secondary schools and high schools in order to present their research (in simple words) and popularize research in computer science.

10 Scientific production

10.1 Major publications

- [1] E. Badouel, L. Bernardinello and P. Darondeau. *Petri Net Synthesis*. Text in Theoretical Computer Science, an EATCS Series. Springer, Nov. 2015, p. 339. DOI: [10.1007/978-3-662-47967-4](https://doi.org/10.1007/978-3-662-47967-4). URL: <https://hal.inria.fr/hal-01237142>.
- [2] C. Baier, N. Bertrand, C. Dubsloff, D. Gburek and O. Sankur. ‘Stochastic Shortest Paths and Weight-Bounded Properties in Markov Decision Processes’. In: *LICS '18 - 33rd Annual ACM/IEEE Symposium on Logic in Computer Science*. Oxford, United Kingdom: ACM Press, July 2018, pp. 86–94. DOI: [10.1145/3209108.3209184](https://doi.org/10.1145/3209108.3209184). URL: <https://hal.archives-ouvertes.fr/hal-01883409>.
- [3] H. Bazille, B. Genest, C. Jegourel and J. Sun. ‘Global PAC Bounds for Learning Discrete Time Markov Chains’. In: *CAV 2020*. Vol. LNCS. CAV 2020 12225. Los Angeles, United States, 2020, pp. 304–326. URL: <https://hal.archives-ouvertes.fr/hal-03065571>.

- [4] N. Bertrand, M. Dewaskar, B. Genest, H. Gimbert and A. Godbole. ‘Controlling a population’. In: *Logical Methods in Computer Science* 15.3 (2019), pp. 1–30. DOI: [10.23638/LMCS-15\(3:6\)2019](https://doi.org/10.23638/LMCS-15(3:6)2019). URL: <https://hal.archives-ouvertes.fr/hal-02350251>.
- [5] P. Bouyer, U. Fahrenberg, K. G. Larsen, N. Markey, J. Ouaknine and J. Worrell. ‘Model Checking Real-Time Systems’. In: *Handbook of model checking*. Springer-Verlag, Apr. 2018, pp. 1001–1046. DOI: [10.1007/978-3-319-10575-8_29](https://doi.org/10.1007/978-3-319-10575-8_29). URL: <https://hal.archives-ouvertes.fr/hal-01889280>.
- [6] E. Fabre, L. Hélouët, E. Lefauchaux and H. Marchand. ‘Diagnosability of Repairable Faults’. In: *13th International Workshop on Discrete Event Systems*. (Version Longue). Xi’an, China, 2016, pp. 256–262. URL: <https://hal.inria.fr/hal-01302562>.
- [7] S. Pinisetty, Y. Falcone, T. Jérón and H. Marchand. ‘Runtime Enforcement of Regular Timed Properties’. In: *Software Verification and Testing, track of the Symposium on Applied Computing ACM-SAC 2014*. Gyeongju, South Korea: ACM, Mar. 2014, pp. 1279–1286. URL: <https://hal.inria.fr/hal-00907571>.

10.2 Publications of the year

International journals

- [8] S. Akshay, L. Hélouët and R. Phawade. ‘Combining Free choice and Time in Petri Nets’. In: *Journal of Logical and Algebraic Methods in Programming* (18th May 2020), pp. 1–36. DOI: [10.1016/j.jlamp.2018.11.006](https://doi.org/10.1016/j.jlamp.2018.11.006). URL: <https://hal.inria.fr/hal-01931728>.
- [9] É. André, D. Lime and N. Markey. ‘Language Preservation Problems in Parametric Timed Automata’. In: *Logical Methods in Computer Science* 16.1 (22nd Jan. 2020). DOI: [10.23638/LMCS-16](https://doi.org/10.23638/LMCS-16). URL: <https://hal.archives-ouvertes.fr/hal-02498022>.
- [10] E. Badouel and R. A. Djeumen Djatcha. ‘A Calculus of Interfaces for Distributed Collaborative Systems: The Guarded Attribute Grammar Approach’. In: *Revue Africaine de la Recherche en Informatique et Mathématiques Appliquées*. Special issue CARI 2018 31 (5th Oct. 2020), pp. 1–16. URL: <https://hal.inria.fr/hal-02145920>.
- [11] O. Beaumont, T. Lambert, L. Marchal and B. Thomas. ‘Performance Analysis and Optimality Results for Data-Locality Aware Tasks Scheduling with Replicated Inputs’. In: *Future Generation Computer Systems* 111 (Oct. 2020), pp. 582–598. DOI: [10.1016/j.future.2019.08.024](https://doi.org/10.1016/j.future.2019.08.024). URL: <https://hal.inria.fr/hal-02275473>.
- [12] N. Bertrand. ‘Model checking randomized distributed algorithms’. In: *ACM SIGLOG News* 7.1 (24th Feb. 2020), pp. 35–45. DOI: [10.1145/3385634.3385638](https://doi.org/10.1145/3385634.3385638). URL: <https://hal.inria.fr/hal-03095637>.
- [13] N. Bertrand, S. Haddad and E. Lefauchaux. ‘Diagnosis and Degradation Control for Probabilistic Systems’. In: *Discrete Event Dynamic Systems* 30.4 (Dec. 2020), pp. 695–723. DOI: [10.1007/s10626-020-00320-2](https://doi.org/10.1007/s10626-020-00320-2). URL: <https://hal.inria.fr/hal-03095652>.
- [14] T. Charrier, A. Queffelec, O. Sankur and F. Schwarzentruber. ‘Complexity of planning for connected agents’. In: *Autonomous Agents and Multi-Agent Systems* 34.2 (Oct. 2020). DOI: [10.1007/s10458-020-09468-5](https://doi.org/10.1007/s10458-020-09468-5). URL: <https://hal.archives-ouvertes.fr/hal-03003144>.
- [15] P. Gardy, P. Bouyer and N. Markey. ‘Dependences in Strategy Logic’. In: *Theory of Computing Systems* 64.3 (Apr. 2020), pp. 467–507. DOI: [10.1007/s00224-019-09926-y](https://doi.org/10.1007/s00224-019-09926-y). URL: <https://hal.archives-ouvertes.fr/hal-02917561>.
- [16] M. Pichené, S. K. Palaniappan, E. Fabre and B. Genest. ‘Modeling Variability in Populations of Cells using Approximated Multivariate Distributions’. In: *IEEE/ACM Transactions on Computational Biology and Bioinformatics* 17.5 (2020), pp. 1691–1702. DOI: [10.1109/TCBB.2019.2904276](https://doi.org/10.1109/TCBB.2019.2904276). URL: <https://hal.archives-ouvertes.fr/hal-02350249>.
- [17] G. M. M. Pinna and E. Fabre. ‘Spreading nets: a uniform approach to unfoldings’. In: *Journal of Logical and Algebraic Methods in Programming* 112 (Apr. 2020), p. 100526. DOI: [10.1016/j.jlamp.2020.100526](https://doi.org/10.1016/j.jlamp.2020.100526). URL: <https://hal.archives-ouvertes.fr/hal-03130461>.

International peer-reviewed conferences

- [18] S. Akshay, B. Genest, L. Helouët and S. Mital. ‘Timed Negotiations’. In: *Lecture Notes in Computer Science*. FOSSACS 2020 - 23rd International Conference on Foundations of Software Science and Computation Structures. Foundations of Software Science and Computation Structures - 23rd International Conference, FOSSACS 2020 12077. Dublin (physical conference cancelled), Ireland, 20th Apr. 2020, pp. 1–36. DOI: [10.1007/978-3-030-45231-5_3](https://doi.org/10.1007/978-3-030-45231-5_3). URL: <https://hal.inria.fr/hal-02964695>.
- [19] H. Bazille, B. Genest, C. Jegourel and J. Sun. ‘Global PAC Bounds for Learning Discrete Time Markov Chains’. In: CAV 2020. Vol. LNCS. CAV 2020 12225. Los Angeles, United States, 2020, pp. 304–326. URL: <https://hal.archives-ouvertes.fr/hal-03065571>.
- [20] N. Bertrand. ‘Concurrent Games with Arbitrarily Many Players’. In: 45th International Symposium on Mathematical Foundations of Computer Science (MFCS 2020). Prag, France, Aug. 2020. URL: <https://hal.inria.fr/hal-03095629>.
- [21] N. Bertrand, P. Bouyer and A. Majumdar. ‘Synthesizing safe coalition strategies’. In: FSTTCS’20. 40th IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science. on line, India, 1st Dec. 2020. URL: <https://hal.archives-ouvertes.fr/hal-02917538>.
- [22] N. Bertrand, M. Lazić and J. Widder. ‘A Reduction Theorem for Randomized Distributed Algorithms Under Weak Adversaries’. In: 22nd International Conference on Verification, Model Checking, and Abstract Interpretation (VMCAI’21). Copenhagen, Denmark, 12th Jan. 2021, pp. 219–239. DOI: [10.1007/978-3-030-67067-2_11](https://doi.org/10.1007/978-3-030-67067-2_11). URL: <https://hal.inria.fr/hal-03150397>.
- [23] N. Bertrand, N. Markey, S. Sadhukhan and O. Sankur. ‘Dynamic network congestion games’. In: FSTTCS 2020 - 40th IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science. Vol. 182. LIPIcs 40. Goa, India, 14th Dec. 2020. URL: <https://hal.archives-ouvertes.fr/hal-02980833>.
- [24] M. Blondin, J. Esparza, B. Genest, M. Helfrich and S. Jaax. ‘Succinct Population Protocols for Presburger Arithmetic’. In: STACS 2020. Vol. LIPICS. STACS 2020 154. Montpellier, France, 2020, 40:1–40:15. DOI: [10.4230/LIPICS.STACS.2020.40](https://doi.org/10.4230/LIPICS.STACS.2020.40). URL: <https://hal.archives-ouvertes.fr/hal-03065533>.
- [25] P. Bourhis, L. Hérouët, Z. Miklos and R. Singh. ‘Data Centric Workflows for Crowdsourcing’. In: Petri Nets 2020 - 41st International Conference on Application and Theory of Petri Nets and Concurrency. Paris, France, 24th June 2020, pp. 1–39. URL: <https://hal.inria.fr/hal-02508838>.
- [26] E. Clement, T. Jérón, N. Markey and D. Mentré. ‘Computing maximally-permissive strategies in acyclic timed automata’. In: FORMATS 2020 - 18th International Conference on Formal Modeling and Analysis of Timed Systems. Vienna, Austria, 1st Sept. 2020, pp. 1–33. URL: <https://hal.archives-ouvertes.fr/hal-02895344>.
- [27] L. Henry, T. Jérón and N. Markey. ‘Active learning of timed automata with unobservable resets’. In: FORMATS 2020 - 18th International Conference on Formal Modeling and Analysis of Timed Systems. Vienna, Austria, 1st Sept. 2020, pp. 1–26. URL: <https://hal.inria.fr/hal-02896517>.
- [28] T. Jérón, N. Markey, D. Mentré, R. Noguchi and O. Sankur. ‘Incremental methods for checking real-time consistency’. In: FORMATS 2020 - 18th International Conference on Formal Modeling and Analysis of Timed Systems. Vienne, Austria, 1st Sept. 2020, pp. 1–18. URL: <https://hal.inria.fr/hal-02894772>.
- [29] R. Singh, L. Hérouët and Z. Miklos. ‘Reducing the Cost of Aggregation in Crowdsourcing’. In: *International Conference on Web Services*. Web Services - ICWS 2020 - 27th International Conference, Held as Part of the Services Conference Federation, SCF 2020. Vol. 12406. Proceedings of Web Services - ICWS 2020 - 27th International Conference, Held as Part of the Services Conference Federation SCF 2020. Honolulu, United States, 20th Sept. 2020, pp. 77–95. DOI: [10.1007/978-3-030-59618-7](https://doi.org/10.1007/978-3-030-59618-7). URL: <https://hal.inria.fr/hal-02964718>.

- [30] G. Zinck, L. S. L. Ricker, H. Marchand and L. Hélouët. ‘Enforcing Opacity in Modular Systems’. In: IFAC 2020 - Ifac world Congress. Virtual, Germany, 30th Nov. 2020, pp. 1–8. URL: <https://hal.inria.fr/hal-02917644>.

Conferences without proceedings

- [31] A. Elmajed, A. Aghasaryan and E. Fabre. ‘Machine Learning Approaches to Early Fault Detection and Identification in NFV Architectures’. In: NetSoft 2020 - 6th IEEE International Conference on Network Softwarization. Ghent, France, 29th June 2020, pp. 200–208. DOI: [10.1109/NetSoft48620.2020.9165361](https://doi.org/10.1109/NetSoft48620.2020.9165361). URL: <https://hal.archives-ouvertes.fr/hal-03129396>.

Doctoral dissertations and habilitation theses

- [32] S. Cherrared. ‘Fault management of programmable multi-tenant networks’. Université Rennes 1, 26th June 2020. URL: <https://tel.archives-ouvertes.fr/tel-03047092>.
- [33] V. Roussanaly. ‘Efficient verification of real time systems’. Université Rennes 1, 30th Nov. 2020. URL: <https://hal.archives-ouvertes.fr/tel-03121385>.

Reports & preprints

- [34] S. Akshay, B. Genest, L. Hélouët, S. Krishna and S. Roychowdhury. *Resilience of Timed Systems*. Dec. 2020. DOI: [10.1145/nnnnnnn.nnnnnnn](https://doi.org/10.1145/nnnnnnn.nnnnnnn). URL: <https://hal.inria.fr/hal-03129402>.
- [35] L. Hélouët, N. Markey and R. Raha. *Reachability games with relaxed energy constraints*. 12th Oct. 2020. URL: <https://hal.inria.fr/hal-02964729>.
- [36] L. Hélouët, Z. Miklos and R. Singh. *Cost and Quality Assurance in Crowdsourcing Workflows*. 12th Oct. 2020. URL: <https://hal.inria.fr/hal-02964736>.
- [37] R. Meira-Góes, S. Lafortune and H. Marchand. *Synthesis of Supervisors Robust Against Sensor Deception Attacks*. 2020. URL: <https://hal.inria.fr/hal-03116690>.

Other scientific publications

- [38] B. W. Watson, E. Badouel and O. Niang. *Foreword of Proceedings of CARI 2020*. Thiès, Senegal, 14th Oct. 2020. URL: <https://hal.inria.fr/hal-02926115>.

10.3 Cited publications

- [39] E. Badouel and U. Schlachter. ‘Incremental Process Discovery using Petri Net Synthesis’. In: *Fundamenta Informaticae* 154.1-4 (June 2017), pp. 1–13. DOI: [10.3233/FI-2017-1548](https://doi.org/10.3233/FI-2017-1548). URL: <https://hal.inria.fr/hal-01599760>.
- [40] N. Bertrand, P. Bouyer, T. Brihaye and P. Carrier. ‘Analysing Decisive Stochastic Processes’. In: *ICALP 2016 - 43rd International Colloquium on Automata, Languages, and Programming*. Vol. 55. LiPIcs. Rome, Italy: LZI, 2016, 101:1–101:14. DOI: [10.4230/LIPIcs.ICALP.2016.101](https://doi.org/10.4230/LIPIcs.ICALP.2016.101). URL: <https://hal.inria.fr/hal-01397794>.
- [41] N. Bertrand, P. Bouyer, T. Brihaye and P. Carrier. ‘When are stochastic transition systems tameable?’ In: *Journal of Logical and Algebraic Methods in Programming* 99 (2018), pp. 41–96. DOI: [10.1016/j.jlamp.2018.03.004](https://doi.org/10.1016/j.jlamp.2018.03.004). URL: <https://hal.inria.fr/hal-01938135>.
- [42] P. Bouyer, N. Markey, N. Perrin and P. Schlehber-Caissier. ‘Timed automata abstraction of switched dynamical systems using control funnels’. In: *Real-Time Systems* 53.3 (May 2017), pp. 327–353. DOI: [10.1007/s11241-016-9262-3](https://doi.org/10.1007/s11241-016-9262-3). URL: <http://dx.doi.org/10.1007/s11241-016-9262-3>.
- [43] J. Desel, J. Esparza and P. Hoffmann. ‘Negotiation as concurrency primitive’. In: *Acta Informatica* 56.2 (2019), pp. 93–159. DOI: [10.1007/s00236-018-0318-9](https://doi.org/10.1007/s00236-018-0318-9). URL: <https://doi.org/10.1007/s00236-018-0318-9>.

- [44] J. Esparza and J. Desel. ‘On Negotiation as Concurrency Primitive’. In: *CONCUR 2013 - Concurrency Theory - 24th International Conference, CONCUR 2013, Buenos Aires, Argentina, August 27-30, 2013. Proceedings*. Ed. by P. R. D’Argenio and H. C. Melgratti. Vol. 8052. Lecture Notes in Computer Science. Springer, 2013, pp. 440–454. DOI: [10.1007/978-3-642-40184-8_31](https://doi.org/10.1007/978-3-642-40184-8_31). URL: https://doi.org/10.1007/978-3-642-40184-8_31.
- [45] J. Esparza, A. Muscholl and I. Walukiewicz. ‘Static analysis of deterministic negotiations’. In: *32nd Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2017, Reykjavik, Iceland, June 20-23, 2017*. IEEE Computer Society, 2017, pp. 1–12. DOI: [10.1109/LICS.2017.8005144](https://doi.org/10.1109/LICS.2017.8005144). URL: <https://doi.org/10.1109/LICS.2017.8005144>.
- [46] O. Grinchtein, B. Jonsson and M. Leucker. ‘Learning of event-recording automata’. In: *Theor. Comput. Sci.* 411.47 (2010), pp. 4029–4054. DOI: [10.1016/j.tcs.2010.07.008](https://doi.org/10.1016/j.tcs.2010.07.008). URL: <https://doi.org/10.1016/j.tcs.2010.07.008>.
- [47] I. V. Konnov, M. Lazic, H. Veith and J. Widder. ‘A short counterexample property for safety and liveness verification of fault-tolerant distributed algorithms’. In: *POPL 2017 - 44th ACM SIGPLAN Symposium on Principles of Programming Languages*. ACM, 2017, pp. 719–734.
- [48] A. Nigam and N. S. Caswell. ‘Business artifacts: An approach to operational specification’. In: *IBM Systems Journal* 42.3 (2003), pp. 428–445.