



Activity Report 2023

Team SUMO

SUpervision of large MOdular and distributed systems

Joint team with Centre Inria de l'Université de Rennes

D4 – Language and Software Engineering



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 Social and environmental responsibility	8
5.1 Footprint of research activities	8
6 Highlights of the year	8
6.1 Highlights	8
6.2 Awards	8
7 New software, platforms, open data	8
7.1 New software	8
7.1.1 MOCHY	8
7.1.2 PyLTA	9
8 New results	9
8.1 New results on Axis 1: Quantitative models	9
8.1.1 Learning for verification of timed systems	9
8.1.2 Runtime verification of timed systems	9
8.1.3 Verification of Real-Time Models	10
8.1.4 Quantitative analysis of timed systems	10
8.1.5 Separation of control and timing constraints in timed concurrent systems	10
8.1.6 Fuzzy logics for games	10
8.2 New results on Axis 2: Large Systems Models	11
8.2.1 Planning Problems	11
8.2.2 Regulation in Metro Systems	11
8.3 New results on Axis 3: Population Models	11
8.3.1 Verification of Distributed Algorithms	11
8.3.2 Network Congestion Games	12
8.3.3 Synchronization	12
9 Bilateral contracts and grants with industry	12
9.1 Bilateral contracts with industry	12

10 Partnerships and cooperations	13
10.1 International initiatives	13
10.1.1 Inria associate team not involved in an IIL or an international program	13
10.1.2 Participation in other International Programs	14
10.2 National initiatives	14
11 Dissemination	16
11.1 Promoting scientific activities	16
11.1.1 Scientific events: selection	16
11.1.2 Journal	16
11.1.3 Invited talks	16
11.1.4 Leadership within the scientific community	16
11.1.5 Research administration	16
11.2 Teaching - Supervision - Juries	17
11.2.1 Teaching	17
11.2.2 Supervision	17
11.2.3 Juries	18
11.2.4 Internal or external Inria responsibilities	18
11.2.5 Interventions	18
12 Scientific production	18
12.1 Major publications	18
12.2 Publications of the year	19
12.3 Other	20
12.4 Cited publications	20

Project-Team SUMO

Creation of the 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 Fabre [INRIA, Senior Researcher, HDR]
- Loic Hélouët [INRIA, Senior Researcher, HDR]
- Thierry Jéron [INRIA, Senior Researcher, HDR]
- Nicolas Markey [CNRS, Senior Researcher, HDR]
- Ocan Sankur [CNRS, Researcher]

PhD Students

- Aymeric Come [INRIA]
- Antoine Thebault [Alstom Transport, CIFRE]
- Bastien Thomas [UNIV RENNES I, until May 2023]
- Nicolas Waldburger [UNIV RENNES I]

Interns and Apprentices

- Hugo Francon [ENS RENNES, Intern, from Sep 2023]
- Pranav Sanjay Ghorpade [INRIA, Intern, from May 2023 until Jul 2023]
- Joseph Lenormand [UNIV PARIS SACLAY, Intern, from Jun 2023 until Jul 2023]
- Thomas Roux [INRIA, Intern, from Jun 2023 until Jul 2023]

Administrative Assistant

- Laurence Dinh [INRIA]

Visiting Scientists

- Ayush Agarwal [IIT BOMBAY, from May 2023 until Jul 2023]
- Prerak Contractor [IIT BOMBAY, from May 2023 until Jul 2023]
- Anand Narasimhan [IIT BOMBAY, from May 2023 until Jul 2023]

External Collaborator

- Ulrich Fahrenberg [EPITA]

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 [27]). 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 [25, 26] 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 (Network Function Virtualization, SDN/NFV) or at the edge (Internet of Things, 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 [28] 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 [29] 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 [24] 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 Social and environmental responsibility

5.1 Footprint of research activities

The Sumo team participated to the Extended Stay Support Scheme (ESSS) set up by TCS4F (Theoretical Computer Science for Future) during the international event that took place in Kassel in summer 2022. The ESSS aimed at exploiting the presence distant researchers (from Asia, America, etc) at these conferences to invite them over for a longer stay in France or Europe. More generally, some members of the team individually take part to the TCS4F initiatives.

Since covid, the carbon footprint of travels related to our research activities decreased a lot. Some members of the team no longer fly to attend conferences, and carefully choose the conferences where they submit not only in terms of reputation, but also taking into account the location or the possibility to attend online.

6 Highlights of the year

6.1 Highlights

- Ocan Sankur successfully defended his habilitation in October [21].
- After 11 fruitful years, SUMO stops at the end of December 2023!

6.2 Awards

Uli Fahrenberg received the [Petri Nets 2023 Outstanding Paper Award](#)

7 New software, platforms, open data

7.1 New software

7.1.1 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. Mochy addresses mainly models of transport networks, their timetables and traffic management techniques. The tool allows for the fast simulation of these regulated models, to measure performance indicators. Since version 2.0, MOCHY proposes a novel traffic management algorithm with neural networks.

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. Traffic management with neural networks.

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

Authors: Loic Helouet, Antoine Thebault, Didier Vojtisek

Contact: Loic Helouet

7.1.2 PyLTA

Keywords: Model Checking, Distributed computing

Functional Description: PyLTA is written in Python. It uses an ad hoc input format to represent distributed algorithms and their specification. The verification process builds on counter-example guided abstraction refinement (CEGAR) principle.

URL: <https://gitlab.com/BastienT/pylta>

Publication: hal-03996060

Contact: Nathalie Bertrand

Participants: Bastien Thomas, Ocan Sankur, Nathalie Bertrand

8 New results

8.1 New results on Axis 1: Quantitative models

8.1.1 Learning for verification of timed systems

Participants: Ocan Sankur.

In [16] present algorithms for model checking and controller synthesis of timed automata, seeing a timed automaton model as a parallel composition of a large finite-state machine and a relatively smaller timed automaton, and using compositional reasoning on this composition. We use automata learning algorithms to learn finite automata approximations of the timed automaton component, in order to reduce the problem at hand to finite-state model checking or to finite-state controller synthesis. We present an experimental evaluation of our approach.

8.1.2 Runtime verification of timed systems

Participants: Thierry Jéron.

Runtime Enforcement (RE) is a monitoring technique aimed at correcting possibly incorrect executions w.r.t. a set of formal requirements (properties) of a system. In this paper, we consider enforcement monitoring of real-time properties. Thus, executions are modelled as timed words and specifications as timed automata. Moreover, we consider that the enforcer has the ability to delay events by storing or buffering them into its internal memory (and releasing them when the property is finally satisfied) and suppressing events when no delaying is appropriate. Practically, in an implementation, the internal memory of the enforcer is finite. In [17], we propose a new RE paradigm for timed properties, where the memory of the enforcer is bounded/finite, to address practical applications with memory constraints and timed specifications. Bounding the memory presents a number of difficulties, e.g., how to accommodate a timed event into the memory when the memory is full, s.t., regardless of the course of action we choose to handle this situation, the behaviour of the bounded enforcer should not significantly differ from that of the unbounded enforcer. The problem of how to optimally discard events when the buffer is full is significantly more difficult in a timed environment where the progress of time affects the satisfaction or violation of a property. We define the bounded-memory RE problem for timed properties and develop a framework for regular timed properties specified as timed automata. The proposed framework is implemented in Python, and its performance is evaluated. From experiments, we discovered that the enforcer has a reasonable execution time overhead.

8.1.3 Verification of Real-Time Models

Participants: Thierry Jéron, Nicolas Markey, Ocan Sankur.

The patent [23] describes a method and system for correcting the operation of a target computer system by using timed requirements. In particular, we show how to modify real-time requirements in order to render them consistent.

8.1.4 Quantitative analysis of timed systems

Participants: Uli Fahrenberg.

In [14], we show how to efficiently solve energy Büchi problems in finite weighted automata and in one-clock weighted timed automata. Solving the former problem is our main contribution and is handled by a modified version of Bellman-Ford interleaved with Couvreur's algorithm. The latter problem is handled via a reduction to the former relying on the corner-point abstraction. All our algorithms are freely available and implemented in a tool based on the open-source platforms TChecker and Spot.

8.1.5 Separation of control and timing constraints in timed concurrent systems

Participants: Loïc Hélouët.

In time Petri nets (TPNs), time and control are tightly connected: time measurement for a transition starts only when all resources needed to fire it are available. Further, upper bounds on duration of enabledness can force transitions to fire (this is called urgency). For many systems, one wants to decouple control and time, i.e. start measuring time as soon as a part of the preset of a transition is filled, and fire it after some delay and when all needed resources are available. Our paper [10] considers an extension of TPN called waiting nets that dissociates time measurement and control. Their semantics allows time measurement to start with incomplete presets, and can ignore urgency when upper bounds of intervals are reached but all resources needed to fire are not yet available. Firing of a transition is then allowed as soon as missing resources are available. It is known that extending bounded TPNs with stopwatches leads to undecidability. Our extension is weaker, and we show how to compute a finite state class graph for bounded waiting nets, yielding decidability of reachability and coverability. We then compare expressiveness of waiting nets with that of other models w.r.t. timed language equivalence, and show that they are strictly more expressive than TPNs. In [22] we extend time Petri nets with continuous variables allowing for the description of objects with linear trajectories. The model is less expressive than hybrid systems, and its continuous part (time and variables) can be abstracted to a finite set of domains. For a subclass that is expressive enough to model the dynamics of metro networks, the verification of quantitative properties is decidable.

8.1.6 Fuzzy logics for games

Participants: Nicolas Markey.

Temporal logics are extensively used for the specification of on-going behaviors of computer systems. Two significant developments in this area are the extension of traditional temporal logics with modalities that enable the specification of on-going strategic behaviors in multi-agent systems, and the transition of temporal logics to a quantitative setting, where different satisfaction values enable the specifier to

formalize concepts such as certainty or quality. In the first class, SL (Strategy Logic) is one of the most natural and expressive logics describing strategic behaviors. In the second class, a notable logic is $LTL[\mathcal{F}]$, which extends LTL with quality operators. In [9], we introduce and study $SL[\mathcal{F}]$, which enables the specification of quantitative strategic behaviors. The satisfaction value of an $SL[\mathcal{F}]$ formula is a real value in $[0,1]$, reflecting “how much” or “how well” the strategic on-going objectives of the underlying agents are satisfied. We demonstrate the applications of $SL[\mathcal{F}]$ in quantitative reasoning about multi-agent systems, showing how it can express and measure concepts like stability in multi-agent systems, and how it generalizes some fuzzy temporal logics. We also provide a model-checking algorithm for $SL[\mathcal{F}]$, based on a quantitative extension of Quantified CTL*. Our algorithm provides the first decidability result for a quantitative extension of Strategy Logic. In addition, it can be used for synthesizing strategies that maximize the quality of the systems’ behavior.

8.2 New results on Axis 2: Large Systems Models

8.2.1 Planning Problems

Participants: Ocan Sankur.

The Connected Multi-Agent Path Finding (CMAPF) problem asks for a plan to move a group of agents in a graph while respecting a connectivity constraint. In [11], we study a generalization of CMAPF in which the graph is not entirely known in advance, but is discovered by the agents during their mission. We present a framework introducing this notion and study the problem of searching for a strategy to reach a configuration in this setting. We prove the problem to be PSPACE-complete when requiring all agents to be connected at all times, and NEXPTIME-hard in the decentralized case.

In [13], we prove the PSPACE-hardness of the problem when the underlying graph is a subgraph of a 3D grid and with range-based connectivity. Moreover, we provide an application of the WHCA* algorithm and show that it outperforms previously given algorithms by an order of magnitude in terms of the sizes of the instances it can handle.

8.2.2 Regulation in Metro Systems

Participants: Loïc Hélouët.

[15] introduces MOCHY, a tool designed for the modeling of concurrent systems with variants of stochastic, timed and hybrid Petri nets. Beyond modeling, the tool serves as a platform for fast simulation, and can be used for statistical verification of properties, controller testing, and learning of control rules. The targeted models are variants of stochastic and timed nets where tokens can be continuous quantities depicting trajectories of moving objects. The architecture of the tool is designed to be as adaptive as possible, and allow the redefinition of objects behaviors or transitions firing through the refinement of a few semantic rules. The framework also allows for the integration of controllers. For any model variant, MOCHY can perform fast simulation, and perform statistical verification, evaluate some quantitative properties of a model, or learn control rules for reachability or quantitative objectives.

8.3 New results on Axis 3: Population Models

8.3.1 Verification of Distributed Algorithms

Participants: Nathalie Bertrand, Nicolas Markey, Ocan Sankur, Bastien Thomas, Nicolas Waldburger.

Distributed algorithms typically run over arbitrary many processes and may involve unboundedly many rounds, making the automated verification of their correctness challenging. In the following papers, we developed parameterized verification techniques to assess the correctness of particular types of distributed algorithms.

In [18] we present the tool PyLTA, which can model check parameterized distributed algorithms against LTL specifications. The parameters typically include the number of processes and a bound on faulty processes, and the considered algorithms are round-based and either synchronous or asynchronous.

In [19] we consider the verification of distributed systems composed of an arbitrary number of asynchronous processes. Processes are identical finite-state machines that communicate by reading from and writing to a shared memory. Beyond the standard model with finitely many registers, we tackle round-based shared-memory systems with fresh registers at each round. In the latter model, both the number of processes and the number of registers are unbounded, making verification particularly challenging. The properties studied are generic presence reachability objectives, which subsume classical questions such as safety or synchronization by expressing the presence or absence of processes in some states. In the more general round-based setting, we establish that the parameterized verification of presence reachability properties is PSPACE-complete. Moreover, for the roundless model with finitely many registers, we prove that the complexity drops down to NP-complete and we provide several natural restrictions that make the problem solvable in polynomial time.

8.3.2 Network Congestion Games

Participants: Nicolas Markey.

Network congestion games are a simple model for reasoning about routing problems in a network. They are a very popular topic in algorithmic game theory, and a huge amount of results about existence and (in)efficiency of equilibrium strategy profiles in those games have been obtained over the last 20 years. In particular, the price of anarchy has been defined as an important notion for measuring the inefficiency of Nash equilibria. Generic bounds have been obtained for the price of anarchy over various classes of networks, but little attention has been put on the effective computation of this value for a given network. The invited talk [12] presents recent results on this problem in different settings.

8.3.3 Synchronization

Participants: Nathalie Bertrand, Nicolas Markey.

Synchronizing a (deterministic, finite-state) automaton is the problem of finding a sequence of actions to be played in the automaton in order to end up in the same state independently of the starting state. In [8], we consider synchronization with LTL constraints on the executions leading to synchronization, extending the results of [Petra Wolf. Synchronization under dynamic constraints. FSTTCS'20] by showing that the problem is PSPACE-complete for LTL as well as for restricted fragments (involving only modality F or G), while it is NP-complete for constraints expressed using only modality X.

9 Bilateral contracts and grants with industry

9.1 Bilateral contracts with industry

Nokia Bell Labs - ADR SAPIENS.

Participants: Éric Fabre, Ocan Sankur.

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 studied control and diagnosis issues for such systems.

Mitsubishi Electric Research Center Europe (MERCE).

Participants: Thierry Jéron, Nicolas Markey, Ocan Sankur.

Several researchers of SUMO are involved in a collaboration on the verification of real-time systems with the "Information and Network Systems" Team (INSv) 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. SUMO collaborates with Reiya Noguchi, a young engineer who was member of MERCE, on leave of a Japanese operational division of Mitsubishi and hosted by the SUMO team one day per week since during his stay in Rennes; Reiya returned in Japan in 2021 but we continue the collaboration with him and David Mentré in MERCE. This year we worked on the consistency of timed requirements and their fixing.

Orange Labs.

Participants: Éric Fabre.

SUMO takes part in I/O Lab, 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.

Alstom Transport

Participants: Loïc Hélouët, Antoine Thébault.

Loïc Hélouët is involved in a long-term collaboration with Alstom Transport, precisely with the ATS division (Automatic Train Supervision), with the objective to evaluate and improve regulation policies of urban train systems. Currently, Antoine Thébault pursues a PhD thesis on this topic.

10 Partnerships and cooperations

10.1 International initiatives

10.1.1 Inria associate team not involved in an IIL or an international program

QuaSL

Title: Quantitative Strategy Logic

Duration: 2020 -> 2024

Coordinator: Aniello Murano (murano@na.infn.it)

Partners:

- University of Naples "Federico II" (Italie)

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.

10.1.2 Participation in other International Programs

We are participating in the activities of the CNRS UMI ReLaX Research Lab in Computer Science, an Indo-French joint research unit dedicated to research in theoretical computer science, its applications and its interactions with mathematics. Within this setting, we are hosting undergraduate students as visitors or interns from Indian universities such as Chennai Mathematical Institute, and IIT Bombay.

We collaborate with Srinivas Pinisetty from Indian Institute of Technology Bhubaneswar on timed enforcement.

10.2 National initiatives

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

Participants: Thierry Jéron, Nicolas Markey, Ocan Sankur.

- [link to website](#)
- Led by Ocan Sankur (SUMO);
- Partners: 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 MAVeriQ: Methods of Analysis for Verification of Quantitative properties (2021-2025)

Participants: Nathalie Bertrand, Éric Fabre, Loïc Hélouët, Nicolas Markey.

- [link to website](#)
- Led by Aldric Degorre (IRIF); Local coordinator Éric Fabre.
- Partners: IRIF, LMF, Inria Rennes/IRISA, LACL, Verimag.

The objective of this project is to develop unified frameworks for quantitative verification of timed, hybrid, and stochastic systems. We believe such a unification is possible because common patterns are used in many cases. The project targets in particular:

- systematization of quantitative properties and their use cases
- substantial progress in the algorithms of quantitative verification;
- practical methodology for stating and verifying quantitative properties of systems.

The aim of MAVeriQ is to progress towards this unification, by gathering skills on timed and stochastic systems and on quantitative verification under a common roof, to jointly address open challenges in quantitative model-checking and quantitative validation. One such challenge we will address is robustness of quantitative models, that is, resilience to small perturbations, which is crucial for implementability. Unified methods developed in the project (such as robustness analysis and simulation techniques) will be showcased in different case studies in the domain of CPS (in particular automotive control), showing that such a system can be verified in different ways without leaving this framework.

ANR BisoUS: Better Synthesis for Underspecified Quantitative Systems (2023-2027)

Participants: Nathalie Bertrand, Loïc Hélouët, Nicolas Markey, Ocan Sankur.

- [link to website](#)
- Led by Didier Lime (LS2N); Local coordinator Nicolas Markey.
- Partners: LS2N, Inria Rennes/IRISA, LIPN, LMF

When designing complex and critical systems (planes, autonomous vehicles, etc.), it is crucial to be able to give guarantees that the system works as intended, which is often done through comprehensive testing. The goal of project BisoUS is to provide stronger guarantees, based on mathematics, and to detect problems as early as possible: solving them is then easier and cheaper.

Unfortunately this is a hard problem because some design choices may not have been done yet, and some key features (e. g. speed of a CPU) are then not known precisely enough.

In project BisoUS we develop formal methods, based on model-checking and synthesis to work with expressive modelling formalisms encompassing parameters, cost/rewards, and games on graphs to meet those challenges.

National informal collaborations The team collaborates with the following researchers:

- Patricia Bouyer (LMF, ENS Paris-Saclay) on quantitative aspects of verification and game models for parameterized systems;
- Yliès Falcone (University Grenoble-Alpes) and Victor Roussanaly (Inria Rhône Alpes) on the distributed timed monitoring.
- Serge Haddad (LMF, ENS Paris-Saclay) on resilience of timed systems;

11 Dissemination

Participants: Nathalie Bertrand, Éric Fabre, Loïc Hélouët, Thierry Jéron, Nicolas Markey, Ocan Sankur.

11.1 Promoting scientific activities

11.1.1 Scientific events: selection

Member of the conference program committees Nathalie Bertrand was on the program committees of FoSSaCS 2023, NETYS 2023 and Highlights 2023.

Thierry Jéron was on the program committee of SAC-SVT 2023.

Loïc Hélouët was on the program committee of Petri Nets 2023.

Nicolas Markey was on the program committee of EUMAS'23, TIME'23, CSL'23.

Ocan Sankur was on the program committee of FORMATS 2023, and on the artifact evaluation committee of ESOP 2023.

Reviewer - reviewing activities All members of the team reviewed a number of papers for international conferences.

11.1.2 Journal

Member of the editorial boards Nathalie Bertrand is an editorial board member for Journal of Logical and Algebraic Methods in Programming (JLAMP) and for Theoretical Computer Science (TCS).

Reviewer - reviewing activities All members of the team reviewed a number of papers for international journals.

11.1.3 Invited talks

Nicolas Markey gave an invited talk at Formats'23 [12].

11.1.4 Leadership within the scientific community

Nathalie Bertrand is the co-head of the French working group on verification of GDR-IM: GT Vérif.

11.1.5 Research administration

Éric Fabre is the co-director of the joint research lab of Nokia Bell Labs and Inria, and member of the Evaluation Committee of Inria;

Nicolas Markey is the head of Department "Language and Software Engineering" of IRISA.

11.2 Teaching - Supervision - Juries

11.2.1 Teaching

Although the team has no faculty members, most researchers do teach a significant number of hours, mostly at Masters level.

- Master: Nathalie Bertrand, Algorithms, and Symbolic AI, 18h, Agrégation, ENS Rennes, France;
- Master: Éric Fabre, Models and Algorithms for Distributed systems, M2, 12h, Master SIF (theoretical computer science)
- Master: Éric Fabre, Information Theory, ENS Rennes, M1, 16h
- Licence: Loïc Hélouët, Algorithms and Java, 40h, INSA Rennes, France;
- Master: Loïc Hélouët, Algorithms and proofs, 16h, Agrégation, ENS Rennes,
- Master: Nicolas Markey, Algorithms, 18h, Agrégation, ENS Rennes, France;
- Master: Nicolas Markey, Computability and Complexity, 18h, Agrégation, ENS Rennes, France;
- Master: Nicolas Markey, Course on Verification of Complex Systems, 10h, M2, ENS Rennes, France;
- Master: Ocan Sankur, Course on Verification of Complex Systems, 10h, M2, ENS Rennes, France;

11.2.2 Supervision

PhD students PhD in progress

- Aymeric Côme, Approximation methods for the soundness of control laws derived by machine learning, started in Dec. 2022, supervised by Éric Fabre and Loïc Hélouët
- Antoine Thébault, CIFRE grant, Efficient learning techniques for management of transport systems, started in Sept. 2022, supervised by Loïc Hélouët
- Bastien Thomas, Automated verification of randomized distributed algorithms, started in Oct. 2019, interrupted in May 2023, supervised by Nathalie Bertrand and Josef Widder (Informal Systems, Austria).
- Nicolas Waldburger, Parameterized verification of distributed algorithms under fairness conditions, started in Oct. 2021, supervised by Nathalie Bertrand, Nicolas Markey and Ocan Sankur.
- Victorien Desbois, on Heuristic search algorithms for vehicle rescheduling problem, started in December 2023, supervised by Ocan Sankur, François Schwarzenruber, Cédric Péloux (NewLogUp).
- Thibaut Le Marre, on Multi-Agent Planning and Reinforcement Learning in an Epistemic Setting, started in October 2023, supervised by François Schwarzenruber, Jilles Dibangoye, Ocan Sankur.
- Mathieu Laurent, on Efficient verification of asynchronous distributed systems, started in October 2023, supervised by Martin Quinson (Myriads/Magellan) and Thierry Jérôme.

Master Students Past master students

- Mathieu Laurent, M2 student at ENS Rennes, was supervised by Thierry Jérôme and Martin Quinson (Myriads project team) from October 2022 to June 2023. His internship topic was about mixing dynamic partial order techniques and directed model-checking for asynchronous distributed programs.
- Thomas Vitry, M1 student at ENS Rennes, was supervised by Nathalie Bertrand and Thierry Jérôme from October 2022 to June 2023 on the parameterized verification of asynchronous distributed systems.

Current master students

- Luca Paparazzo, M2 student at ENS Rennes, is supervised by Loïc Hélouët since october 2023. His internship topic is addresses energy saving in Metro systems using new enrgy games models.
- Gaëtan Regaud, M1 student at ENS Rennes, is supervised by Uli Fahrenberg and Nicolas Markey since october 2023 on history-determinism for timed systems.
- Hugo Francon, prelab student at ENS Rennes, is supervised by Nathalie Bertrand and Thierry Jéron since september 2023, on parameterized verification of asynchronous distributed systems.

11.2.3 Juries

Habilitation and PhD committees

- Nathalie Bertrand was on the PhD defense committees of Gaëtan Douéneau-Tabot (Université Paris Cité), Soumyajit Paul (Université Bordeaux and Chennai Mathematical Institute), Santiago Bautista (Université Rennes), and Nathan Tomasset (Université Paris-Saclay). She was reviewer for the PhD thesis of Olivier Stietel (Université Paris-Cité) and external examiner for the one of Oscar Darwin (Oxford University). Nathalie Bertrand was on the HDR defense committee of Ocan Sankur (Université Rennes).
- Thierry Jéron was on the PhD defense committee of Dylan Marinho (Université Nancy)
- Nicolas Markey was on the PhD defense committee of Lénaïg Cornanguer (Université Rennes).
- Loïc Hélouët was president of the PhD jury of Nicolas Amat (Université de Toulouse)

Hiring committees

Nathalie Bertrand was president of the joint hiring committee for three MCF positions at ISTIC (Université Rennes).

11.2.4 Internal or external Inria responsibilities

- Thierry Jéron is référent chercheur for Inria Rennes and IRISA.
- Nicolas Markey is co-head of the gender-equality committee of Inria Rennes and IRISA.

11.2.5 Interventions

Nicolas Markey and Loïc Hélouët took part in the Chiche! Inria program.

12 Scientific production

12.1 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. Dubslaff, 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. Lefaucheux 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>.

12.2 Publications of the year

International journals

- [8] N. Bertrand, H. Francon and N. Markey. ‘Synchronizing words under LTL constraints’. In: *Information Processing Letters* 182 (Aug. 2023), p. 106392. DOI: [10.1016/j.ipl.2023.106392](https://doi.org/10.1016/j.ipl.2023.106392). URL: <https://hal.science/hal-04219120>.
- [9] P. Bouyer, O. Kupferman, N. Markey, B. Maubert, A. Murano and G. Perelli. ‘Reasoning about Quality and Fuzziness of Strategic Behaviors’. In: *ACM Transactions on Computational Logic* 24.3 (31st July 2023), pp. 1–38. DOI: [10.1145/3582498](https://doi.org/10.1145/3582498). URL: <https://inria.hal.science/hal-04219131>.
- [10] L. Hélouët and P. Agrawal. ‘Waiting Nets: State Classes and Taxonomy’. In: *Fundamenta Informaticae*. Fundamenta Informaticae (2023), pp. 1–45. URL: <https://inria.hal.science/hal-04310411>.
- [11] A. Queffelec, O. Sankur and F. Schwarzenrüber. ‘Complexity of planning for connected agents in a partially known environment’. In: *Theoretical Computer Science* 941 (Jan. 2023), pp. 202–220. DOI: [10.1016/j.tcs.2022.11.015](https://doi.org/10.1016/j.tcs.2022.11.015). URL: <https://hal.science/hal-03930625>.

Invited conferences

- [12] N. Markey. ‘Computing the Price of Anarchy in Atomic Network Congestion Games (Invited Talk)’. In: *FORMATS 2023 - International Conference on Formal Modeling and Analysis of Timed Systems*. Vol. 14138. Lecture Notes in Computer Science. Antwerp, Belgium: Springer Nature Switzerland, 29th Aug. 2023, pp. 3–12. DOI: [10.1007/978-3-031-42626-1_1](https://doi.org/10.1007/978-3-031-42626-1_1). URL: <https://hal.science/hal-04214722>.

International peer-reviewed conferences

- [13] I. Calviac, O. Sankur and F. Schwarzenrüber. ‘Improved Complexity Results and an Efficient Solution for Connected Multi-Agent Path Finding’. In: *Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023)*. AAMAS 2023 - 22nd International Conference on Autonomous Agents and Multiagent Systems. London, United Kingdom, 2023, pp. 1–9. URL: <https://hal.science/hal-04075393>.
- [14] S. Dziadek, U. Fahrenberg and P. Schlehuber-Caissier. ‘Energy Büchi Problems’. In: *LNCS - Lecture Notes in Computer Science*. FM 2023 - 25th International Symposium on Formal Methods. Vol. 14000. Lübeck, Germany: Springer, 2023, pp. 222–239. DOI: [10.1007/978-3-031-27481-7_14](https://doi.org/10.1007/978-3-031-27481-7_14). URL: <https://inria.hal.science/hal-04344167>.

- [15] L. Hélouët and A. Thébaud. ‘Mochy : a tool for the modeling of concurrent hybrid systems’. In: *Lecture Notes in Computer Science*. Petri Nets 2023 - 44th International Conference on Application and Theory of Petri Nets and Concurrency. Vol. 13929. Lisbonne, Portugal: Springer, 2023, pp. 1–11. DOI: [10.1007/978-3-031-33620-1_11](https://doi.org/10.1007/978-3-031-33620-1_11). URL: <https://inria.hal.science/hal-04311218>.
- [16] O. Sankur. ‘Timed Automata Verification and Synthesis via Finite Automata Learning’. In: *LNCS*. TACAS 2023 - 29th International Conference on Tools and Algorithms for the Construction and Analysis of Systems. Vol. 13994. Paris, France: Springer, 22nd Apr. 2023, pp. 329–349. URL: <https://hal.science/hal-03947462>.
- [17] S. Shankar, S. Pinisetty and T. Jérón. ‘Bounded-Memory Runtime Enforcement of Timed Properties’. In: *LIPICs*. TIME 2023 - 30th International Symposium on Temporal Representation and Reasoning. Demokritos - Athènes, Greece: Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2023, pp. 1–22. DOI: [10.4230/LIPICs.TIME.2023.6](https://doi.org/10.4230/LIPICs.TIME.2023.6). URL: <https://inria.hal.science/hal-04213422>.
- [18] B. Thomas and O. Sankur. ‘PyLTA: A Verification Tool for Parameterized Distributed Algorithms’. In: *LNCS*. TACAS 2023 - 29th International Conference on Tools and Algorithms for the Construction and Analysis of Systems. Vol. 13994. Paris, France: Springer, 27th Apr. 2023, pp. 28–35. URL: <https://hal.science/hal-03996060>.
- [19] N. Waldburger. ‘Checking Presence Reachability Properties on Parameterized Shared-Memory Systems’. In: *48th International Symposium on Mathematical Foundations of Computer Science (MFCS 2023)*. MFCS 2023 - Mathematical Foundations of Computer Science. Vol. 272. Bordeaux, France: Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 28th Aug. 2023, 88:1–88:15. DOI: [10.4230/LIPICs.MFCS.2023.88](https://doi.org/10.4230/LIPICs.MFCS.2023.88). URL: <https://hal.science/hal-04394124>.

Scientific books

- [20] N. Fijalkow, N. Bertrand, P. Bouyer, R. Brenguier, A. Carayol, J. Fearnley, H. Gimbert, F. Horn, R. Ibsen-Jensen, N. Markey, B. Monmege, P. Novotny, M. Randour, O. Sankur, S. Schmitz, O. Serre and M. Skomra. *Games on Graphs: From Logic and Automata to Algorithms*. 2024, pp. 1–491. DOI: [10.48550/arXiv.2305.10546](https://doi.org/10.48550/arXiv.2305.10546). URL: <https://hal.science/hal-04273394>.

Doctoral dissertations and habilitation theses

- [21] O. Sankur. ‘Contributions on Formal Methods for Timed and Probabilistic Systems’. Université de Rennes, 6th Nov. 2023. URL: <https://hal.science/tel-04310389>.

Reports & preprints

- [22] L. Hélouët and P. Contractor. *Symbolic domains and reachability for nets with trajectories*. 2023. URL: <https://inria.hal.science/hal-04330097>.

12.3 Other

Patents

- [23] R. Noguchi, T. Jérón, N. Markey and O. Sankur. ‘Method and system for correcting the operation of a target computer system by using timed requirements’. EP 4 064 057 B1 (France). 26th July 2023. URL: <https://inria.hal.science/hal-04306235>.

12.4 Cited publications

- [24] 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>.

- [25] N. Bertrand, P. Bouyer, T. Brihaye and P. Carlier. ‘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>.
- [26] N. Bertrand, P. Bouyer, T. Brihaye and P. Carlier. ‘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>.
- [27] P. Bouyer, N. Markey, N. Perrin and P. Schlehuber-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>.
- [28] 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.
- [29] A. Nigam and N. S. Caswell. ‘Business artifacts: An approach to operational specification’. In: *IBM Systems Journal* 42.3 (2003), pp. 428–445.