RESEARCH CENTRE

Inria Center at Rennes University

IN PARTNERSHIP WITH:

Université Rennes 1, CNRS

2022

ACTIVITY REPORT

Project-Team

SUMO

SUpervision of large MOdular and distributed systems

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

DOMAIN

Algorithmics, Programming, Software and Architecture

THEME

Proofs and Verification

Inria

Contents

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]
- Eric Badouel [INRIA, Researcher, until Dec.31, 2022, HDR]
- Eric Fabre [INRIA, Senior Researcher, HDR]
- Blaise Genest [CNRS, Senior Researcher, until Jul 2022, HDR]
- Loïc Hélouët [INRIA, Senior Researcher, HDR]
- Thierry Jéron [INRIA, Senior Researcher, HDR]
- Hervé Marchand [INRIA, Researcher, HDR]
- Nicolas Markey [CNRS, Senior Researcher, HDR]
- Ocan Sankur [CNRS, Researcher]

Post-Doctoral Fellow

- Aline Goeminne [CNRS, until Aug 2022]

PhD Students

- Abdul Majith Noordheen [INRIA, until Mar 2022]
- Emily Clément [Mitsubishi Electric R&D Centre Europe, CIFRE, until Jan 2022]
- Aymeric Côme [INRIA, from Dec 2022]
- Bastien Thomas [UNIV RENNES I]
- Antoine Thébault [Alstom, CIFRE, from Sep 2022]
- Nicolas Waldburger [UNIV RENNES I]

Technical Staff

- Antoine Thébault [INRIA, Engineer, until May 2022]

Administrative Assistant

- Laurence Dinh [INRIA]

Visiting Scientists

- Willy Kengne Kungne [University Yaoundé 1, from Aug 2022 until Sep 2022]
- Joskel Ngoufou Tagueu [University Yaoundé 1, from Jun 2022 until Sep 2022]

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 [33]). 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 [31, 32] 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 developping 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 [34] 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 [35] 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, CESPA 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 [30] 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. Exemples 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, decisons 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 reasonning 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 events ICALP and Highlights that took place in Paris 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 Retirement of Éric Badouel

Éric Badouel is internationally renowned for his contributions on concurrency models. He wrote a reference book on Petri nets synthesis, co-authored with his close colleagues Luca Bernardinello and Philippe Darondeau. Motivated by public health issues in Africa, over the last years, he developed a new research axis, proposing declarative models for collaborative systems.

Throughout his carreer, Éric Badouel has strongly committed to cooperation with Africa. He co-supervized PhD theses, was director of the computer science department at École Nationale Supérieure Polytechnique de Yaoundé from 1999 to 2003, and was involved in the following events and institutions: CARI (Colloque Africain sur la Recherche en Informatique), ARIMA (Revue Africaine de Recherche en Informatique et Mathématiques Appliquées) and LIRIMA (Laboratoire International de Recherche en Informatique et Mathématiques Appliquées) as a co-director. Within Inria, he was responsible of the Africa and Middle-East area at DRI.

After these great achievements, the whole Sumo team wishes Éric Badouel to enjoy his retirement!

6.2 Awards

- Nicolas Markey received the LICS test-of-time award for his paper "Temporal logic with forgettable past", co-authored by François Laroussinie (IRIF, Paris) and Philippe Schnoebelen (LMF, Paris-Saclay). This paper is one of the most influential papers among those that appeared at LICS twenty years ago.
- Loïc Hélouët received the best paper award at the international conference Petri Nets for his paper [21].

7 New software and platforms

7.1 New software

7.1.1 MOCHY

Name: MOdels for Concurrent and HYbrid systems

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

Scientific Description: En 2022, nous avons publié une nouvelle version de l'outil (v 2.0.2), et ajouté un model checker pour la logique Signal LTL.

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.

News of the Year: In 2022, we have released versions 2.0.2 of MOCHy. We also have included a statistical model checker for the Signal LTL Logic in the tool.

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

Contact: Nathalie Bertrand

8 New results

8.1 New results on Axis 1: Quantitative models

8.1.1 Verification of probabilistic systems

Participants: Ocan Sankur.

The stochastic shortest path problem (SSPP) asks to resolve the non-deterministic choices in a Markov decision process (MDP) such that the expected accumulated weight before reaching a target state is maximized. [23] addresses the optimization of the variance-penalized expectation (VPE) of the accumulated weight, which is a variant of the SSPP in which a multiple of the variance of accumulated weights is incurred as a penalty. It is shown that the optimal VPE in MDPs with non-negative weights as well as an optimal deterministic finite-memory scheduler can be computed in exponential space. The threshold problem whether the maximal VPE exceeds a given rational is shown to be EXPTIME-hard and to lie in NEXPTIME. Furthermore, a result of interest in its own right obtained on the way is that a variance-minimal scheduler among all expectation-optimal schedulers can be computed in polynomial time.

8.1.2 Dynamical systems

Participants: Blaise Genest.

The Skolem problem is a long-standing open problem in linear dynamical systems: can a linear recurrence sequence (LRS) ever reach 0 from a given initial configuration? Similarly, the positivity problem asks whether the LRS stays positive from an initial configuration. Deciding Skolem (or positivity)

has been open for half a century: The best known decidability results are for LRS with special properties (e.g., low order recurrences). On the other hand, these problems are much easier for "uninitialized" variants, where the initial configuration is not fixed but can vary arbitrarily: checking if there is an initial configuration from which the LRS stays positive can be decided by polynomial time algorithms (Tiwari in 2004, Braverman in 2006).

In [12], we consider problems that lie between the initialized and uninitialized variant. More precisely, we ask if 0 (resp. negative numbers) can be avoided from every initial configuration in a neighborhood of a given initial configuration. This can be considered as a robust variant of the Skolem (resp. positivity) problem. We show that these problems lie at the frontier of decidability: if the neighborhood is given as part of the input, then robust Skolem and robust positivity are Diophantine-hard, i.e., solving either would entail major breakthrough in Diophantine approximations, as happens for (non-robust) positivity. Interestingly, this is the first Diophantine-hardness result on a variant of the Skolem problem, to the best of our knowledge. On the other hand, if one asks whether such a neighborhood exists, then the problems turn out to be decidable in their full generality, with PSPACE complexity. Our analysis is based on the set of initial configurations such that positivity holds, which leads to new insights into these difficult problems, and interesting geometrical interpretations.

8.1.3 Verification of Real-Time Models

Participants: Thierry Jéron, Nicolas Markey, Ocan Sankur, Loïc Hélouët.

In [22], we consider the problem of repairing inconsistent real-time requirements with respect to two consistency notions: non-vacuity, which means that each requirement can be realized without violating other ones, and rt-consistency, which means that inevitable violations are detected immediately. We provide an iterative algorithm, based on solving SMT queries, to replace designated parameters of real-time requirements with new Boolean expressions and time constraints, so that the resulting set of requirements becomes consistent.

Timed automata algorithms are based on the traversal of the state-space using zones as a symbolic representation. Since the state-space is infinite, termination relies on finite abstractions that yield a finite representation of the reachable states. The standard approach is based on extrapolations of the zones that can be efficiently represented. In [17], we present a new approach, based on simulations, that has emerged in the last ten years. This approach has led to new efficient algorithms for reachability and liveness verification of timed automata. It has also been extended to richer models such as timed automata with weights or timed automata with diagonal constraints and updates. We also present the new verification tool TChecker that implements the simulation-based approach.

In [21, 28], we have proposed verification techniques for a new timed variant of Petri nets called Waiting nets. This model extends Time Petri nets (TPNs), where 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. This is in particular the case for urban trains, where a vehicle leaves a platform when its dwell time has expired *and* a departure order has been given. This paper 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 have shown how to compute a finite state class graph for bounded waiting nets, yielding decidability of reachability and coverability. We have compare expressiveness of waiting nets with that of other models w.r.t. timed language equivalence, and shown that they are strictly more expressive than TPNs.

8.1.4 Expressiveness of Timed Models

Participants: Nicolas Markey.

A chronicle is a temporal model introduced by Dousson et al. for situation recognition. In short, a chronicle consists of a set of events and a set of real-valued temporal constraints on the delays between pairs of events. In [20], we investigate the relationship between chronicles and classical temporal-model formalisms, namely TPTL and MTL. More specifically, we answer the following question: is it possible to find an equivalent formula in such formalisms for any chronicle? This question arises from the observation that a single chronicle captures complex temporal behaviours, without imposing a particular order of the events in time. For our purpose, we introduce the subclass of linear chronicles, which set the order of occurrence of the events to be recognized in a temporal sequence. Our first result is that any chronicle can be expressed as a disjunction of linear chronicles. Our second result is that any linear chronicle has an equivalent TPTL formula. Using existing expressiveness results between TPTL and MTL, we show that some chronicles have no equivalent in MTL. This confirms that the model of chronicle has interesting properties for situation recognition.

8.2 New results on Axis 2: Large Systems Models

8.2.1 Planning Problems

Participants: Ocan Sankur.

In [10], we study the Connected Multi-Agent Path Finding (CMAPF) problem which asks for a plan to move a group of agents in a graph while respecting a connectivity constraint. 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.

8.2.2 Supervisory Control

Participants: Hervé Marchand.

In [9], we consider feedback control systems where sensor readings and actuator commands may be compromised by an attacker intending to damage the system. We study this problem at the supervisory layer of the control system, using discrete event systems techniques. The attacker can edit the outputs from the sensors of the system before they reach the supervisory controller as well as it can edit actuator commands before they reach the system. In this context, we formulate the problem of synthesizing a supervisor that is robust against a large class of edit attacks on the sensor readings and actuator commands. Intuitively, we search for a supervisor that guarantees the safety of the system even when sensor readings and actuator commands are compromised. Given the similarities of the investigated problem to the standard supervisory control problem, our solution methodology reduces the problem of synthesizing a robust supervisor against deception attacks to a supervisory control problem. This new and intuitive solution methodology improves upon prior work on this topic.

8.2.3 Regulation in Metro Systems

Participants: Éric Fabre, Loïc Hélouët.

In [11], we analyse how train delays propagate in a metro network due to disturbances and disruptions when different recovery strategies are implemented. Metro regulators use traffic management policies to recover from delays as fast as possible, return to a predefined schedule, or achieve an expected regularity of train arrivals and departures. We use as a metro traffic simulator SIMSTORS, which is based on a Stochastic Petri Net variant and simulates a physical system controlled by traffic management algorithms. To model existing metro networks, SIMSTORS has been mainly used with rule-based traffic management algorithms. In this work, we enhance traffic management strategies. We integrate SIMSTORS and the AGLIBRARY optimization solver in a closed-loop framework. AGLIBRARY is a deterministic solver for managing complex scheduling and routing problems. We formulate the real-time train rescheduling problem by means of alternative graphs, and use the decision procedures of AGLIBRARY to obtain rescheduling solutions. Several operational issues have been investigated throughout the use of the proposed simulation-optimization framework, among which how to design suitable periodic or event-based rescheduling strategies, how to setup the traffic prediction horizon, how to decide the frequency and the length of the optimization process. The Santiago Metro Line 1, in Chile, is used as a practical case study. Experiments with this framework in various settings show that integrating the optimization algorithms provided by AGLIBRARY to the rule-based traffic management embedded in SIMSTORS optimizes performance of the network, both in terms of train delay minimization and of service regularity.

This year, we have proposed a new control framework for urban train system with quantitative objectives [27]. We have used the MOCHY platform to train a Neural Network (NN) in charge of implementing a clever startegy to recover from a bunching mechanism. The obtained NN-based controller showed good results with respect to this objective. This initial work raises several interesting question. The first one is whether such a simple approach extends to any quantitative objective. The second important question, that will be addressed during A. Thebault's thesis is how to provide an understandable feedback on management rules implemented by a NN to operators.

8.2.4 Control and verification of reconfigurable systems

Participants: Abdul Majith, Hervé Marchand, Ocan Sankur.

In this work, we considered the verification and control of reconfigurable systems with a focus on SDN applications. From a theoretical point of point, we provided new techniques to analyse efficiently modular systems as well as some benchmarks on an SNN Platform showing the efficiency of our methods from both the point of view of verification and control [26]. This work was done in the context of the ADR Inria-Nokia Bell Labs.

8.2.5 Resilience in discrete event systems

Participants: Éric Fabre.

Modern software systems are assemblings of ready-made components. To ensure the reusability of components and the robustness of systems, numerous failsafes are incorporated to guarantee a minimal functioning: the system should behave well even if some component fails to perfectly fulfill its task (see the chaos engineering paradigm). This work [18] examines possible formalisms to express such a resilience property. Resilience is defined as the ability to return to a normal behavior in bounded time after a fault, and in a way that goes unnoticed by an external observer. Verifying this property is proved to be PSPACE-complete. This work examines the connections of the notion of resilience with classical properties like diagnosability (detectability of a fault) and opacity (undetectability of a secret behavior). It is proved that resilience is a strong form of opacity, or conversely that weak resilience is equivalent to opacity. The weakness degree can be preciselely quantified. To go around the difficulty of verifying resilience in practice, we consider an adaptation of k-width automata, that form a hierarchy of weakly non-deterministic systems covering all NFA. It is shown that the problem becomes PTIME in k-width automata, with a level of complexity that can be finely adjusted by the height k in the hierarchy.

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 [16], motivated by an asynchronous binary consensus algorithm by Aspnes, we consider round-based distributed algorithms communicating with shared memory. A particular challenge in these systems is that 1) the number of processes is unbounded, and, more importantly, 2) there is a fresh set of registers at each round. A verification algorithm thus needs to manage both sources of infinity. In this setting, we prove that the safety verification problem, which consists in deciding whether all possible executions avoid a given error state, is PSPACE-complete. For negative instances of the safety verification problem, we also provide exponential lower and upper bounds on the minimal number of processes needed for an error execution and on the minimal round on which the error state can be covered.

Blockchain has recently attracted the attention of the industry due, in part, to its ability to automate asset transfers. It requires distributed participants to reach a consensus on a block despite the presence of malicious (a.k.a. Byzantine) participants. Malicious participants exploit regularly weaknesses of these blockchain consensus algorithms, with sometimes devastating consequences. In fact, these weaknesses are quite common and are well illustrated by the flaws in various blockchain consensus algorithms. Paradoxically, until now, no blockchain consensus has been holistically verified. In the brief announcement [14] and in the longer paper [13], we remedy this paradox by model checking for the first time a blockchain consensus used in industry. We propose a holistic approach to verify the consensus algorithm of the Red Belly Blockchain, for any number n of processes and any number $f < \frac{n}{3}$ of Byzantine processes. We decompose directly the algorithm pseudocode in two parts-an inner broadcast algorithm and an outer decision algorithm-each modelled as a threshold automaton, and we formalize their expected properties in linear-time temporal logic. We then automatically check the inner broadcasting algorithm, under a carefully identified fairness assumption. For the verification of the outer algorithm, we simplify the model of the inner algorithm by relying on its proven properties. Doing so, we formally verify, for any parameter, not only the safety properties of the Red Belly Blockchain consensus but also its liveness in less than 70 seconds.

In 7.1.2, we present the software 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.

8.3.2 Network Congestion Games

Participants: Nathalie Bertrand, Aline Goeminne, Nicolas Markey, Ocan Sankur.

Network congestion games are a convenient model for reasoning about routing problems in a network: agents have to move from a source to a target vertex while avoiding congestion, measured as a cost depending on the number of players using the same link. Network congestion games have been extensively studied over the last 40 years, while their extension with timing constraints were considered more recently. Most of the results on network congestion games consider blind strategies: they are static, and do not adapt to the strategies selected by the other players. In [19], we extend the recent results of [Bertrand et al., Dynamic network congestion games. FSTTCS'20] to timed network congestion games, in which the availability of the edges depend on (discrete) time. We prove that computing Nash equilibria satisfying some constraint on the total cost (and in particular, computing the best and worst

Nash equilibria), and computing the social optimum, can be achieved in exponential space. The social optimum can be computed in polynomial space if all players have the same source and target.

In [15], we consider atomic congestion games on series-parallel networks, and study the structure of the sets of Nash equilibria and social local optima on a given network when the number of players varies. We establish that these sets are definable in Presburger arithmetic and that they admit semilinear representations whose all period vectors have a common direction. As an application, we prove that the prices of anarchy and stability converge to 1 as the number of players goes to infinity, and show how to exploit these semilinear representations to compute these ratios precisely for a given network and number of players.

9 Bilateral contracts and grants with industry

9.1 Bilateral contracts with industry

Nokia Bell Labs - ADR SAPIENS.

Participants: Éric Fabre, Abdul Majith, Hervé Marchand, 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 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). He defended is PhD in June 2022.

Mitsubishi Electric Research Center Europe (MERCE).

Participants: Emily Clement, 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. The SUMO team and MERCE jointly supervised a Cifre PhD student (Emily Clément) funded by MERCE from fall 2018 until March 2022; the thesis was about robustness of reachability in timed automata and will be defended in the beginning of 2022. Moreover we collaborate 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 with them on the consistency of timed requirements and their fixing (see Section 8.1).

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.

IPSCO (Intelligent Support Processes and Communities)

Participants: Éric Badouel.

- Duration: 2021 -> 2023
- Partners: Academy: University of Rennes I/IRISA with Diverse team (Coordinator) and SUMO team; Industry: Jamespot and Logpickr

The IPSCO project aims to develop a new customer support platform for digital companies and public services. Both by setting up intelligent mechanisms for filtering and processing requests from the public (customers and partners) and by providing a reflective vision of the processes implemented in the responses to these requests. In addition, to provide a robust response to small teams, the solution will enable the effective management of expert user communities to foster their autonomy and the emergence of best practices.

10 Partnerships and cooperations

10.1 International initiatives

10.1.1 Associate Teams in the framework of an Inria International Lab or in the framework of an Inria International Program

FUCHSIA

Participants: Éric Badouel.

Title: Flexible user-centric higher-order systems for collective intelligence in agencies

Duration: 2019 -> 2022

Coordinator: Georges-Edouard Kouamou

Partners:

- Université de Yaoundé Yaoundé (Cameroun)

Inria contact: Eric Badouel

Summary: Develop methods and tools, based on guarded attribute grammars, to design flexible and adaptive systems for information gathering and deliberation in order to collaboratively build expertise in health emergency situations.

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

QuaSL

Participants: Nathalie Bertrand, Nicolas Markey, Ocan Sankur.

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 vakues 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.

10.1.3 Participation in other International Programs

We are participating in the activities of 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 Mathametical Institute, and IIT Bombay.

10.1.4 Visits of international scientists

In October, S. Akshay and Goving Rajanbabu from IIT Bombay visited the SUMO Team. This was the occasion to continue our work on resilience of timed systems [29], and to start a new thread of research on timed hyperproperties.

10.2 National initiatives

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

  Participants: Emily Clement, Aline Goeminne, Thierry Jéron, Nicolas Markey, Ocan Sankur.

  – link to web site

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

  Participants: Éric Badouel, Loïc Hélouët.

  - link to website
  - Led by David Gross-Amblard (Université Rennes 1);
  - 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 facilite development, deployment, and monitoring of crowd-based participative applications. This requires handling complex workflows with multiple participants, incertainty 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.

- **ANR MAVeriQ**: Methods of Analysis for Verification of Quantitative properties (2021-2025)

  Participants: Nathalie Bertrand, Éric Fabre, Blaise Genest, 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.

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;

The team also has ongoing interactions with engineers at Alstom transports, in particular through the co-supervised CIFRE PhD of Antoine Thébault.

11 Dissemination

Participants: Éric Badouel, Nathalie Bertrand, Éric Fabre, Blaise Genest, Loïc Hélouët, Thierry Jéron, Hervé Marchand, Nicolas Markey, Ocan Sankur.

11.1 Promoting scientific activities

11.1.1 Scientific events: selection

Member of conference program committees

- Nathalie Bertrand served on the Program Committees of Concur'22, Formats'22 and GandALF'22;
- Éric Fabre served on the Program Commitee of Wodes'22;
- Loïc Hélouët served in the Program Committee of Petri Nets'22;
- Thierry Jéron served in the Program Committee of SAC-SVT'22;
- Hervé Marchand served on the Program Commitee of Wodes'22.

Reviewing activities All members of the team regularly write reviews for the main international conferences in their expertise areas: LICS, ICALP, CAV, Concur, FSTTCS, STACS, FoSSaCS, CDC.

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);
- Hervé Marchand is associate editor of the journal Discrete Event Dynamical Systems - Theory and applications (JDEDS).

Reviewing activities All members of the team reviewed a number of papers for international journals.

11.1.3 Invited talks

Nathalie Bertrand gave an invited talk at *Journées Nationales du GDR-IM* in March 2022 in Lille.

Nathalie Bertrand was invited speaker at FoSSaCS'22, the 25th International Conference on Foundations of Software Science and Computation Structures, part of ETAPS which took place in Munich in April 2022.

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.
- Hervé Marchand is a member of the IFAC Technical Committee (TC 1.3 on Discrete Event and Hybrid Systems) as well as the IEEE DES CDC committee. He is the president of the steering committee of MSR (Modélisation de Systèmes Réactifs).

11.1.5 Scientific expertise

- Éric Fabre is reviewer for the CIR program of the French Ministry of Research.
- Loïc Hélouët is reviewer for the ANR.

11.1.6 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: 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, Travaux pratiques, Analyse et Conception Formelle (ACF), 26h, M1, Univ Rennes 1, France;
- Master: Ocan Sankur, Course on Verification of Complex Systems, 10h, M2, ENS Rennes, France;
- Master: Ocan Sankur, Logic, 12h, Agrégation, ENS Rennes, France.

11.2.2 Supervision

Post-Doc

- Aline Goeminne (supervised by Nicolas Markey and Ocan Sankur), working on Timed Network Congestion Games (funded by ANR project TickTac) until September 2022.

PhD Students
Completed PhD

- Emily Clement, Verification and synthesis of control systems: efficiency and robustness, defended in March 2022, supervised by Thierry Jéron, Nicolas Markey and David Mentré (Mitsubishi Electric) [25].
- Abdul Majith Noordheen, Control of Adaptive Systems, defended in June 2022, supervised by Hervé Marchand, Ocan Sankur and Dinh Thai-Bui (Nokia Bell Labs) [26].

PhD in progress

- Aymeric Côme, Approximation methods for the soundness of control laws derived by machine learning, supervised by Éric Fabre and Loïc Hélouët
- Antoine Thébault, CIFRE grant, Efficient learning techniques for management of transport systems, supervised by Loïc Hélouët
- Bastien Thomas, Automated verification of randomized distributed algorithms, started in Oct. 2019, 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.

Master Students
Past Master Students.

- Gregory Gobin, M2 student at ENS Rennes, was supervised by Thierry Jéron and Martin Quinson (Myriads project team) since October 2021. His internship topic is the mixing of unfolding based dynamic partial order reduction and bounded model-checking for asynchronous distributed programs.
- Naïm Moussaoui, M1 student at ENS Rennes, was supervised by Thierry Jéron and Martin Quinson (Myriads project team) 2h/week during 6 months, since October 2021. The topic is unfolding-based dynamic partial order reduction for checking liveness properties on asynchronous distributed programs.

Current Master Students.

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

11.2.3 Juries

Habilitation and PhD committees

- Nathalie Bertrand was on the PhD defense committees of Xavier Badin de Montjoye (Université Versailles-St-Quentin-en-Yvelines), Guillaume Ambal (Université Rennes 1), Mathieu Hilaire (ENS Paris-Saclay), Cassius Puodzius (Université Rennes 1). She was reviewer of the PhD thesis of Émile Hazard (ENS Lyon). Nathalie Bertrand was on the HDR defense committees of Uli Fahrenberg (Paris-Saclay) and Arnaud Sangnier (Université Paris Cité).
- Thierry Jéron was in the PhD jury of Amrita Suresh (ENS Paris-Saclay)
- Nicolas Markey was a reviewer for the HDR thesis of Benjamin Monmege (Aix-Marseille Université) and for the PhD thesis of Florian Renkin (EPITA). He was on the PhD defense committee of Bastien Sérée (Ecole Centrale de Nantes).

- Ocan Sankur was in the PhD defense committees of Clément Tamines (Université de Mons), Sayan Mukherjee (Chennai Mathematical Institute).

Hiring committees

- Nathalie Bertrand was president of the hiring committee for an MCF position at ENSEIRB-MATMECA in Bordeaux and participant to hiring committees of professor positions at ENS Paris-Saclay and *Université de Lorraine*.
- Éric Fabre was in the hiring committee of Inria Lyon for young researchers, and of Inria at large for researchers with a disability.
- Hervé Marchand was member of the selection committee for “maître de conférence” (assistant professor) in computer science at Univeristy of Aix-Marseille.

Other

- Éric Fabre was in the oral examination committee for student selection of the *Écoles Normales Supérieures*.

11.3 Popularization

11.3.1 Internal or external Inria responsibilities

- Loïc Hélouët and Hervé Marchand are elected member of the "Comité de Centre" at Inria Rennes.
- Loïc Hélouët is a member of the CHSCT of Inria, and responsible for the "Young researchers" mission for Inria Rennes. He is also representative of Inria in the Matisse doctoral school council.
- Thierry Jéron is *référent chercheur* for Inria Rennes since 2016.
- Nicolas Markey is co-head of the gender-equality group of IRISA and Inria Rennes since 2021.

11.3.2 Education

Nicolas Markey is one of the main organizers of “J’peux pas, j’ai informatique”, a 1-day event for maths and computer-science teachers in secondary schools and high schools, about gender stereotypes in computer science.

Nathalie Bertrand, 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. 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. 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. 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. 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. Jéron 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] L. Hélouët, N. Markey and R. Raha. 'Reachability games with relaxed energy constraints'. In: *Information and Computation* 285B (May 2022). DOI: 10.1016/j.ic.2021.104806. URL: https://hal.inria.fr/hal-03482420.

[9] R. Meira-Góes, H. Marchand and S. Lafortune. 'Dealing with sensor and actuator deception attacks in supervisory control'. In: *Automatica* 147 (Jan. 2023). DOI: 10.1016/j.automatica.2022.110736. URL: https://hal.inria.fr/hal-03878794.

[10] A. Queffelec, O. Sankur and F. Schwarzentruber. '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. URL: https://hal.science/hal-03930625.

[11] M. L. Tessitore, M. Sama, A. D'Ariano, L. Hélouët and D. Pacciarelli. 'A Simulation-Optimization Framework for Traffic Disturbance Recovery in Metro Systems'. In: *Transportation research. Part C, Emerging technologies* 136 (Mar. 2022), pp. 1–23. DOI: 10.1016/j.trc.2021.103525. URL: https://hal.inria.fr/hal-03482456.

International peer-reviewed conferences

[12] S. Akshay, H. Bazille, B. Genest and M. Vahanwala. 'On Robustness for the Skolem and Positivity Problems'. In: STACS 2022 - 39th International Symposium on Theoretical Aspects of Computer Science. Vol. 219. Proceedings of the 39th International Symposium on Theoretical Aspects of Computer Science (STACS 2022) 5. Marseille, France, 1st Mar. 2022, pp. 1–20. DOI: 10.4230/LIPIcs.STACS.2022.5. URL: https://hal.archives-ouvertes.fr/hal-03695798.

[13] N. Bertrand, V. Gramoli, I. Konnov, M. Lazić, P. Tholoniat and J. Widder. 'Holistic Verification of Blockchain Consensus'. In: DISC 2022 - 36th International Symposium on Distributed Computing. Augusta, United States, 25th Oct. 2022, pp. 1–24. DOI: 10.4230/LIPIcs.DISC.2022.10. URL: https://hal.inria.fr/hal-03819724.

[14] N. Bertrand, V. Gramoli, I. Konnov, M. Lazić, J. Widder and P. Tholoniat. 'Brief Announcement: Holistic Verification of Blockchain Consensus'. In: PODC 2022 - 41st ACM Symposium on Principles of Distributed Computing. Salerno, Italy, 20th July 2022, pp. 1–2. DOI: 10.1145/3519270.3538468. URL: https://hal.inria.fr/hal-03819744.

[15] N. Bertrand, N. Markey, S. Sadhukhan and O. Sankur. 'Semilinear Representations for Series-Parallel Atomic Congestion Games'. In: *Leibniz International Proceedings in Informatics*. FSTTCS 2022 - 42nd IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science. Vol. 250. Chennai, France, 18th Dec. 2022, pp. 1–20. DOI: 10.4230/LIPIcs.FSTTCS.2022.32. URL: https://hal.science/hal-03937259.

[16] N. Bertrand, N. Markey, O. Sankur and N. Waldburger. 'Parameterized safety verification of round-based shared-memory systems'. In: ICALP 2022 - International Colloquium on Automata, Languages and Programming. ICALP 2022 - International Colloquium on Automata, Languages and Programming. Paris, France, 4th July 2022, pp. 1–32. DOI: 10.4230/LIPIcs.ICALP.2022.113. URL: https://hal.archives-ouvertes.fr/hal-03777440.

[17] P. Bouyer, P. Gastin, F. Herbreteau, O. Sankur and B. Srivathsan. 'Zone-based verification of timed automata: extrapolations, simulations and what next?' In: FORMATS 2022 - 20th International Conference on Formal Modeling and Analysis of Timed Systems. Warsaw, Poland, 12th Sept. 2022, pp. 1–27. URL: https://hal.archives-ouvertes.fr/hal-03654350.

[18] E. Fabre. 'Resilience in Discrete Event Systems'. In: WODES 2022 - 16th IFAC Workshop on Discrete Event Systems. Prague, Czech Republic, 7th Sept. 2022, pp. 1–6. URL: https://hal.science/hal-03777858.

[19] A. Goeminne, N. Markey and O. Sankur. 'Non-blind Strategies in Timed Network Congestion Games'. In: FORMATS 2022 - International Conference on Formal Modeling and Analysis of Timed Systems. Vol. 13465. Lecture Notes in Computer Science. Warsaw, Poland: Springer International Publishing, 29th Aug. 2022, pp. 183–199. DOI: 10.1007/978-3-031-15839-1_11. URL: https://hal.archives-ouvertes.fr/hal-03776452.

[20] T. Guyet and N. Markey. 'Logical forms of chronicles'. In: *Leibniz International Proceedings in Informatics*. TIME 2022 - 29th International Symposium on Temporal Representation and Reasoning. Virtual, Heard & McDonald Islands, 7th Nov. 2022, pp. 1–15. URL: https://hal.archives-ouvertes.fr/hal-03777471.

[21] L. Hélouët and P. Agrawal. 'Waiting Nets'. In: PETRI NETS 2022 - 43rd International Conference on Application and Theory of Petri Nets and Concurrency. Vol. 13288. Lecture Notes in Computer Science. Bergen, Norway: Springer International Publishing, 13th June 2022, pp. 67–89. DOI: 10.1007/978-3-031-06653-5_4. URL: https://hal.inria.fr/hal-03777443.

[22] R. Noguchi, O. Sankur, T. Jéron, N. Markey and D. Mentré. 'Repairing Real-Time Requirements'. In: *Lecture Notes in Computer Science*. ATVA 2022 - 20th International Symposium on Automated Technology for Verification and Analysis. Vol. 13505. Beijing, China, 25th Oct. 2022, pp. 1–16. URL: https://hal.science/hal-03777464.

[23] J. Piribauer, O. Sankur and C. Baier. 'The variance-penalized stochastic shortest path problem'. In: ICALP 2022 - 49th International Colloquium on Automata, Languages, and Programming. Vol. 229. Paris / Hybrid, France, 4th July 2022, pp. 1–19. URL: https://hal.archives-ouvertes.fr/hal-03776449.

[24] C. K. Sharpe, S. L. Ricker and H. Marchand. 'Mutual Opacity between Multiple Adversaries'. In: Wodes 2022 - 16th IFAC Workshop on Discrete Event Systems. Prague, Czech Republic, 7th Sept. 2022, pp. 1–7. URL: https://hal.inria.fr/hal-03773490.

Doctoral dissertations and habilitation theses

[25] E. Clément. 'Robustness of timed automata : computing the maximally-permissive strategies'. Université Rennes 1, 11th Mar. 2022. URL: https://theses.hal.science/tel-03666840.

[26] A. M. Noordheen. 'Automated verification and synthesis of distributed systems : in particular applied to SDN-IoT platform'. Université Rennes 1, 7th June 2022. URL: https://theses.hal.science/tel-03882284.

Reports & preprints

[27] E. Fabre, L. Hélouët and A. Thébault. *Optimization of traffic management with learning machines*. 14th Sept. 2022. URL: https://hal.inria.fr/hal-03777459.

[28] L. Hélouët and P. Agrawal. *Waiting Nets (Extended Version)*. 19th Nov. 2022. URL: https://hal.archives-ouvertes.fr/hal-03860955.

12.3 Cited publications

[29] S. Akshay, B. Genest, L. Hélouët, S. N. Krishna and S. Roychowdhury. 'Resilience of Timed Systems'. In: *41st IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science, FSTTCS 2021, December 15-17, 2021, Virtual Conference*. Ed. by M. Bojanczyk and C. Chekuri. Vol. 213. LIPIcs. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021, 33:1–33:22. DOI: 10.4230/LIPIcs.FSTTCS.2021.33. URL: https://doi.org/10.4230/LIPIcs.FSTTCS.2021.33.

[30] 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. URL: https://hal.inria.fr/hal-01599760.

[31] 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. URL: https://hal.inria.fr/hal-01397794.

[32] 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. URL: https://hal.inria.fr/hal-01938135.

[33] 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. URL: http://dx.doi.org/10.1007/s11241-016-9262-3.

[34] 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.

[35] A. Nigam and N. S. Caswell. 'Business artifacts: An approach to operational specification'. In: *IBM Systems Journal* 42.3 (2003), pp. 428–445.