IN PARTNERSHIP WITH:
CNRS

Université Rennes 1

Activity Report 2014

Project-Team SUMO

SUpervision of large MOdular and distributed systems

RESEARCH CENTER
Rennes - Bretagne-Atlantique

THEME
Proofs and Verification

Table of contents

Project-Team SUMO

Keywords: Distributed Systems, Formal Methods, Discrete Event Systems, Verification, Self-management

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

1. Members

Research Scientists

Éric Fabre [Team leader, Inria, Senior Researcher, HdR]
Éric Badouel [Inria, Researcher, HdR]
Nathalie Bertrand [Inria, Researcher]
Blaise Genest [CNRS, Researcher]
Loïc Hélouët [Inria, Researcher, HdR]
Thierry Jéron [Inria, Senior Researcher, HdR]
Hervé Marchand [Inria, Researcher]

Faculty Members

Christophe Morvan [Univ. Paris Est, Associate Professor, 1/2 Inria delegation, HdR]
Samy Abbes [Univ. Paris VII, Associate Professor, 1/2 Inria delegation from Sep 2014, HdR]

Engineer

Nicolas Berthier [Inria, granted by ANR CTRL-GREEN project]

PhD Students

Sébastien Chédor [Inria, until Jan 2014, granted by ANR CTRL-GREEN project]
Mohamadou Diouf [Inria, until Jan 2014]
Paulin Fournier [Univ. Rennes I]
Matthieu Pichene [Inria, from Dec 2014]
Srinivas Pinisetty [Inria]
Abd El Karim Kecir [Inria, from May 2014]

Visiting Scientists

Christel Baier [Invited Prof, Univ. Rennes I, Aug 2014]
Valentin Goranko [Invited Prof, Univ. Rennes I, Apr 2014]
Robert Nsaibirni [University of Yaoundé, from Jul until Aug 2014]
Doron Peled [Invited Prof, Univ. Rennes I, from Feb until Jun 2014]
Shauna Laurene Ricker [Mount Allison University, Canada, from May until June 2014]
Akshay Sundararaman [Invited Prof, Univ. Rennes I, from May until Jun 2014]

Administrative Assistant

Laurence Dinh [Inria]

Others

Rishika Garg [Inria, intern, from May until Jul 2014]
Engel Lefaucheux [Normale Sup Cachan, intern, from Mar until Jul 2014]
Ayush Maheshwari [Inria, intern, from May until Jul 2014]
Sanaa Mairouch [Inria, intern, from May until Aug 2014]
Aminatou Mohamadou [Inria, intern, from May until Jul 2014]
Dhananjay Raju [Inria, intern, from Mar until Jul 2014]
Maroua Maalej [Inria, intern, from Apr until Jul 2014]

2. Overall Objectives

2.1. Overall objectives

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 features of these pervasive applications are size, complexity, heterogeneity, and modularity (or concurrency). Besides, several of such systems are out and running before they are fully mastered, or they have grown so much that they now raise new problems that are hardly manageable by human operators. And while these systems and applications become more essential, or even critical, the demand for their *reliability, efficiency* and *manageability* becomes a central concern for computer science. The main objective of SUMO is to develop theoretical tools to address such systems, according to the following axes.

2.1.1. 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, as time, probabilities, costs, and combinations of them. This drastically changes the nature of questions that are addressed. 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. Discrete event systems approaches follow the same trend towards quantitative models. For diagnosis aspects, one is interested in the most likely explanations to observed failures, in the identification of the most informative tests to perform, in the optimal placement of sensors, and 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.

2.1.2. 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, and in particular 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.

2.1.3. 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 build online a part of their model, on demand 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 systems and problems have connections with other approaches for the management of large structured stochastic systems, 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. Model expressivity and quantitative verification

The overall objective of this axis is to combine the quantitative aspects of models with a distributed/modular setting, while maintaining the tractability of verification and management objectives.

There is first an issue of modeling, to nicely weave time, costs and probabilities with concurrency and/or asynchronism. Several approaches are quite natural, as time(d) Petri nets, networks of timed automata, communicating synchronously or through FIFO, etc. But numerous bottlenecks remain. For example, so far, no probabilistic model nicely fits the notion of concurrency: there is no clean way to express that two components are stochastically independent between two rendez-vous.

Second, the models we want to manipulate should allow for quantitative verification. This covers two aspects: either the verification question is itself quantitative (compute an optimal scheduling policy) or boolean (decide whether the probability is greater than a threshold). Our goal is to explore the frontier between decidable and undecidable problems, or more pragmatically tractable and untractable problems. Of course, there is a tradeoff between the expressivity and the tractability of a model. Models that incorporate distributed aspects, probabilities, time, etc, are typically untractable. In such a case, abstraction or approximation techniques are a work around that we will explore.

In more details, our research program on this axis covers the following topics:

- the verification of distributed timed systems,
- the verification of large scale probabilistic (dynamic) systems, with a focus on approximation techniques for such systems,
- the evaluation of the opacity/diagnosability degree of stochastic systems,
- the design of modular testing methods for large scale modular systems.

3.2. Management of large distributed systems

The generic terms of "supervision" or "management" of distributed systems cover problems like control (and controller synthesis), diagnosis, sensor placement, planning, optimization, (state) estimation, parameter identification, testing, etc. These questions have both an offline and an online facet. The literature is abundant for discrete event systems (DES), even in the distributed case, and for some quantitative aspects of DES in the centralized case (for example partially observed Markov decision processes (POMDP), probabilistic diagnosis/diagnosers, (max,+) approaches to timed automata). And there is a strong trend driving formal methods approaches towards quantitative models and questions like the most likely diagnosis, control for best average reward or for best QoS, optimal sensor placement, computing the probability of failure (un)detection, estimating the average impact of some failure or of a decision, etc. This second research axis focuses on these issues, and aims at developing new concepts and tools to master some already existing large scale systems, as telecommunication networks, cloud infrastructures, web-services, etc. (see the Application Domains section).

The objective being to address large systems, our work will be driven by two considerations: how to take advantage of the modularity of systems, and how to best approximate/abstract too complex systems by more tractable ones. We mention below main topics we will focus on:

- Approximate management methods. We will explore the relevance of ideas developed for large scale stochastic systems, as turbo-algorithms for example, in the setting of modular dynamic systems.
- Self-modeling, which consists in managing large scale systems that are known by their building rules, but which specific managed instance is only discovered at runtime, and on the fly. The model of the managed system is built on-line, following the needs of the management algorithms.
- Distributed control. We will tackle issues related to asynchronous communications between local controllers, and abstraction techniques to address large systems.
- Test and enforcement. We will tackle coverage issues for the test of large systems, and the test and enforcement of properties for timed models, or for systems handling data.

3.3. Data driven systems

The term data-driven systems refers to systems the behavior of which depends both on explicit workflows (scheduling and durations of tasks, calls to possibly distant services,...) and on the data processed by the system (stored data, parameters of a request, results of a request,...). This family of systems covers workflows that convey data (business processes or information systems), transactional systems (web stores), large databases managed with rules (banking systems), collaborative environments (health systems), etc. These systems are distributed, modular, and open: they integrate components and sub-services distributed over the web and accept requests from clients. Our objective is to provide validation and supervision tools for such systems. To achieve this goal, we have to solve several challenging tasks:

- provide realistic models, and sound automated abstraction techniques, to reason on models that are reasonable abstractions of real implemented systems designed in low-level languages (for instance BPEL (Business Process Execution Language)). These models should be able to encompass modularity, distribution, in a context where workflows and data aspects are tightly connected.
- provide tractable solutions for validation of models. Important questions that are frequently addressed (for instance safety properties or coverability) should remain decidable on our models, but also with a decent complexity.
- address design of data driven systems in a declarative way: declarative models are another way to handle data-driven systems. Rather than defining the explicit workflows and their effects on data, rule-based models state how actions are enacted in terms of the shape (pattern matching) or value of the current data. Such declarative models are well accepted in business processes (Companies such as IBM use their own model of business rules [53] to interact with their clients). Our approach is to design collaborative activities in terms of distributed structured documents, that can be seen as communicating rewriting systems. This modeling paradigm also includes models such as distributed Active XML [48], [51]. We think that distributed rewriting rules or attributed grammars can provide a practical but yet formal framework for maintenance, by providing a solution to update mandatory documentation during the lifetime of an artifact.
- address QoS management in large reconfigurable systems:

Data driven distributed systems such as web services often have constraints in terms of QoS. This calls for an analysis of quantitative features, and for reconfiguration techniques to meet QoS contracts. We will build from the experience in our team on QoS contracts composition [54] and planning [47], [49] to propose optimization and reconfiguration schemes.

4. Application Domains

4.1. Telecommunication network management

The domain of autonomic network management, will remain an important playground for SUMO. It covers a wide variety of problems, ranging from distributed (optimal) control to distributed diagnosis, optimization, reconfiguration, provisioning, etc. We have a long experience in model-based diagnosis, in particular distributed (active) diagnosis, and have recently proposed promising techniques for self-modeling. It consists in building the model of the managed network on the fly, guided by the needs of the diagnosis algorithm. This approach allows one to deal with potentially huge models, that are only described by their construction grammar, and discovered at runtime. Another important research direction concerns the management of "multi-resolution" models, that can be considered at different granularity levels. This feature is central to network design, but has no appropriate modeling formalism nor management approaches. This is a typical investigation field for abstraction techniques. Technology is ahead of theory in this domain since networks are already driven or programmed through management policies, that assign high level objectives to an abstract view of the network, leaving open the question of their optimal implementation. As a last topic of investigation, today management issues are no longer isolated within one operator, but range accross several of them, up to the supported services, which brings game theory aspects into the picture.

4.2. Control of data centers

Data centers are another example of a large scale reconfigurable and distributed system: they are composed of thousands of servers on which Virtual Machines (VM) can be (de)activated, migrated, etc. depending on the requests of the customers, on the load of the servers and on the power consumption. Autonomic management functionalities already exist to deploy and configure applications in such a distributed environment. They can also monitor the environment and react to events such as failures or overloads and reconfigure applications and/or infrastructures accordingly and autonomously. To supervise these systems, Autonomic Managers (AM) can be deployed in order to apply administration policies of specific aspects to the different entities of a data center (servers, VM, web services, power supply, etc). These AMs may be implemented in different layers: the hardware level, the operating system level or the middleware level. Therefore several control loops may coexists, and they have to take globally consistent decisions to manage the trade-off between availability, performance, scalability, security and energy consumption. This leads to multi-criteria optimization and control problems in order to automatically derive controllers in charge of the coordination of the different AMs. We are relatively new on this topic, that will require more technical investment from us. But we are driven to it by both the convergence of IT and networking, by virtualization techniques that reach networks (see the growing research effort about network operating systems), and by the call for more automation in the management of clouds. We believe our experience in network management can help. Some members of SUMO are already involved in the ANR Ctrl-Green, which addresses the controller coordination problem. We are also in contact with the Myriads team, which research interests moved from OS for grids/clouds to autonomic methods. This is supported as well by the activities of b<>com, the local IRT, where some projects in cloud management and in networking may start joint activities.

4.3. Web services and distributed active documents

Data centric systems are already deployed, and our goal is not to design new languages, architectures, or standards for them, but rather to propose techniques for the verification and monitoring of the existing systems. A bottleneck is the complexity and heterogeneity of web-based systems, that make them difficult to model and analyze. However, one can still hope for some lightweight verification or monitoring techniques for some specific aspects, for example to check the absence of conflict of interest in a transaction system, to verify (off line) and maintain (on line) the QoS, to prevent security breaches, etc. Safety aspects of WS are little addressed; any progress in that area would be useful. Besides, modeling issues are central for some applications of data centric systems. Collaborative work environments with shared active documents can be found in many

domains ranging from banking, maintenance of critical systems, webstores... We consider that models for data driven systems can find applications in most of these application areas. Our approach will be to favor purely declarative approaches for the specification of such collaborative environments. We have contacts with Centre Pasteur in Yaoundé on the design of diseases monitoring systems in developing countries. Diseases monitoring systems can be seen as a collaborative edition work, where each actor in the system reports and aggregates information about cases he or she is aware of. This collaboration is an opportunity to confront our models to real situations and real users needs. Formally modeling such a large distributed system can be seen as a way to ensure its correctness. We also envision to promote this approach as a support for maintenance operations in complex environments (train transportation, aeronautics,...). We believe this framework can be useful both for the specification of distributed maintenance procedures, for circulating information and sharing processes across teams, but also for the analysis of the correctness of procedures, possibly for their optimization or redesign, and finally to automatically elaborate logs of maintenance operations. We are in contact with several major companies on these topics, for the maintenance application side. Other industrial contacts need to be built: we have preliminary contact with IBM (leader in business artifacts), and would like to establish relations with SAP (leader in service architectures).

5. New Software and Platforms

5.1. Sigali

Participants: Hervé Marchand, Nicolas Berthier.

Sigali is a model-checking tool that operates on ILTS (Implicit Labeled Transition Systems, an equational representation of an automaton), an intermediate model for discrete event systems. It offers functionalities for verification of reactive systems and discrete controller synthesis. It is developed jointly by the TEA and SUMO teams. The techniques used consist in manipulating the system of equations instead of the set of solutions, which avoids the enumeration of the state space. Each set of states is uniquely characterized by a predicate and the operations on sets can be equivalently performed on the associated predicates. Therefore, a wide spectrum of properties, such as liveness, invariance, reachability and attractivity, can be checked. Algorithms for the computation of predicates on states are also available. Sigali is connected with the Polychrony environment (Tea project-team) as well as the Matou environment (VERIMAG), thus allowing the modeling of reactive systems by means of Signal Specification or Mode Automata and the visualization of the synthesized controller by an interactive simulation of the controlled system. Sigali is registered at APP under the identification number IDDN.FR.001.370006.S.P.1999.000.10600.

Sigali is also integrated as part of the compiler of the language BZR (web site).

We are currently developing a new version of Sigali that will be able to handle numerical variables.

5.2. Tipex

Participants: Thierry Jéron, Hervé Marchand, Srinivas Pinisetty.

We are implementing a prototype tool named Tipex (TImed Properties Enforcement during eXecution) for the enforcement of timed properties, in collaboration with Ylies Falcone (LIG, Grenoble). Tipex is based on the theory and algorithms that we develop for the synthesis of enforcement monitors for properties specified by timed automata (TA). The prototype is developped in python, and uses the PyUPPAAL and DBMpyuppaal libraries of the UPPAAL tool. It is currently restricted to safety and co-safety timed property. The property provided as input to the tool is a TA that can be specified using the UPPAAL tool, and is stored in XML format. The tool synthesizes an enforcement monitor from this TA, which can then be used to enforce a sequence of timed events to satisfy the property. Experiments have been conducted on a set of case studies. This allowed to validate the architecture and feasibility of enforcement monitoring in a timed setting and to have a first assessment of performance (and to what extent the overhead induced by monitoring is negligible).

5.3. DAXML

Participant: Loïc Hélouët.

DAXML is an implementation of Distributed Active Documents, a formalism for data centric design of Web Services proposed by Serge Abiteboul. This implementation is based on a REST framework, and can run on a network of machines connected to internet and equipped with JAVA. This implementation was realized during the post doc of Benoit Masson in 2011. A demo of the software is available at this web page. This year, the source code of DAXML has been submitted at the APP, and a distribution with free ad-hoc licence will follow in 2015.

6. New Results

6.1. Highlights of the Year

We started our first industrial collaboration "Project P22" with Alstom Transport, in the context of a common laboratory between Inria and Alstom. The project started in March 2014 and tackles robustness issues and regulation in urban train systems. The second phase of the project will start in march 2015, for a duration of three years. Most of the researchers of Sumo are involved in this project.

6.2. Control and enforcement

6.2.1. Runtime enforcement of timed properties

Participants: Thierry Jéron, Hervé Marchand, Srinivas Pinisetty.

Runtime enforcement is a powerful technique to ensure that a running system satisfies some desired properties. Using an enforcement monitor, an (untrustworthy) input execution (in the form of a sequence of events) is modified into an output sequence that complies with a property. Over the last decade, runtime enforcement has been mainly studied in the context of untimed properties. The contributions [26] and [34] deal with runtime enforcement of timed properties by revisiting the founda-tions of runtime enforcement when time between events matters. We propose a new enforcement paradigm where enforcement mechanisms are time retardants: to produce a correct output sequence, additional delays are introduced between the events of the input sequence. We consider runtime enforcement of any regular timed property defined by a timed automaton. We prove the correctness of enforcement mechanisms and prove that they enjoy two usually expected features, revisited here in the context of timed properties. The first one is soundness meaning that the output sequences (eventually) satisfy the required property. The second one is transparency, meaning that input sequences are modified in a minimal way. We also introduce two new features, i) physical constraints that describe how a time retardant is physically constrained when delaying a sequence of timed events, and ii) optimality, meaning that output sequences are produced as soon as possible. To facilitate the adoption and implementation of enforcement mechanisms, we describe them at several complementary abstraction levels. Our enforcement mechanisms have been implemented and our experimental results demonstrate the feasibility of runtime enforcement in a timed context and the effectiveness of the mechanisms. Finally, in [33], we considered more practical applications. Indeed, in network security, RE monitors can detect and prevent Denial-of-Service attacks. In resource allocation, RE monitors can ensure fairness. Specifications in these domains express data-constraints over the received events where the timing between events matters. To formalize these requirements, we introduce Parameterized Timed Automata with Variables (PTAVs), an extension of Timed Automata (TAs) with internal and external variables. We then extend enforcement for TAs to enforcement for PTAVs for safety properties. We model requirements from the considered application domains and show how enforcement monitors can ensure system correctness w.r.t. these requirements.

6.2.2. Enforcing opacity

Participant: Hervé Marchand.

In [22], we have been interested in enforcing opacity of regular predicates on modal transition systems. Intuitively, a labelled transition system $\mathcal{T}$ partially observed by an attacker, and a regular predicate S over the runs of $\mathcal{T}$, enforcing opacity of the secret S in $\mathcal{T}$ means computing a supervisory controller K such that an attacker who observes a run of the controlled system $K \setminus \mathcal{T}$ cannot ascertain that the trace of this run belongs to S based on the knowledge of $\mathcal{T}$ and K. We lift the problem from a single labelled transition system $\mathcal{T}$ to the class of all labelled transition systems specified by a *Modal Transition System* $\mathcal{M}$. The lifted problem is to compute the maximally permissive controller K such that S is opaque in $K/\mathcal{T}$ for every labelled transition system $\mathcal{T}$ which is a model of $\mathcal{M}$. The situations of the attacker and of the controller are asymmetric: at run time, the attacker may fully know $\mathcal{T}$ and K whereas the controller knows only $\mathcal{M}$ and the sequence of actions executed so far by the unknown $\mathcal{T}$.

In [23], we provided a different solution by enforcing and validate ay runtime various notion of opacity. More specically, we studied how we can model-check, verify and enforce at system runtime, several levels of opacity. Besides existing notions of opacity, we also introduce K-step strong opacity, a more practical notion of opacity that provides a stronger level of confidentiality.

6.2.3. Discrete Controller Synthesis for Infinite State Systems with ReaX

Participants: Nicolas Berthier, Hervé Marchand.

This year, we investigated the control of infinite reactive synchronous systems modeled by arithmetic symbolic transition systems for safety properties handling numerical variable. We provide effective algorithms allowing to solve the safety control problem, and report on experiments based on ReaX, our tool implementing these algorithms [28].

6.3. Model expressivity and quantitative verification

6.3.1. Diagnosis

Participants: Nathalie Bertrand, Sébastien Chédor, Éric Fabre, Loïc Hélouët, Blaise Genest, Hervé Marchand, Christophe Morvan.

Diagnosis of a system consists in providing explanations to a supervisor from a partial observation of the system and a model of possible executions. This year, we have extended results on diagnosis algorithm from scenarios. Systems are modeled using High-level Message Sequence Charts (HMSCs), and the diagnosis is given as a new HMSC, which behaviors are all explanations of the partial observation. The results published this year are first an offline centralized diagnosis algorithm (a single process in a network collects an observation, and emits a diagnosis) that has then been extended to a decentralized version of this algorithm. This allows us to give a complete diagnosis framework for infinite state systems, with a strong emphasis on concurrency and causal ordering in behaviors. HMSC-based diagnosis showed nice properties w.r.t. compositionality. We have also considered solutions for online diagnosis from scenarios, but came to the conclusion that online solutions are memory consuming, and need too many restrictions to run with finite memory. The last contribution of this work is an application of diagnosis techniques to anomaly detection, that is a comparison of observation of the system with a model of usual behaviors to detect security attacks. This work has been published this year [24].

In [21] we have been interested in the analysis of discrete event systems under partial observation which is an important topic, with major applications such as the detection of information flow and the diagnosis of faulty behaviors. These questions have, mostly, not been addressed for classical models of recursive systems, such as pushdown systems and recursive state machines. In this paper, we consider recursive tile systems, which are recursive infinite systems generated by a finite collection of finite tiles, a simplified variant of deterministic graph grammars (slightly more general than pushdown systems). Since these systems are infinite-state in general powerset constructions for monitoring do not always apply. We exhibit computable conditions on recursive tile systems and present non-trivial constructions that yield effective computation of the monitors.We apply these results to the classic problems of state-based opacity and diagnosability (off-line verification of opacity and diagnosability, and also run-time monitoring of these properties). For a decidable subclass of recursive tile systems, we also establish the decidability of the problems of state-based opacity and diagnosability.

In discrete event systems prone to unobservable faults, a diagnoser must eventually detect fault occurrences. The diagnosability problem consists in deciding whether such a diagnoser exists. We laid the foundations of diagnosis and predicatability for probabilistic systems represented by partially observed Markov chains (denoted pLTS) [32]. In particular, we studied different specifications of diagnosability and establish their relations both in finite and infinite pLTS. Then we analyzed the complexity of the diagnosability problem for finite pLTS: we showed that the polynomial time procedure proposed earlier is erroneous and that in fact for all considered specifications, the problem is PSPACE-complete. We also established tight bounds for the size of diagnosers. Afterwards we considered the dual notion of predictability which consists in predicting that in a safe run, fault will eventually occur. Predictability is easier than diagnosability: it is NLOGSPACE-complete. Yet the predictor synthesis is as hard as the diagnoser synthesis.

When a system is not diagnosable, the active diagnosis problem consists in controlling the system in order to ensure its diagnosability. In the same probabilistic setting, the active diagnosis problem consists in deciding whether there exists some observation-based strategy that makes the system diagnosable with probability one. We proved that this problem is EXPTIME-complete, and that the active diagnosis strategies are belief-based. The *safe* active diagnosis problem is similar, but aims at enforcing diagnosability while preserving a positive probability to non faulty runs, i.e. without enforcing the occurrence of a fault. We prove that this problem requires non belief-based strategies, and that it is undecidable. However, it belongs to NEXPTIME when restricted to belief-based strategies. Our work also refines the decidability/undecidability frontier for verification problems on partially observed Markov decision processes [30].

6.3.2. Probabilistic model checking

Participants: Nathalie Bertrand, Blaise Genest, Paulin Fournier.

In [16], we considered the verification of Markov chains against properties talking about distributions of probabilities. Even though a Markov chain is a very simple formalism, by discretizing in a finite number of classes the space of distributions through some symbolics, we proved that the language of trajectories of distribution (one for each initial distribution) is not regular in general, even with 3 states. We then proposed a parametrized algorithm which approximate what happens to infinity, such that each symbolic block in the approximate language is at most ϵ away from the concrete distribution.

Parameterized verification aims at validating a model of a system irrespective of the value of a parameter. This year, we studied verification problems for a model of network with the following characteristics: the number of entities is parametric, communication is performed through broadcast with adjacent neighbors, entities can change their internal state probabilistically and reconfiguration of the communication topology can happen at any time. The semantics of such a model is given in term of an infinite state system with both non deterministic and probabilistic choices. We are interested in qualitative problems like whether there exists an initial topology and a resolution of the non determinism such that a configuration exhibiting an error state is almost surely reached. We showed in [44] that all the qualitative reachability problems are decidable and some proofs are based on solving a 2 player game played on the graphs of a reconfigurable network with broadcast with parity and safety objectives.

On a different topic, we considered a control problem for stochastic sytems specified by timed automata with distributions over delays. In [29] we considered reachability objectives on such decision stochastic timed automata (DSTA). Given a reachability objective, the value 1 problem asks whether a target can be reached with probability arbitrarily close to 1. Simple examples show that the value can be 1 and yet no strategy ensures reaching the target with probability 1. In this paper, we prove that, the value 1 problem is decidable for single clock DSTA by non-trivial reduction to a simple almost-sure reachability problem on a finite Markov decision process. The ε-optimal strategies are involved: the precise probability distributions, even if they do not change the winning nature of a state, impact the timings at which ε-optimal strategies must change their decisions, and more surprisingly these timings cannot be chosen uniformly over the set of regions.

6.3.3. Distributed timed systems

Participants: Blaise Genest, Loïc Hélouët.

We have proposed and considered properties of a new timed variant of Petri nets [42], namely Timed Petri Nets with Urgency, that extend Timed Petri Nets with the main features of TPNs. Time Petri Nets (TPN) [52] and Timed Petri Nets [45] are two incomparable classes of concurrent models with timing constraints: urgency cannot be expressed using Timed Petri Nets, while TPNs can only keep track of a bounded number of continuous values (clocks). The work performed this year provides up to-our-knowledge the first decidability results for Petri Net variants combining time, urgency and unbounded places. We have obtained decidability of control-state reachability for the subclass of Timed Petri Nets with Urgency where urgency constraints can only be used on bounded places. By restricting this class to use a finite number of clocks, we have shows decidability of (marking) reachability. Formally, this class corresponds to TPNs under a new, yet natural, timed semantics where urgency constraints are restricted to bounded places. Further, under their original semantics, reachability for a more restricted class of TPNs is decidable.

6.3.4. Test Generation from Recursive Tile Systems

Participants: Sébastien Chédor, Christophe Morvan, Thierry Jéron.

In [20] we explore the generation of conformance test cases for *Recursive Tile Systems* in the framework of the classical **ioco** testing theory. The RTS model allows the description of reactive systems with recursion, and is very similar to other models like Pushdown Automata, Hyperedge Replacement Grammars or Recursive State Machines. Test generation for this kind of infinite state labelled transition systems is seldom explored in the literature. The first part presents an off-line test generation algorithm for *Weighted* RTSs, a determinizable subclass of RTSs, and the second one, an on-line test generation algorithm for the full RTS model. Both algorithms use test purposes to guide test selection through targeted behaviours. Additionally, essential properties relating verdicts produced by generated test cases with both the soundness with respect to the specification, and the precision with respect to a test purpose, are proved.

6.4. Management of large distributed systems

6.4.1. Distributed optimal planning

Participant: Éric Fabre.

Planning problems consist in organizing actions in a system in order to reach one of some target states. The actions consume and produce resources, can of course take place concurrently, and may have costs. We have a collection of results addressing this problem in the setting of distributed systems. This takes the shape of a network of components, each one holding private actions operating over its own resources, and shared/synchronized actions that can only occur in agreement with its neighbors. The goal is to design in a distributed manner a tuple of local plans, one per component, such that their combination forms a consistent global plan of minimal cost.

Our previous solutions to this problem modeled components as weighted automata. In collaboration with Loig Jezequel (TU Munich) and Victor Khomenko (Univ. of Newcastle), we have extended this approach to the case of components modeled as safe Petri nets [50]. This allows one to benefit from the internal concurrency of actions within a component. Benchmarks have shown that this method can lead to significant time reductions to find feasible plans, in good cases. In the least favorable cases, performances are comparable to those obtained with components modeled as automata. The method does not apply to all situations however, as computations require to perform ϵ-reductions on Petri-nets (our work also contains a contribution to this difficult question). This work has been accepted by the ACM Transactions in Embedded Computing Systems, to appear in 2015.

6.5. Data driven systems

6.5.1. Web services

Participants: Blaise Genest, Loïc Hélouët.

This year, we considered transactional properties (ACID) for web services. In particular, we focused on the atomicity (A of ACID) property, obtained in case of a failure inside an atomic block through compensation of the executed actions of the block. To do so, logs need to be kept. We were interested in maintaining the maximal amount of privacy. We proposed modular algorithms [19] which maintain privacy between modules, with minimal information shared among modules, both in the logging and the compensation phases. Furthermore, each module logs a small number of information, such that the sum of all actions logged is guaranteed minimal. Last, modularity allows fast algorithms, as they need to consider only what happens in the module itself, and not the exact structure of its parent module nor of its sub-modules.

We also published results on our model of sessions systems [27]. This models allows for the modeling of distributed web-based systems that are running an arbitrary number of transactions among arbitrarily many participants. We have shown how simple restrictions can guarantee decidability of simple coverability properties, and then be used to detect violation of buisness rules such as conflict of interest, or a more complex property called the chinese wall.

We are currently considering new models that manage at the same time explicit workflows and structured data. This model can be seen as a combination of AXML [46] and Petri nets.

6.5.2. An Artifact-centric Process Model

Participants: Éric Badouel, Loïc Hélouët, Christophe Morvan.

In [37] we present a purely declarative approach to artifact-centric case management systems, and a decentralization scheme for this model. Each case is presented as a tree-like structure; nodes bear information that combines data and computations. Each node belongs to a given stakeholder, and semantic rules govern the evolution of the tree structure, as well as how data values derive from information stemming from the context of the node. Stakeholders communicate through asynchronous message passing without shared memory, enabling convenient distribution.

7. Bilateral Contracts and Grants with Industry

7.1. Bilateral Contracts with Industry

Several researchers of Sumo are involved in the joint research lab of Alstom and Inria, in a common research team called P22. On Alstom side, this joint research team involves researchers of the ATS division (Automatic Train Supervision). The objective of this joint team is to evaluate regulation policies of urban train systems, to assess their robustness to perturbations and failures, to design more efficient regulation policies and finally to provide decision support for human regulators. The project started in march 2014, and a second phase of the project will start in march 2015, for a duration of three years. This covers in particular the PhD of Karim Kecir.

8. Partnerships and Cooperations

8.1. National Initiatives

8.1.1. ANR

ANR VACSIM: Validation of critical control-command systems by coupling simulation and formal analysis, 2011-2015, web site
Partners: EDF R&D, Dassault Systèmes, LURPA, I3S, LaBRI, and Inria SUMO.
The project aims at developing both methodological and formal contributions for the simulation and validation of control-command systems. SUMO contributes to quantitative analysis and its application to testing, monitoring of timed systems, and verification of communicating timed automata.

ANR Ctrl-Green: Autonomic management of green data centers, 2011-2014, web site
Partners: UJF/LIG, INPT/IRIT, Inria SUMO, EOLAS, Scalagent.
This project aims at developing techniques for the automatic optimal management of reconfigurable systems in the context of data centers using discrete controller synthesis methodology applied in the synchronous paradigm.

ANR ImpRo: Implementability and Robustness of Timed Systems, 2010-2014, web site
Partners: IRCCyN, LIP6, LSV, LIAFA, LIF, and Inria SUMO.
This project addresses the issues related to the practical implementation of formal models for the design of communicating embedded systems: such models abstract many complex features or limitations of the execution environment. The modeling of time, in particular, is usually ideal, with infinitely precise clocks, instantaneous tests or mode commutations, etc. Our objective is thus to study to what extent the practical implementation of these models preserves good properties that are satisfied by idealized models. Within ImpRo, members of SUMO mainly focus on robustness issues for timed models (timed automata, timed Petri nets,...), and diagnosis.

ANR STOCH-MC: Model-Checking of Stochastic Systems using approximated algorithms, 2014-2018, web site.
Led by SUMO.
Partners: Inria Project Team CONTRAINTES (Rocquencourt), LaBRI (Bordeaux), and LIAFA (Paris).
The aim of STOCH-MC is to perform model-checking of large stochastic systems, using controlled approximations. Two formalisms will be considered: Dynamic Bayesian Networks, which represent compactly large Markov Chains; and Markov Decision Processes, allowing non deterministic choices on top of probabilities.

8.1.2. National informal collaborations

We collaborate with Yliès Falcone (VaSCO - LIG) and Antoine Rollet (Labri) on the enforcement of timed properties.

We collaborate with Arnaud Sangnier (LIAFA) on the parameterized verification of probabilistic systems.

8.2. International Initiatives

8.2.1. Inria International Labs

Eric Badouel is member of the team Aloco (Architecture logicielle à composants) of LIRIMA, the Inria International Lab in Africa. This collaboration is on the development of artifact-centric business process models.

8.2.2. Inria Associate Teams

DISTOL (web site) is a joint project between the SUMO Team at Inria Rennes, the LogicA team at IRISA Rennes, the Chennai Mathematical Institute, the Institute of Mathematical Sciences at Chennai and the National University of Singapore.

The DISTOL project (Distributed systems, stochastic models and logics) aims at gathering researchers from Inria Rennes, two institutes in Chennai, India (CMI and IMSC) and National University of Singapore, working on formal modeling and verification of distributed systems. This project covers four main research directions. Each of these directions rely on specific and complementary competences:

- Robustness and time issues in distributed systems models (members of SUMO consider this problem with the Chennai Mathematical Institute)
- Applications of formal models & techniques to Web Services (members of SUMO consider this problem with the Chennai Mathematical Institute)
- Quantitative verification for distributed systems (members of SUMO consider this problem with researchers at NUS)
- Unification of Control Theory of Distributed Systems (This part is mainly addressed by the LogicA team in collaboration with the Institute of Mathematical Sciences)

8.2.3. Inria International Partners

8.2.3.1. Informal International Partners

We have long lasting relations with indian labs : The Chennai Mathematical Institute in Chennai (M. Mukund, N. Kumar), the Institute for Mathematical Sciences in Chennai (R. Ramanujam, K. Lodaya). We are extending these relations in India. S. Akshay holds a permanent position in IIT Bombay after his postdoc at IRISA. Our relation with our Indian partners has been formalized as associated teams (currently EA DISTOL 2012-2015).

We have started a collaboration with J. Mullins from Université Polytechnique de Montréal. The main theme of this collaboration is security properties in concurrency models. We have submitted a joint paper of variants of interference properties (information leakage) for partial order models.

We collaborate with Laurie Ricker (Mount Allison University, Canada) on the control of distributed systems and the enforcement of opacity.

8.2.4. Participation In other International Programs

AVeRTS is an Indo-French project on the algorithmic verification of real-time systems. The project is funded by CNRS on the french side, and by DST on the Indian side, under the CEFIPRA - Indo-French Program in ICST 2014-2016. From SUMO, Nathalie Bertrand and Blaise Genest are involved and contribute on stochastic timed games.

8.3. International Research Visitors

8.3.1. Visits of International Scientists

This year, S. Akshay, from IIT Bombay visited us for a one month stay, from end of May to July. This visit was funded by Rennes 1 University. During this visit, he has worked with B. Genest and L. Hélouët on verification of extensions of Petri nets calle time Petri nets with restricted urgency, that can be used to model communication systems with threshold and latency in messages. The work performed this summer is currently under submission.

Christel Baier, professor at Dresden University, was also invited for a 2-week stay paid by Rennes 1 University. She has worked during her visit with N. Bertrand on long-run quantiles in Markov decision processes.

Doron Peled visited our team for a total duration of a month in Spring 2014. He worked with B. Genest on knowledge computation in distributed systems, a work currently under review.

Valentin Goranko, professor at Stockholm University, was invited for a 2-week stay paid by Rennes 1 University. He has worked during his visit with C. Morvan on first order properties of Rational Graphs.

Laurie Ricker (Mount Allison University) visited us during for 2 months [Mai-June 2014] on the control of distributed systems and the enforcement of opacity.

Robert Nsaibirni (University of Yaoundé) visited SUMO from July to August 2014 on the use of the Guarded Attribute Grammar formalism for the description of the workspaces of actors of a disease surveillance system.

8.3.1.1. Internships

Rishika Garg

Date: May 2014 - Jul 2014

Institution: IIT Kampur (India)

Engel Lefaucheux

Date: March 2014 - July 2014

Institution: ENS Cachan (France)

Ayush Maheshwari

Date: May 2014 - July 2014

Institution: IIT Kanpur (India)

Maroua Maalej

Date: Apr 2014 - July 2014

Institution: ENSI Tunis (Tunisia)

Sanaa Mairouch

Date: May 2014 - Aug 2014

Institution: ISTIC (France)

Aminatou Mohamadou

Date: Jun 2014 - July 2014

Institution: ISTIC (France)

Dhananjay Raju

Date: March 2014 - July 2014

Institution: CMI (India)

8.3.1.2. Research stays abroad

N. Bertrand spent two visits of one month each at Mons University (Belgium), pursuing a collaboration with Thomas Brihaye, and funded by the FNRS. The resulting work on stochastic timed automata with decisions was presented at the QEST conference [29].

9. Dissemination

9.1. Aministrative duties

Éric Badouel is secretary of the Permanent Committee of CARI, co-Director of LIRIMA, and the Scientific Manager for Africa and Middle-East region at Inria DRI.

Nathalie Bertrand is secretary and committee member of Gilles Kahn PhD prize.

Éric Fabre is the co-director of the joint research lab of Alcatel-Lucent Bell Labs France and Inria, together with Olivier Audouin on Bell Labs' side. This lab has currently five active joint research teams. Éric Fabre is also member of the Scientific Committee of the joint lab of Alstom and Inria.

Blaise Genest is member of the *Comité de Centre* of Inria Rennes, and of the CNRS CRFP of Bretagne.

Loïc Hélouët was member of the hiring comittee a associate professor at INSA Rennes. Loïc Hélouët is currently *référent chercheur* (advisor) for Inria Rennes sharing this responsibility with Christine Morin.

Thierry Jéron is member of the IFIP Working Group 10.2 on Embedded Systems. He is vice-chairman of the project committee of Inria Rennes - Bretagne Atlantique. He gave an invited lecture on "Model-based conformance test generation for timed systems" at WODES'2014. He was member of 3 committees to hire Inria junior scientists (CR2 at Lillle and Saclay and CR1)

9.2. Promoting Scientific Activities

9.2.1. Scientific events organisation

9.2.1.1. general chair, scientific chair

Nathalie Bertrand is an elected steering committee member for the international conference QEST (Quantitative Evaluation of Systems).

Thierry Jéron was member of the steering committee of Movep 2014 in Nantes (July 2014).

Hervé Marchand is member of the IFAC Technical Committees (TC 1.3 on Discrete Event and Hybrid Systems) since 2005. He is member of the steering committee of MSR (Modélisation de systèmes réactifs).

9.2.2. Scientific events selection

9.2.2.1. chair/member of conference program committee

Éric Badouel has served in the Programme Committee of BPM'2014 Demos and CARI'2014.

Nathalie Bertrand was PC co-chair of QAPL'14 with L. Bortolussi. She was on the program committee of workshops PRUV'14 and QAPL'14, and of the international conferences QEST'14, TACAS'14.

Loïc Hélouët was member of the program committee of SAM'2014.

Thierry Jéron was PC member of ICTSS 2014, PECCS 2014, SAC-SVT 2014, TAP 2014, WODES 2014, and the forthcoming TAP 2015.

Hervé Marchand was PC member of the WODES conference and IFAC World Congress in 2014 and PC member of the forthcoming DCDS 2015 conference.

9.2.3. Journal

9.2.3.1. reviewer

Éric Badouel is reviewer for AMS (MathSciNet) and reviewed papers for Science of Computer Programming, Fundamenta Informaticae, Theoretical Computer Science, and Transaction on Petri Nets and Other Models of Concurrency.

Nathalie Bertrand has written reviews for Logical Methods in Computer Science, Information Processing Letters, Information and Computation.

Éric Fabre is a regular reviewer for the Ministry of Education and Research, through the Credit Impot Recherche program (support to industrial research through tax reductions). He reviewed for several journals and conferences: Automatica, Journal of discrete event dynamical systems, CDC, ACC,...

Hervé Marchand was this year reviewer for Automatica and Journal of discrete event dynamical systems.

Thierry Jéron was reviewer for Annals of telecommunications, Evolving Systems, Science of Computer Programming and Software Testing, Verification and Reliability.

9.3. Teaching - Supervision - Juries

9.3.1. Teaching

Christophe Morvan

License: Compilation, System, Advanced Algorithmics Université de Paris-Est, Marne-la-Vallée, France.

Nathalie Bertrand

Agreg: Responsible of computer science organization, 16h (eq. TD), M2, Ecole Normale Supérieure de Rennes, France.

Loïc Hélouët

Licence: JAVA programming, 35h, INSA Rennes, France.

Agreg: Algorithms, M2, ENS Rennes, France.

Master: Initiation to research, M1, ENS Rennes, France.

Éric Fabre

Master: "ASR: introduction to distributed systems and algorithms," 12h (eq. TD), M2, Univ. Rennes 1, France.

Master: "Information theory," 30h (eq. TD), M1, Ecole Normale Superieure de Rennes, France.

Master: Preparation of Agrégation, 6h (eq. TD), M1, Ecole Normale Superieure de Rennes, France.

Master: Initiation to research, 12h (eq. TD), M1, ENS Rennes, France.

Licence : "Information theory for computer scientists," 4.5h (eq. TD), L3, Ecole Normale Supérieure de Cachan, France.

9.3.2. Supervision

HdR: Christophe Morvan, Familles de graphes de présentation finie, propriétés et applications, Université de Paris-Est Marne-La-Vallée, 19 novembre 2014.

PhD: Sébastien Chédor, *Diagnostic, opacité et test de conformité pour des systèmes récursifs*, Université Rennes 1, 7th january 2015, supervised by Thierry Jéron and Christophe Morvan.

PhD: Mohamadou Lamine Diouf, *Opacité des artefacts d'un système workflow*, Rennes I University and University Cheikh Anta Diop (UCAD) in Dakar, 10th October 2014, supervised by Éric Badouel.

PhD in progress: Srinivas Pinisetty, *Runtime validation of critical control-command systems*, started in December 2011, supervised by Thierry Jéronand Hervé Marchand. Defense scheduled 23rd of January 2015.

PhD in progress: Bruno Karelovic, *Approximated analysis for checking Stochastic Models and Games*, started in November 2012, supervised by Blaise Genest and Wieslaw Zielonka.

PhD in progress : Paulin Fournier, *Parameterized verification of networks of probabilistic processes*, started in september 2012, supervised by Nathalie Bertrand and Thierry Jéron.

PhD in progress: Mathieu Pichene, Stochastic models to model apoptosis, started in December 2014, supervised by Blaise Genest.

PhD in progress: Karim Kecir, *Supervision of urban railway systems with advanced regulation techniques*, starting early 2015 (currently in pre-thesis period before a CIFRE convention is agreed by the ANRT), supervised by Loïc Hélouët.

9.3.3. Juries

Nathalie Bertrand was member of the private and public PhD defenses committee of Aaron Bohy, *Antichain based algorithms for the synthesis of reactive systems*, Université Mons, Belgium, 30/05/2014 and 10/06/2014.

This year, Thierry Jéron has been:

- member of the PhD defense jury of Axel Belinfante, *JTorX: Exploring Model-based Testing*, University of Twente, Netherlands, 18/09/2014.
- member and president of the PhD defense jury of Hernan Ponce de Léon, *Testing concurrent systems through event structures*, LSV, ENS Cachan, 7/11/2014.
- member and president of the PhD defense jury of Anaïs Guignard, *Validation fonctionnelle de contrôleurs logiques: contribution au test de conformité et à l'analyse en boucle fermée*, Lurpa, ENS Cachan, 4/12/2014.

Éric Fabre participated to the Habilitation committee (HDR) of Samy Abbes, University Paris Diderot (Paris 7), "Asynchronous Probabilistic Processes," Sept. 2014.

10. Bibliography

Major publications by the team in recent years

[1] S. AKSHAY, B. GENEST, L. HÉLOUËT, S. YANG. *Regular Set of Representatives for Time-Constrained MSC Graphs*, in "Information Processing Letters", 2012, vol. 112, n° 14-15, pp. 592-598, http://hal.inria.fr/hal-00879825

[2] E. BADOUEL, M. A. BEDNARCZYK, A. M. BORZYSZKOWSKI, B. CAILLAUD, P. DARONDEAU. *Concurrent Secrets*, in "Discrete Event Dynamic Systems", 2007, vol. 17, n° 4, pp. 425-446

[3] A. BENVENISTE, E. FABRE, S. HAAR, C. JARD. *Diagnosis of Asynchronous Discrete Event Systems: A Net Unfolding Approach*, in "IEEE Transactions on Automatic Control", November 2003, vol. 48, n^{o} 5, pp. 714-727, RNRT project MAGDA [*DOI :* 10.1109/TAC.2003.811249], http://hal.inria.fr/inria-00638224

[4] N. BERTRAND, B. GENEST, H. GIMBERT. *Qualitative Determinacy and Decidability of Stochastic Games with Signals*, in "Proceedings of LICS'09", Los Angeles, États-Unis, August 2009, http://hal.archives-ouvertes.fr/hal-00356566

[5] N. BERTRAND, T. JÉRON, A. STAINER, M. KRICHEN. *Off-line Test Selection with Test Purposes for Non-Deterministic Timed Automata*, in "Logical Methods in Computer Science", October 2012, vol. 8, n^{o} 4:8, pp. 1-33, http://hal.inria.fr/hal-00744074

[6] J. DUBREIL, P. DARONDEAU, H. MARCHAND. *Supervisory Control for Opacity*, in "IEEE Transactions on Automatic Control", May 2010, vol. 55, n^{o} 5, pp. 1089-1100 [*DOI :* 10.1109/TAC.2010.2042008]

[7] E. FABRE, A. BENVENISTE. *Partial Order Techniques for Distributed Discrete Event Systems: why you can't avoid using them*, in "Journal of Discrete Events Dynamical Systems", 2007, vol. 17, n^{o} 3, pp. 357-403

[8] E. FABRE. *Trellis Processes: a Compact Representation for Runs of Concurrent Systems*, in "Journal of Discrete Event Dynamical Systems", 2007, vol. 17, n^{o} 3, pp. 267-306

[9] E. FABRE, L. JEZEQUEL. *Distributed optimal planning: an approach by weighted automata calculus*, in "CDC", 2009, pp. 211-216

[10] B. GAUDIN, H. MARCHAND. *An Efficient Modular Method for the Control of Concurrent Discrete Event Systems: A Language-Based Approach*, in "Discrete Event Dynamic System", 2007, vol. 17, n^{o} 2, pp. 179-209

[11] T. GAZAGNAIRE, B. GENEST, L. HÉLOUËT, P. THIAGARAJAN, S. YANG. *Causal Message Sequence Charts*, in "Theoretical Computer Science", 2009, 38 p. , EA DST, http://hal.inria.fr/inria-00429538

[12] C. JARD, T. JÉRON. *TGV: theory, principles and algorithms*, in "STTT", 2005, vol. 7, n^{o} 4, pp. 297-315

[13] B. JEANNET, T. JÉRON, V. RUSU, E. ZINOVIEVA. *Symbolic Test Selection Based on Approximate Analysis*, in "TACAS", Edinburgh, Royaume-Uni, 2005, http://hal.inria.fr/inria-00564617

Publications of the year

Doctoral Dissertations and Habilitation Theses

[14] S. CHÉDOR. *Diagnosis, opacity and conformance testing for recursive tile systems*, Université Rennes 1, January 2014, https://tel.archives-ouvertes.fr/tel-00980800

[15] C. MORVAN. *Familles de graphes de présentation finie, propriétés et applications*, Université Paris-Est, November 2014, Habilitation à diriger des recherches, https://hal.archives-ouvertes.fr/tel-01094616

Articles in International Peer-Reviewed Journals

[16] M. AGRAWAL, S. AKSHAY, B. GENEST, P. THIAGARAJAN. *Approximate Verification of the Symbolic Dynamics of Markov Chains*, in "Journal of the ACM (JACM)", January 2014, forthcoming, https://hal.inria.fr/hal-00920793

[17] N. BERTRAND, P. BOUYER, T. BRIHAYE, Q. MENET, C. BAIER, M. GROESSER, M. JURDZINSKI. *Stochastic timed automata*, in "Logical Methods in Computer Science", December 2014, vol. 10, n° 4, 73 p. [*DOI :* 10.2168/LMCS-10(4:6)2014], https://hal.inria.fr/hal-01102368

[18] N. BERTRAND, A. STAINER, T. JÉRON, M. KRICHEN. *A game approach to determinize timed automata*, in "Formal Methods in System Design", December 2014, 39 p. [*DOI :* 10.1007/s10703-014-0220-1], https://hal.inria.fr/hal-01102472

[19] D. BISWAS, B. GENEST. *Minimal Observability and Privacy Preserving Compensation for Transactional Services*, in "Discrete Event Dynamic Systems", 2014, vol. 24, n° 4, pp. 611-646 [*DOI :* 10.1007/s10626-013-0177-z], https://hal.inria.fr/hal-00916645

[20] S. CHÉDOR, T. JÉRON, C. MORVAN. *Test Generation from Recursive Tile Systems*, in "Journal of Software Testing, Verification, and Reliability", November 2014, vol. 24, n° 7, pp. 532-557 [*DOI :* 10.1002/STVR.1525], https://hal.inria.fr/hal-01091672

[21] S. CHÉDOR, C. MORVAN, S. PINCHINAT, H. MARCHAND. *Diagnosis and Opacity Problems for Infinite State Systems Modeled by Recursive Tile Systems*, in "Discrete Event Dynamic Systems", 2014 [*DOI :* 10.1007/s10626-014-0197-3], https://hal.inria.fr/hal-00994970

[22] P. DARONDEAU, H. MARCHAND, S. L. RICKER. *Enforcing Opacity of Regular Predicates on Modal Transition Systems*, in "Discrete Event Dynamic Systems", 2014, 20 p. [*DOI :* 10.1007/s10626-014-0193-7], https://hal.inria.fr/hal-00987988

[23] Y. FALCONE, H. MARCHAND. *Enforcement and Validation (at runtime) of Various Notions of Opacity*, in "Discrete Event Dynamic Systems", 2014, 42 p. [*DOI :* 10.1007/s10626-014-0196-4], https://hal.inria.fr/hal-00987985

[24] L. HÉLOUËT, H. MARCHAND, B. GENEST, T. GAZAGNAIRE. *Diagnosis from Scenarios*, in "Discrete Event Dynamic Systems", 2014, vol. 24, n° 4, pp. 353–415 [*DOI :* 10.1007/s10626-013-0158-2], https://hal.inria.fr/hal-00879441

[25] G. KALYON, T. LE GALL, H. MARCHAND, T. MASSART. *Symbolic Supervisory Control of Distributed Systems with Communications*, in "IEEE Transactions on Automatic Control", 2014, vol. 59, n° 2, pp. 396-408 [*DOI :* 10.1109/TAC.2013.2283093], https://hal.inria.fr/hal-00903452

[26] S. PINISETTY, Y. FALCONE, T. JÉRON, H. MARCHAND, A. ROLLET, O. L. NGUENA TIMO. *Runtime enforcement of timed properties revisited*, in "Formal Methods in System Design", 2014, vol. 45, n° 3, pp. 381-422 [*DOI :* 10.1007/s10703-014-0215-y], https://hal.inria.fr/hal-01088136

International Conferences with Proceedings

[27] S. AKSHAY, L. HÉLOUËT, M. MUKUND. *Sessions with an unbounded number of agents*, in "14th International Conference on Application of Concurrency to System Design", Tunis, Tunisia, 14th International Conference on Application of Concurrency to System Design, IEEE, June 2014, https://hal.inria.fr/hal-01088994

[28] N. BERTHIER, H. MARCHAND. *Discrete Controller Synthesis for Infinite State Systems with ReaX*, in "IEEE International Workshop on Discrete Event Systems", Cachan, France, May 2014, pp. 420-427, https://hal.inria.fr/hal-00974553

[29] N. BERTRAND, T. BRIHAYE, B. GENEST. *Deciding the Value 1 Problem for Reachability in 1-Clock Decision Stochastic Timed Automata*, in "Quantitative Evaluation of Systems (QEST'14)", Florence, Italy, September 2014, pp. 313 - 328 [*DOI :* 10.1007/978-3-319-10696-0_25], https://hal.inria.fr/hal-01088113

[30] N. BERTRAND, E. FABRE, S. HAAR, S. HADDAD, L. HÉLOUËT. *Active diagnosis for probabilistic systems*, in "17th International Conference on Foundations of Software Science and Computation Structures (FoSSaCS'14)", Grenoble, France, A. MUSCHOLL (editor), Proceedings of FOSSACS 2014, Springer, April 2014 [*DOI :* 10.1007/978-3-642-54830-7_2], https://hal.inria.fr/hal-00930919

[31] N. BERTRAND, P. FOURNIER, A. SANGNIER. *Playing with Probabilities in Reconfigurable Broadcast Networks*, in "Foundations of Software Science and Computation Structures (FoSSaCS)", Grenoble, France, April 2014 [*DOI :* 10.1007/978-3-642-54830-7_9], https://hal.inria.fr/hal-01082129

[32] N. BERTRAND, S. HADDAD, E. LEFAUCHEUX. *Foundation of Diagnosis and Predictability in Probabilistic Systems*, in "IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS'14)", New Delhi, India, December 2014, https://hal.inria.fr/hal-01088117

[33] S. PINISETTY, Y. FALCONE, T. JÉRON, H. MARCHAND. *Runtime Enforcement of Parametric Timed Properties with Practical Applications*, in "IEEE International Workshop on Discrete Event Systems", Cachan, France, May 2014, pp. 46-53, https://hal.inria.fr/hal-00974548

[34] S. PINISETTY, Y. FALCONE, T. JÉRON, 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, March 2014, pp. 1279-1286, https://hal.inria.fr/hal-00907571

Books or Proceedings Editing

[35] N. BERTRAND, L. BORTOLUSSI (editors). *Proceedings Twelfth International Workshop on Quantitative Aspects of Programming Languages and Systems (QAPL'14)*, Electronic Proceedings in Theoretical Computer Science (EPTCS), April 2014 [*DOI :* 10.4204/EPTCS.154], https://hal.inria.fr/hal-01082133

[36] M. SELLAMI, E. BADOUEL, M. LO (editors). *Actes du CARI 2014 (Colloque africain sur la recherche en informatique et mathématiques appliquées)*, Colloques CARI, Inria, October 2014, 376 p. , https://hal.inria.fr/hal-01062320

Research Reports

[37] E. BADOUEL, L. HÉLOUËT, G.-E. KOUAMOU, C. MORVAN. *A Grammatical Approach to Data-centric Case Management in a Distributed Collaborative Environment*, May 2014, n° RR-8528, 32 p. , https://hal.inria.fr/hal-00990007

[38] E. FABRE, K. KECIR, L. HÉLOUËT. *Deliverable L3.1 : Model for regulated urban railway systems.*, Projet P22 Inria SUMO- Alstom Transport, September 2014, 21 p. , https://hal.inria.fr/hal-01094381

[39] B. GENEST, C. MORVAN, E. FABRE, H. MARCHAND, K. KECIR, L. HÉLOUËT. *Deliverable L2: Objectives of Phase 1 and use cases denition.*, Projet P22 Inria SUMO- Alstom Transport, June 2014, 5 p. , https://hal.inria.fr/hal-01094374

[40] L. HÉLOUËT, E. FABRE, C. MORVAN, B. GENEST, K. KECIR, H. MARCHAND. *Deliverable L1 : State of the art and selected references*, Projet P22 Inria SUMO- Alstom Transport, September 2014, 32 p. , https://hal.inria.fr/hal-01094366

[41] L. HÉLOUËT. *Robustness in Scenarios. ANR IMPRO Deliverable L 4.3*, ANR IMPRO, November 2014, 23 p. , https://hal.inria.fr/hal-01094363

Other Publications

[42] S. AKSHAY, B. GENEST, L. HÉLOUËT. *Timed Petri Nets with (restricted) Urgency*, November 2014, https://hal.inria.fr/hal-01088997

[43] S. AKSHAY, L. HÉLOUËT, M. MUKUND. *Sessions with an unbounded number of agents*, April 2014, https://hal.inria.fr/hal-00979409

[44] N. BERTRAND, P. FOURNIER, A. SANGNIER. *Playing with probabilities in Reconfigurable Broadcast Networks*, 2014, https://hal.archives-ouvertes.fr/hal-00929857

References in notes

[45] P. A. ABDULLA, A. NYLÉN. *Timed Petri Nets and BQOs*, in "ICATPN", Springer, 2001, pp. 53-70

[46] S. ABITEBOUL, L. SEGOUFIN, V. VIANU. *Static analysis of active XML systems*, in "Proceedings of the Twenty-Seventh ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems, PODS 2008, June 9-11, 2008, Vancouver, BC, Canada", ACM, 2008, pp. 221–230

[47] E. FABRE, L. JEZEQUEL. *Distributed optimal planning: an approach by weighted automata calculus*, in "CDC", 2009, pp. 211-216

[48] L. HÉLOUËT, A. BENVENISTE. *Document Based Modeling of Web Services Choreographies Using Active XML*, in "IEEE International Conference on Web Services, ICWS 2010", IEEE Computer Society, 2010, pp. 291-298

[49] L. JEZEQUEL, E. FABRE. *A#: A distributed version of A* for factored planning*, in "CDC", 2012, pp. 7377-7382

[50] L. JEZEQUEL, E. FABRE, V. KHOMENKO. *Factored Planning: From Automata to Petri Nets*, in "Proceedings of the 13th International Conference on Application of Concurrency to System Design, ACSD 2013", IEEE, 2013, pp. 130–139

[51] B. MASSON, L. HÉLOUËT, A. BENVENISTE. *Compatibility of Data-Centric Web Services*, in "WS-FM", Lecture Notes in Computer Science, Springer, 2011, vol. 7176, pp. 32-47

[52] P. M. MERLIN. *A Study of the Recoverability of Computing Systems*, University of CaliforniaIrvine, CA, USA, 1974

[53] A. NIGAM, N. S. CASWELL. *Business artifacts: An approach to operational specification*, in "IBM Systems Journal", 2003, vol. 42, n^{o} 3, pp. 428-445, http://dx.doi.org/10.1147/sj.423.0428

[54] S. ROSARIO. *Quality of Service issues in compositions of Web services*, Université de Rennes 1, 2009