Activity Report 2013

Team SUMO

SUpervision of large MOdular and distributed systems

RESEARCH CENTER
Rennes - Bretagne-Atlantique

THEME
Proofs and Verification

Table of contents

Team SUMO

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

Creation of the Team: 2013 January 01.

1. Members

Research Scientists

Éric Fabre [Team leader, Inria, Senior Researcher, from Jan 2013, HdR]
Éric Badouel [Inria, Researcher, from Jan 2013, HdR]
Nathalie Bertrand [Inria, Researcher, from Jan 2013]
Philippe Darondeau [Inria, Senior Researcher, from Jan 2013 until Mar 2013, HdR]
Blaise Genest [CNRS, Researcher, from Jan 2013]
Loïc Hélouët [Inria, Researcher, from Jan 2013, HdR]
Thierry Jéron [Inria, Senior Researcher, from Jan 2013, HdR]
Hervé Marchand [Inria, Researcher, from Jan 2013]

Faculty Member

Christophe Morvan [Univ. Paris Est, Associate Professor, Inria delegation from Sep 2013]

Engineer

Nicolas Berthier [Inria, granted by ANR CTRL-GREEN project, from Jun 2013]

PhD Students

Rouwaida Abdallah [Inria, granted by FP7 UNIVERSELF project, from Jan 2013 until Mar 2013]
Sébastien Chédor [Inria, granted by Université Rennes 1 till August 2013 and by ANR CTRL-GREEN project from September 2013]
Mohamadou Diouf [Inria, from Jan 2013 until Jul 2013]
Paulin Fournier [ENS Cachan Antenne de Ker Lann till August 2013 and Univ. Rennes I from September 2013]
Carole Hounkonnou [Inria, granted by FP7 UNIVERSELF project, from Jan 2013 until Apr 2013]
Aurore Junier [Inria, granted by Alcatel-Lucent Bell Labs, from Jan 2013 until Aug 2013]
Srinivas Pinisetty [Inria, granted by ANR VACSIM project, from Jan 2013]
Amélie Stainer [Univ. Rennes I, from Jan 2013 until Aug 2013]

Visiting Scientists

Luca Bernardinello [Professor at Universitat degli studi di Milano Bicocca Italy, Aug 2013]
Ylies Falcone [Associate professor at Université Joseph Fourier Grenoble, Aug 2013]
Shibashis Guha [PhD student from IIT Delhi India, from Jan 2013 until Mar 2013]
Georges Edouard Kouamou Wambo [Associate Professor at ENSP Yaoundé Cameroun, from Sep 2013 until Oct 2013]
Madhavan Mukund [Professor at Chennai Mathematical Institute India, from Sep 2013 until Oct 2013]
Akshay Sundararaman [Professor at IIT Mumbai India, from Sep 2013 until Oct 2013]

Administrative Assistant

Laurence Dinh [Inria]

Others

Baptiste Lefebvre [Inria, Internship, from Jun 2013 until Aug 2013]
Raphael Struk [Inria, Internship, from May 2013 until Jul 2013]

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 controlers, 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 controlers. The potential of this connection is largely unexplored, but it suggests that one could rather easily derive from it good approximate management methods for large distributed dynamic systems.

2.2. Highlights of the Year

- Loïc Hélouët and Hervé Marchand were co-chairs of the conference MSR 2013 (Modélisation des Systèmes Réactifs), located in Rennes this year and organized by SUMO (Laurence Dinh, Loïc Hélouët, Hervé Marchand and Paulin Fournier).
- ANR Stoch-MC has been accepted in 2013, led by SUMO (Blaise Genest (PI), Nathalie Bertrand and Éric Fabre). Its aim is to provide scalable algorithms to analyse stochastic 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 questions like:

- 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 controler 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 several objectives that we consider as more important.

- Approximate management methods: this consists in exploring the relevance of ideas developed for large scale stochastic systems, as turbo-algorithms for example, in the setting of modular dynamic systems.
- Self-modeling: 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: explores issues related to asynchronous communications between local controllers, and abstraction techniques to address large systems.
- Test and enforcement: considers 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). 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 [60] to interact with their clients). Our approach, initiated in [21], 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 [57], [59]. We think that distributed rewriting rules or attributed grammars can provide a practical but yet formal framework for maintenance documents.

- 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 that appear mostly with QoS properties, and for reconfiguration techniques to meet QoS contracts, building from the experience in our team on QoS contracts composition [61] and planning [56], [58] to propose optimization and reconfiguration schemes.

4. Application Domains

4.1. Telecommunication network management

The domain of autonomic network management, under its new hype names, 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 controler 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 (see above), 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 then 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, initiated in [21], 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. 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 Espresso/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 (Espresso/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) [45] (see Subsection 6.2.3). 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. SOFAT

Participants: Loïc Hélouët, Rouwaida Abdallah.

SOFAT is the acronym for Scenario Oracle and Formal Analysis Toolbox. As this name suggests it is a formal analysis toolbox for scenarios. Scenarios are informal descriptions of behaviors of distributed systems. SOFAT allows the edition and analysis of distributed systems specifications described using Message Sequence Charts, a scenario language standardized by the ITU [54]. The main functionalities proposed by SOFAT are the textual edition of Message Sequence Charts, their graphical visualization, the analysis of their formal properties, and their simulation. The analysis of the formal properties of a Message Sequence Chart specification determines if a description is regular, local choice, or globally cooperative. Satisfaction of these properties allows respectively for model-checking of logical formulae in temporal logic, implementation, or comparison of specifications. All these applications are either undecidable problems or unfeasible if the Message Sequence Chart description does not satisfy the corresponding property. The SOFAT toolbox implements most of the theoretical results obtained on Message Sequence Charts this last decade. It is regularly updated and re-distributed. The purpose of this software is twofold: provide a scenario based specification tool for developers of distributed applications; serve as a platform for theoretical results on scenarios and partial orders SOFAT provides several functionalities, that are: syntactical analysis of scenario descriptions, formal analysis of scenario properties, interactive simulation of scenarios when possible, and diagnosis. See also the web page.

This year, SOFAT has been extended with model transformation techniques that allow to transform non-implementable HMSCs into implementable ones [49].

APP: IDDN.FR.001.080027.000.S.P.2003.00.10600
Programming language: Java

5.4. 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. We plan to maintain this prototype as a demonstrator for our Web Services activities, and to distribute the sources.

6. New Results

6.1. Model expressivity and quantitative verification

6.1.1. Diagnosis from scenarios

Participants: Loïc Hélouët, Blaise Genest, Hervé Marchand.

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 is already available online in [25], and will soon be published.

6.1.2. Probabilistic model checking

Participants: Nathalie Bertrand, Blaise Genest, Paulin Fournier.

In [20], 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.

With the objective of model checking infinite state probabilistic systems, we proved a general finite-time convergence theorem for fixpoint expressions over a well-quasi-ordered set [22]. This has immediate applications for the verification of well-structured systems, where a main issue is the computability of fixpoint expressions, and in particular for game-theoretical properties and probabilistic systems where nesting and alternation of least and greatest fixpoints are common [35].

Parameterized verification aims at validating a system's model irrespective of the value of a parameter. In [34] we introduced a model for networks of an arbitrary number of probabilistic timed processes, communicating by broadcasting. This model is suitable for distributed protocols, and can be applied to wireless sensor networks or peer-to-peer applications. The number of processes is unknown and either is constant (static case), or evolves over time through random disappearances and creations (dynamic case). On the one hand, most parameterized verification problems turn out to be undecidable in the static case (even for untimed processes). On the other hand, we prove their decidability in the dynamic case.

6.1.3. Distributed timed systems

Participants: Nathalie Bertrand, Amélie Stainer.

We study the reachability problem for communicating timed processes, both in discrete and dense time. Our model comprises automata with local timing constraints communicating over unbounded FIFO channels. Each automaton can only access its set of local clocks; all clocks evolve at the same rate. Our main contribution is a complete characterization of decidable and undecidable communication topologies, for both discrete and dense time. We also obtain complexity results, by showing that communicating timed processes are at least as hard as Petri nets; in the discrete time, we also show equivalence with Petri nets. Our results follow from mutual topology-preserving reductions between timed automata and (untimed) counter automata. To account for urgency of receptions, we also investigate the case where processes can test emptiness of channels. This resut is published in [39] and is a part of Amélie Stainer's PhD manuscript [18]. It also constitutes a contribution to ANR VACSIM.

We also studied a model for distributed systems composed of stochastic and timed processes that interact via broadcasting. For these networks of stochastic timed automata (NSTA), we provided a precise performance evaluation algorithm, without resorting to simulation techniques. The idea is to characterize the general state space Markov chain through transient stochastic state classes that represent the system's state after each action. This yields an algorithmic approach to the transient analysis of NSTA models, with fairly general termination conditions [32].

6.2. Management of large distributed systems

6.2.1. Test generation from Recursive Tile Systems

Participants: Sébastien Chédor, Thierry Jéron, Christophe Morvan.

We explore the generation of conformance test cases for Recursive Tile Systems (RTSs) 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 models is seldom explored in the literature. We first propose an off-line test generation algorithm for Weighted RTSs, a determinizable sub-class of RTSs, and second, 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 on an implementation with both the conformance with respect to its specification, and the precision with respect to a test purpose, are proved. This work is published in [51], and a journal version will appear in 2014. It is also a part of Sébastien Chédor's PhD manuscript.

6.2.2. Distributed control

Participants: Blaise Genest, Hervé Marchand.

We focused this year on the control of distributed systems modeled as *asynchronous automata*, that is asynchronous network of automata communicating through peer to peer synchronizations. First, we considered the case where all events are controllable, and the objective is to accept exactly a given language. Here, a famous result is the Zielonka theorem [62], stating that every regular language closed under commutation can be turned into an asynchronous automaton. However, the construction is plagued with deadends and final state of the network are decided by a global controller monitoring every process at the same time and perfectly, which is unrealistic and defeat the distribution idea. This year, we characterized the languages which can be controlled realistically (no deadends, local final states and local decision on each process), and give algorithms to obtain the associated distributed machines in [30]. The case where some events are uncontrollable is reputed very difficult. We made a progress this year in [42], showing that we can decide whether a reachability objective can be ensured, granted that the communication between the processes follow a tree: siblings can not communicate directly together, they need to go through their common parent.

In [27], we consider an alternative model for the control of distributed systems; the aim is to build local controllers that restrict the behavior of a distributed system in order to satisfy a global state avoidance property. We model distributed systems as communicating finite state machines with reliable unbounded FIFO queues between subsystems. Local controllers can only observe the behavior of their proper subsystem and do not see the queue contents. To refine their control policy, controllers can use the FIFO queues to communicate by piggy-backing extra information (some timestamps and their state estimates) to the messages sent by the subsystems. We provide an algorithm that computes, for each local subsystem (and thus for each controller), during the execution of the system, an estimate of the current global state of the distributed system. We then define a synthesis algorithm to compute local controllers. Our method relies on the computation of (co-)reachable states. Since the reachability problem is undecidable in our model, we use abstract interpretation techniques to obtain overapproximations of (co-)reachable states. Similarly, in [46], we have been interested in the control of distributed systems with synchronous communications (called decentralized Discrete Event Systems). We introduced a novel architecture that extends the class of problems that can be solved in decentralized DES control in the absence of communication. In this architecture, unlike previous architectures that use either conjunction or disjunction to fuse local control decisions, the fusion rule is exclusive or. We characterized the new architecture, where controllers take a single decision, with respect to the recently-proposed multi-decision framework of Chakib and Khoumsi. Unlike previous architectures, parity-based controllers cannot predetermine their local control decision based solely on their local observations. Instead, the local control decisions are calculated a priori.

6.2.3. Enforcement of timed and security properties

Participants: Thierry Jéron, Hervé Marchand, Srinivas Pinisetty.

Runtime enforcement is a verification/validation technique aiming at correcting (possibly incorrect) executions of a system of interest. This year, we first consider enforcement monitoring for systems with timing specifications (modeled as timed automata). We consider runtime enforcement of any regular timed property specified by a timed automaton [45]. To ease their design and their correctness-proof, enforcement mechanisms are described at several levels: enforcement functions that specify the input-output behavior, constraints that should be satisfied by such functions, enforcement monitors that implement an enforcement function as a transition system, and enforcement algorithms that describe the implementation of enforcement monitors. The feasibility of enforcement monitoring for timed properties is validated by prototyping the synthesis of enforcement monitors. This work is also a contribution to ANR Vacsim. In [41], we studied an alternative enforcement problem of security properties, namely, the enforcement of K-step opacity at runtime. In K-step opacity, the knowledge of the secret is of interest to the attacker within K steps after the secret occurs and becomes obsolete afterwards. We introduce the mechanism of runtime enforcer that is placed between the output of the system and the attacker and enforces opacity using delays. If an output event from the system violates K-step opacity, the enforcer stores the event in the memory, for the minimal number of system steps until the secret is no longer interesting to the attacker (or, K-step opacity holds again)

6.2.4. Discrete control of computing systems administration

Participants: Hervé Marchand, Nicolas Berthier.

We address the problem of using Discrete Controller Synthesis for the administration of Computing Systems, following an approach supported by a programming language [24]. We present a mixed imperative/declarative programming language, where declarative contracts are enforced upon imperatively described behaviors. Its compilation is based on the notion of supervisory control of discrete event systems. More precisely, our language can serve programming closed-loop adaptation controllers, enabling flexible execution of functionalities w.r.t. changing resource and environment conditions. DCS is integrated into a1 programming language compiler, which facilitates its use by users and programmers, performing executable code generation. The tool is concretely built upon the basis of a reactive programming language compiler, where the nodes describe behaviors that can be modeled in terms of transition systems. Our compiler integrates this with a DCS tool, making it a new environment for formal methods. We apply our method to the problem of coordinating several administration loops in a data center (number of servers, repair, and local processor frequencies) [40]. We formulate this problem as an invariance controller synthesis problem. We are currently working on an extension of the controller synthesis tool so that it can handle the use of numerical variables in order to model both the system and the properties to be ensured by control.

6.2.5. Distributed 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 consistent local plans, one per component, such that their combination forms a global plan of minimal cost.

Our previous solutions to this problem modeled components as weighted automata. In collaboration with Loïg Jezequel (TU Munich) and Victor Khomenko (Univ. of Newcastle), we have extended this approach to the case of components modeled as safe Petri nets [44]. 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.

6.2.6. *Diagnosis based on self-modeling*

Participants: Éric Fabre, Carole Hounkonnou.

Model-based approaches have been proved to provide the best results for fault diagnosis in telecommunication networks, with various kinds of models. They suffer however from several difficulties: one has to build a model adequate to the supervised network (and possibly adapt it as the network evolves), one has to find the correct abstraction level for this model, and one has to deal with size issues of such models. In Carole Hounkonnou's thesis [15], we have proposed an approach that addresses these three limitations, under the generic name of self-modeling. It consists modeling a network in a generic manner, through its building rules. The actual instance one has to manage is then discovered on the fly, when some malfunction explanation request is triggered. Starting from the identified malfunction, the network model instance is discovered/revealed progressively, as requested by the needs of the diagnosis procedure. The latter progressively extends a Bayesian network model of the network, in order to collect more information and identify the malfunction rootcause. The model extension is guided by an information theory criterion: it seeks access to the new observations that are be the most informative (on the average) given previous observations taken into account. This approach allows to deal with potentially large models, as the supervised system needs not be entirely modeled before the diagnosis starts. We are currently working on the extension of this setting to model refinement, and to a framework of dynamic systems rather than static systems.

6.2.7. *Graceful restart methods for link state routing protocols*

Participants: Éric Fabre, Carole Hounkonnou.

Link state routing protocols are ubiquitous in the internet. OSPF (Open Shortest Path First) is one of them within an Autonomous System. In collaboration with Alcatel-Lucent, we have proposed an extension of graceful restart procedures, that allow to shut down the control plane of routers while maintaining the data plane active, and thus the packet forwarding activity. A drawback of existing procedures was that frozen routers had to be removed from the network as soon as topology evolved. We have shown that this pessimistic precaution could be damageable to the network and was not necessary [43]. Frozen routers may still be useful, even if they do not forward packets in an optimal manner. And even if they create routing loops, the latter can be easily detected, and optimally patched, which is often more efficient than declaring these routers as dead. Experiments on classical topologies of the topology zoo, as well as on random topologies, have confirmed these results.

6.3. Data driven systems

6.3.1. *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 [23] 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 have extended the *session system* model originally proposed in [55]. We have deisgned a mode for Web-based systems that allows to describe systems running an arbitrary number of transactions over an arbitrary number of agents. For these systems, syntactic restrictions allow to decide coverability properties, and then more elaborated business rules, such as conflict of interest (the fact that a participant to a system can be involved in two exclusive services), or the Chinese Wall Property (that prevents users of a system to use benefits or information right they may have obtained from a privileged role at later instant of any execution of the system. These results were obtained with M. Mukund and S. Akshay within the context of the DISTOL associated team, and should lead to a publication next year.

6.3.2. Implementation of scenarios

Participants: Loïc Hélouët, Rouwaida Abdallah.

We have revisited the problem of program synthesis from specifications described by High-level Message Sequence Charts. The main objective is to obtain a distributed implementation (for instance described with communicating automata) from a global specification given as High-level MSCS. In the general case, synthesis by a simple projection on each component of the system allows more behaviors in the implementation than in the specification. The differences arise from loss of ordering among messages, but we have shown that for a subclass of HMSCs (the *local HMSCs*) behaviors can be preserved by addition of communication controllers, that intercept messages to add stamping information before resending them, and deliver messages to processes in the order described by the specification. This work was published in [19].

The second aspect of our work on scenarios implementability has considered implementation of requirements expressed as non-local HMSCs. We have proposed a new technique to transform an arbitrary HMSC specification into a local HMSC, hence allowing implementation. This transformation can be automated as a constraint optimization problem, and the impact of modifications brought to the original specification minimized w.r.t. a cost function. The approach was evaluated on a large number of randomly generated HMSCs, and the results show an average runtime of a few seconds, which demonstrates applicability of the technique. These results were published in [28]. Both results mentionned in this sections are part of the PhD thesis of Rouwaida Abdallah, defended this year [14].

6.3.3. Attribute grammars

Participant: Éric Badouel.

Evaluation of attributes w.r.t. an attribute grammar can be obtained by inductively computing a function expressing the dependencies of the synthesized attributes on inherited attributes. This higher-order functional approach to attribute evaluation can straightforwardly be implemented in a higher-order lazy functional language like Haskell. The resulting evaluation functions are, however, not easily amenable to optimization when we want to compose two attribute grammars. In [21], we present an alternative first-order functional interpretation of attribute grammars where the input tree is replaced by an extended cyclic tree each node of which is aware of its context viewed as an additional child tree. These cyclic representations of zippers (trees with their context) are natural generalizations of doubly-linked lists to trees over an arbitrary signature. Then we show that, up to that representation, descriptional composition of attribute grammars reduces to the composition of tree transducers.

7. Bilateral Contracts and Grants with Industry

7.1. Bilateral Contracts with Industry

HiMa. The SUMO team was involved in the common research lab between Alcatel-Lucent Bell Labs France and Inria, through the High-Manageability team, co-headed by Pierre Peloso (Alcatel) and Éric Fabre. This joint team involved other Inria teams: Madynes and Mexico. In the last years of its existence, most of the activity of this joint team was redirected to the UniverSelf FP7 integrated project. Both the joint team and the project ended in 2013 (see the UniverSelf description for more details). This joint team supported two PhD students of SUMO, who defended their thesis in 2013: Aurore Junier (network calculus for early malfunction detection) and Carole Hounkonnou (self-diagnosis for large scale services and networks). Éric Fabre is member of the scientific board of the joint lab ALBLF-Inria, which is now entering in its second round of 5 year common teams.

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, 2011-2014, 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
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.
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) on the enforcement of timed properties and Tristan Le Gall (CEA) on the control of distributed systems.

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

8.2. European Initiatives

8.2.1. FP7 Projects

Participant: Éric Fabre.

Univerself is a FP7 IP, with 19 partners, among which Alcatel-Lucent, Orange Labs, Thales Comunications, Telefonica, Telecom Italia as industrial partners. It lasted from Sept. 2010 to Nov. 2013. See also http://www.univerself-project.eu/ Univerself aimed at developing self-management methods for telecommunication networks, regardless of technological boundaries (wireless, wireline, services) and at providing tools for their integrability and acceptability. The focus was first on the development of network empowerment methods (NEM), that address specific needs in automating management functions, for example power tuning in SONs (Self-Organizing Networks), network and/or service diagnosis, vulnerability detection and correction, knowledge acquisition and elaboration, optimal resource usage and allocation, etc. A second set of results was on a methodology to deploy and coordinate such NEMs, through a Universal Management Framework (UMF).

8.3. International Initiatives

8.3.1. 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)

8.3.2. Inria International Partners

8.3.2.1. Declared Inria International Partners

Éric Badouelis member of the team ALOCO (Architecture logicielle à Composants) of LIRIMA lab (Laboratoire international de recherche en informatique et mathématiques appliquées). LIRIMA is an african lab with headquarters in Yaoundé (Cameroun) partially funded by Inria. Within the team ALOCO, Éric collaborates on artifact-centric business process models.

8.3.2.2. Informal International Partners

We collaborate with Thomas Brihaye (UMONS, Brihaye) on the verification of stochastic timed systems.

We collaborated with Laurie Ricker (Mount Allison University, Canada) and Thierry Massart (ULB,Belgium) on the control of distributed systems.

8.3.3. Participation in other International Programs

Several researchers of the SUMO team are members of the LIA Informel. The Indo-French Formal Methods Lab is a CNRS International Associated Laboratory fostering the scientific collaboration between India and France in the domain of formal methods and applications to the verification of complex systems.

The research within LIA Informel focuses on theoretical foundations of games, automata, and logics, three important tools in formal methods. Members of Informel work on the verification of safety-critical systems, with an emphasis on quantitative aspects (time, cost, energy, etc.), concurrency, control, and security protocols.

8.4. International Research Visitors

8.4.1. Visits of International Scientists

S. Akshay from IIT Bombay visited the SUMO team one week in autumn.

Luca Bernardinello, professor at the University of Milano Bicocca (Italy).

Thomas Brihaye, professor at Mons University (Belgium), spent one month in SUMO team as ISTIC (University Rennes 1) invited professor.

Georges-Edouard Kouamou, junior professor at ENSP Yaoundé (Cameroun).

Madhavan Mukund, from the Chennai Mathematical Institute, visited SUMO in May 2013 and was part of Loïc Hélouët's habilitation jury. He also stayed one week in autumn.

Laurie Ricker (Mount Allison University) visited us during for 2 weeks in March 2013.

8.4.1.1. Internships

Shibashis Guha, PhD student at IIT Delhi, spent two months in SUMO team, supervised by Nathalie Bertrand.

Baptiste Lefebvre (L3 student, ENS Ulm), was an intern from June to Aug. 2013, on the experimental evaluation of an enhanced graceful shutdown method for the OSPF routing protocol, supervised by Éric Fabre.

Raphael Struk (L3 student, ENS Rennes), did an internship supervised by Blaise Genest and Loïc Hélouët.

9. Dissemination

9.1. Scientific Animation

Éric Fabre was evaluator of the first round of the ANR call for projects, 2013. He is also a regular reviewer for the Ministry of Research and Innovation, through the Credit Impot Recherche program (support to industrial research through tax reductions).

Éric Badouel is Associate Editor of the ARIMA journal, member of the Board of SARIMA and of the Steering Committee of LIRIMA. He is the Secretary of the Permanent Committee of the CARI.

Nathalie Bertrand is elected member of the Steering Committee of QEST, international conference on Quantitative Evaluation of Systems. She is also on the Steering Committee of the international workshop QAPL. She has served this year on the Programme Committee of MSR'13, MFCS'13, QEST'13 and QAPL'13. She is member and scientific secretary of the Gilles Kahn PhD award committee.

Thierry Jéron was PC member of ACM SAC-SVT 2014, PECCS 2014, MAROC 2013, TAP 2013, RV 2013. He was co-chairman of a Dagstuhl seminar on Symbolic Methods in Testing (January 2013). He is member of the steering committee of Movep'2014 in Nantes (July 2014). He is member of the IFIP Working Group 10.2 on Embedded Systems. He gave an invited lecture on "Model-based conformance test generation for timed systems" at the workshop MAROC'2013.

Loïc Hélouët was co-organizer with Hervé Marchandof the MSR 2013 Conference. He is member of the program committee of the SDL conference. In 2013, he was also reviewer for the following conferences: SDL, TACAS, CONCUR, MOVEP, and journals : SOSYM, TCS. He is also scientific coordinator of the DISTOL associated team. He co-organizes (with N. Bertrand, F. Schwarzentruber, D. Cachera and J.-P. Talpin) the 68NQRT seminar at IRISA/Inria Rennes, a weekly event that proposes talks in the domain of theoretical computer science (around 40 talks each year from worldwide participants). He is *référent chercheur* for the Inria Rennes research center, helping researchers that face difficulties during their carreer. He was part of the committee for the selection of a *maître de conférences* position at ISTIC in May 2013.

Hervé Marchand 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). He was co-organizer with Loïc Hélouëtof MSR 2013 in Rennes (13-15 November 2013). He was PC member of DCDS Conference and PC chair with Loïc Hélouëtof MSR 2013. He is PC member of the forthcoming WODES conference and IFAC World Congress in 2014. He was reviewer for Automatica, Transaction automatic and control, Discrete Event dynamical systems as well as CDC, ACC conference.

9.2. Teaching - Supervision - Juries

9.2.1. Teaching

Éric Fabre

Master: ASR: introduction to distributed systems and algorithms, 12h, M2, Univ. Rennes 1, France.

Master: Information theory, 30h, M1, Ecole Normale Superieure de Rennes, France.

Nathalie Bertrand

Master: Advanced verification techniques, 15h (eq. TD), M2, ISTIC, Université de Rennes 1, France.

Agreg: Formal languages, 27h (eq. TD), M2, Ecole Normale Superieure de Rennes, France.

Loïc Hélouët

Licence: JAVA programming, 37h, INSA Rennes, France.

Agreg: Finite automata, and flow algorithms, 8h (eq. TD), M2, ENS Rennes, France.

9.2.2. Supervision

HdR: Loïc Hélouët, *Scenario Automata: theory and applications* , Rennes 1 University, 17th May 2013.

PhD: Carole Hounkonnou, *Auto-diagnostic actif dans les réseaux de télécommunications*, Rennes 1 University, 12th July 2013, supervised by Éric Fabre.

PhD: Rouwaida Abdallah, *Implémentabilité de systèmes distribués décrits à l'aide de scénarios*, ENS Cachan antenne de Bretagne, 16th July 2013, supervised by Loïc Hélouët and Claude Jard.

PhD: Amélie Stainer, *Contribution to the Verification of Timed Automata: Determinization, Quantitative Verification and Reachability in Networks of Automata*, Rennes 1 University, 25th November 2013, supervised by Thierry Jéron and Nathalie Bertrand.

PhD: Aurore Junier, *Performance and stability analysis in telecommunication networks*, Rennes 1 University, 16th December 2013, supervised by Anne Bouillard and Claude Jard.

PhD in progress (Defense on January 7th, 2014): Sébastien Chédor, *Diagnostic, opacité et test de conformité pour des systèmes récursifs.*, started in September 2009, supervised by Thierry Jéron and Christophe Morvan.

PhD in progress: Mohamadou Lamine Diouf, *Opacité des artefacts dans un système workflow*, started in spring 2011, supervised by Éric Badouel and colocated with Dakar Université Cheik Anta Diop (Senegal).

PhD in progress: Srinivas Pinisetty, *Runtime validation of critical control-command systems*, started in December 2011, supervised by Hervé Marchand and Thierry Jéron.

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

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

9.2.3. Juries

Éric Fabre was reviewer of the PhD thesis of Fabien Kuntz, "Optimization of the monitoring of avionic systems through enhanced diagnosis performances," LABRI and University Bordeaux 1, July 2013. He also reviewed the PhD thesis of Leila Bennacer, "Contribution to self-diagnosis methods in large scale communication networks," University Paris-Est Creteil, Dec. 2013.

Thierry Jéron was member of the PhD defense jury of Aymeric Hervieu (Dec. 2013, Université Rennes 1).

Hervé Marchand was a member of the PhD defense juries of Mohammed Ali Kammoun (LAGIS, Metz) in July 2013 and of Xin An (Inria Rhones Alpes, Grenoble) in October 2013.

9.3. Popularization

Loïc Hélouët contributed to a vist of young pupils (3e) visiting IRISA to discover a research environment in February 2012. He did a short interactive presentation (1 hour) of his research theme, and of the duties of a researcher.

Éric Fabre gave a survey presentation about failure diagnosis in telecommunication networks to the 2nd year students of ENST Bretagne (Brest) engaged in a research training track. This was followed by informal discussions about the every day life of a researcher.

10. Bibliography

Major publications by the team in recent years

[1] S. AKSHAY, B. GENEST, L. HELOUET, S. YANG. *Regular Set of Representatives for Time-Constrained MSC Graphs*, in "Information Processing Letters", 2012, vol. 112, n^{o} 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^{o} 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], http://hal.inria.fr/inria-00483891

[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. HELOUET, 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] R. ABDALLAH. , *Implementability of distributed systems described with scenarios*, École normale supérieure de Cachan - ENS Cachan, July 2013, http://hal.inria.fr/tel-00919684

[15] C. HOUNKONNOU. , *Auto-diagnostic actif dans les réseaux de télécommunications*, Université Rennes 1, July 2013, http://hal.inria.fr/tel-00931853

[16] L. HÉLOUËT. , *Automates d'ordres : théorie et applications*, Université Rennes 1, May 2013, Habilitation à Diriger des Recherches, http://hal.inria.fr/tel-00926742

[17] A. JUNIER. , *Analyse de performance et de stabilité des réseaux de télécommunication*, École normale supérieure de Cachan - ENS Cachan, December 2013, http://hal.inria.fr/tel-00932176

[18] A. STAINER. , *Contribution à la vérification des automates temporisés : déterminisation, vérification quantitative et accessibilité dans les réseaux d'automates*, Université Rennes 1, November 2013, http://hal.inria.fr/tel-00926316

Articles in International Peer-Reviewed Journals

[19] R. ABDALLAH, L. HÉLOUËT, C. JARD. *Distributed Implementation of Message Sequence Charts*, in "Software and Systems Modeling", June 2013, pp. 10-32, http://hal.inria.fr/hal-00840372

[20] 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, accepted, http://hal.inria.fr/hal-00920793

[21] E. BADOUEL, R. TCHOUGONG, C. NKUIMI-JUGNIA, B. FOTSING. *Attribute grammars as tree transducers over cyclic representations of infinite trees and their descriptional composition*, in "Theoretical Computer Science", February 2013, vol. 480, pp. 1-25 [*DOI :* 10.1016/J-TCS.2013.02.007], http://hal.inria.fr/hal-00915031

[22] N. BERTRAND, P. SCHNOEBELEN. *Computable fixpoints in well-structured symbolic model checking*, in "Formal Methods in System Design", October 2013, vol. 43, n^{o} 2, pp. 233-267 [*DOI :* 10.1007/S10703-012-0168-Y], http://hal.inria.fr/hal-00906826

[23] D. BISWAS, B. GENEST. *Minimal Observability and Privacy Preserving Compensation for Transactional Services*, in "Discrete Event Dynamic Systems", December 2013, vol. accepted, http://hal.inria.fr/hal-00916645

[24] G. DELAVAL, É. RUTTEN, H. MARCHAND. *Integrating discrete controller synthesis into a reactive programming language compiler*, in "Discrete Event Dynamic Systems", 2013, vol. 23, n^{o} 4, pp. 385-418 [*DOI :* 10.1007/s10626-013-0163-5], http://hal.inria.fr/hal-00863286

[25] L. HELOUET, H. MARCHAND, B. GENEST, T. GAZAGNAIRE. *Diagnosis from Scenarios*, in "Discrete Event Dynamic Systems", March 2013, n^{o} 10626 [*DOI :* 10.1007/s10626-013-0158-2], http://hal.inria.fr/hal-00879441

[26] T. JÉRON, M. VEANES, B. WOLFF. *Symbolic Methods in Testing (Dagstuhl Seminar 13021)*, in "Dagstuhl Reports", 2013, vol. 3, n^{o} 1, pp. 1–29 [*DOI :* 10.4230/DAGREP.3.1.1], http://hal.inria.fr/hal-00945878

[27] 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^{o} 2, pp. 396-408 [*DOI :* 10.1109/TAC.2013.2283093], http://hal.inria.fr/hal-00903452

International Conferences with Proceedings

[28] R. ABDALLAH, A. GOTLIEB, L. HÉLOUËT, C. JARD. *Scenario Realizability with Constraint Optimization*, in "FASE 2013", Rome, France, LNCS, March 2013, vol. 7793, pp. 194-209, http://hal.inria.fr/hal-00840393

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

[30] S. AKSHAY, I. DINCA, B. GENEST, A. STEFANESCU. *Implementing Realistic Asynchronous Automata*, in "IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science, FSTTCS 2013", Guwahati, India, December 2013, pp. 213-224, http://hal.inria.fr/hal-00920776

[31] S. AKSHAY, L. HELOUET, C. JARD, D. LIME, O. H. ROUX. *Robustness of Time Petri Nets under architectural constraints*, in "Formal Modeling and Analysis of Timed Systems", Warwik, United Kingdom, M. JURDZINSKI, D. NICKOVIC (editors), LNCS, Springer, September 2013, vol. 7595, pp. 11-26, http://hal.inria.fr/hal-00879818

[32] P. BALLARINI, N. BERTRAND, A. HORVATH, M. PAOLIERI, E. VICARIO. *Transient Analysis of Networks of Stochastic Timed Automata using Stochastic State Classes*, in "QEST - 10th International Conference on Quantitative Evaluation of Systems", Buenos Aires, Argentina, LNCS, Springer, 2013, vol. 8054, pp. 355-371 [*DOI :* 10.1007/978-3-642-40196-1_30], http://hal.inria.fr/hal-00915026

[33] N. BERTRAND, E. FABRE, S. HAAR, S. HADDAD, L. HELOUET. *Active diagnosis for probabilistic systems*, in "FOSSACS'2014", Grenoble, France, A. MUSCHOLL (editor), Springer, April 2014, http://hal.inria.fr/hal-00930919

[34] N. BERTRAND, P. FOURNIER. *Parameterized verification of many identical probabilistic timed processes*, in "FSTTCS - 33rd IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science", Guwahati, India, 2013, http://hal.inria.fr/hal-00914263

[35] N. BERTRAND, P. SCHNOEBELEN. *Solving Stochastic Büchi Games on Infinite Decisive Arenas*, in "QAPL - 11th International Workshop on Quantitative Aspects of Programming Languages and Systems", Rome, Italy, EPTCS, 2013, vol. 117, pp. 116-131 [*DOI :* 10.4204/EPTCS.117.8], http://hal.inria.fr/hal-00906831

[36] B. BOLLIG, A. CYRIAC, L. HELOUET, A. KARA, T. SCHWENTICK. *Dynamic Communicating Automata and Branching High-Level MSCs*, in "LATA 2013", bilbao, Spain, A. HORIA DEDIU, C. MARTÍN-VIDE, B. TRUTHE (editors), LNCS, Springer, April 2013, vol. 7810, pp. 177-189, http://hal.inria.fr/hal-00879353

[37] A. BOUILLARD, C. JARD, A. JUNIER. *Some Synchronization Issues in OSPF Routing*, in "DCNET13: the 4th International Conference on Data Communication Networking", Reykjavik, Iceland, M. S. OBAIDAT, J. L. SEVILLANO, Z. ZHANG (editors), SciTePress, July 2013, pp. 5-14, http://hal.inria.fr/hal-00840350

[38] A. BOUILLARD, A. JUNIER, B. RONOT. *Impact of Rare Alarms on Event Correlation*, in "CNSM - 9th international Conference on Network and Service Management", Zürich, Switzerland, October 2013, http://hal.inria.fr/hal-00920685

[39] L. CLEMENTE, F. HERBRETEAU, A. STAINER, G. SUTRE. *Reachability of Communicating Timed Processes*, in "FoSSaCS - 16th International Conference on Foundations of Software Science and Computation Structures, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS - 2013", Rome, Italy, F. PFENNING (editor), Lecture Notes in Computer Science, Springer, March 2013, vol. 7794, pp. 81-96, Extended version [*DOI :* 10.1007/978-3-642-37075-5_6], http://hal.inria.fr/hal-00744085

[40] G. DELAVAL, N. DE PALMA, S. M.-K. GUEYE, H. MARCHAND, É. RUTTEN. *Discrete Control of Computing Systems Administration: a Programming Language supported Approach*, in "European Control Conference", Zurich, Switzerland, M. MORARI (editor), July 2013, pp. 117-124, http://hal.inria.fr/hal-00863276

[41] Y. FALCONE, H. MARCHAND. *Runtime Enforcement of K-step Opacity*, in "52nd IEEE Conference on Decision and Control", Florence, Italy, T. PARISINI, R. TEMPO (editors), 2013, http://hal.inria.fr/hal-00863223

[42] B. GENEST, H. GIMBERT, A. MUSCHOLL, I. WALUKIEWICZ. *Asynchronous Games over Tree Architectures*, in "ICALP - 40th International Colloquium on Automata, Languages, and Programming- 2013", RIGA, Latvia, F. V. FOMIN, R. FREIVALDS, M. KWIATKOWSKA, D. PELEG (editors), LNCS, Springer, July 2013, vol. 7966, pp. 275-286 [*DOI :* 10.1007/978-3-642-39212-2_26], http://hal.inria.fr/hal-00916615

[43] C. HOUNKONNOU, E. FABRE. *Enhanced OSPF Graceful Restart*, in "IFIP/IEEE International Symposium on Integrated Network Management", Ghent, Belgium, May 2013, http://hal.inria.fr/hal-00931798

[44] L. JEZEQUEL, E. FABRE, V. KHOMENKO. *Factored Planning: From Automata to Petri Nets*, in "International Conference on Application of Concurrency to System Design (ACSD)", Barcelone, Spain, July 2013, http://hal.inria.fr/hal-00931844

[45] S. PINISETTY, Y. FALCONE, T. JÉRON, H. MARCHAND. *Runtime Enforcement of Timed Response Properties*, in "Software Verification and Testing, track of the Symposium on Applied Computing ACM-SAC 2014", Gyeongju, Korea, Republic Of, 2014, http://hal.inria.fr/hal-00907571

[46] L. RICKER, H. MARCHAND. *A parity-based architecture for decentralized discrete-event control*, in "American Control Conference", Washigton, United States, 2013, pp. 5678 - 5684, http://hal.inria.fr/hal-00828714

Books or Proceedings Editing

[47] T. Jéron, M. Veanes, B. Wolff (editors). , *Symbolic Methods in Testing*, Dagstuhl Publishing, May 2013, vol. 3, 29 p. , Report from Dagstuhl Seminar 13021 [*DOI :* 10.4230/DagRep.3.1.1], http://hal.inria.fr/hal-00823739

[48] H. Marchand, L. Hélouët (editors). , *Modélisation des systèmes réactifs*, Journal européen des systèmes automatisés, Lavoisier, 2013, vol. 47, 260 p. , http://hal.inria.fr/hal-00906737

Research Reports

[49] R. Abdallah, A. Gotlieb, L. Hélouët, C. Jard. , *Scenario realizability with constraint optimization*, January 2013, http://hal.inria.fr/hal-00769656

[50] A. Bouillard, A. Junier, B. Ronot. , *Alarms correlation in telecommunication networks*, Inria, June 2013, n^{o} RR-8321, 17 p. , http://hal.inria.fr/hal-00838969

[51] S. Chédor, T. Jéron, C. Morvan. , *Test Generation from Recursive Tile Systems*, Inria, January 2013, n^{o} RR-8206, 32 p. , http://hal.inria.fr/hal-00778134

[52] G. Kalyon, T. Le Gall, H. Marchand, T. Massart. , *Symbolic Supervisory Control of Distributed Systems with Communications*, Inria, March 2013, n^{o} RR-8260, http://hal.inria.fr/hal-00801840

Other Publications

[53] N. Bertrand, P. Fournier, A. Sangnier. , *Playing with probabilities in Reconfigurable Broadcast Networks*, 2014, http://hal.inria.fr/hal-00929857

References in notes

[54] , *Z.120: Message Sequence Charts (MSC)*, International Telecommunication Union, 2011

[55] P. Darondeau, L. Hélouët, M. Mukund. *Assembling Sessions*, in "Automated Technology for Verification and Analysis, 9th International Symposium, ATVA 2011", Lecture Notes in Computer Science, Springer, 2011, vol. 6996, pp. 259-274

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

[57] 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

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

[59] 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

[60] 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

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

[62] W. ZIELONKA. *Notes on Finite Asynchronous Automata*, in "R.A.I.R.O - Informatique Théorique et Applications", 1987, vol. 21, n^{o} 2, pp. 99-135