



INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET EN AUTOMATIQUE

Project-Team Distribcom

*Distributed Models and Algorithms for the
Management of Telecommunication
Systems*

Rennes - Bretagne-Atlantique

Theme : Networks and Telecommunications

Activity
R *eport*

2009

Table of contents

1. Team	1
2. Overall Objectives	1
2.1. Objectives of the team	1
2.2. Highlights	2
3. Scientific Foundations	2
3.1. Overview of the needed paradigms	2
3.2. Models of concurrency: nets, scenarios, event structures, and their variants	3
3.2.1. Scenarios.	3
3.2.2. Event structures.	3
3.2.3. Nets and languages of scenarios.	3
3.2.4. Extensions and variants.	3
3.2.5. Handling dynamic changes in the systems.	4
4. Software	4
5. New Results	4
5.1. Fundamentals results and algorithms: distributed unfoldings	4
5.2. Fundamentals results and algorithms: distributed planning	5
5.3. Fundamentals and algorithms: communication with messages and Scenarios	5
5.4. Fundamentals and algorithms: timed models	6
5.5. Fundamentals and algorithms: probabilistic models	7
5.6. Specific studies: Web services orchestrations	7
5.7. Specific studies: Active XML documents	8
5.8. Specific studies: security	9
5.9. Specific studies: Distributed optimization in network management	9
6. Other Grants and Activities	10
6.1. ANR Docflow	10
6.2. CREATE ActiveDoc	10
6.3. ANR Dots	11
6.4. Associated team DST	11
6.5. PHC FAST	11
6.6. DISC	12
6.7. Joint Bell Labs INRIA Laboratory	12
6.8. ADR HiMa	12
6.9. ANR Pegase	13
7. Dissemination	13
7.1. Partner-Countries	13
7.2. Scientific animation	13
7.3. Teaching	14
7.4. Visits and invitations	14
8. Bibliography	15

1. Team

Research Scientist

Albert Benveniste [DR INRIA, part-time, HdR]
Éric Fabre [CR INRIA , HdR]
Loïc Hérouët [CR INRIA]
Blaise Genest [CR CNRS]

Faculty Member

Claude Jard [PROF. ENS CACHAN , HdR]
Anne Bouillard [ASSISTANT PROF. ENS CACHAN]

Technical Staff

Deepak Bhatia [Associate Engineer, until June 30th]

PhD Student

Sidney Rosario [INRIA grant, until September 30th]
Bartosz Grabiec [INRIA grant]
Ajay Kattepur [INRIA grant, since August 1st]
Loïc Jézéquel [ENS Cachan grant, since September 1st]
Carole Hounkonnou [INRIA grant, since November 1st]

Post-Doctoral Fellow

Benoît Masson [Since October 15th]

Administrative Assistant

Laurence Dinh

2. Overall Objectives

2.1. Objectives of the team

The DistribCom team addresses models and algorithms for distributed network and service management, and the distributed management of Web services and business processes.

Today, research on network and service management as well as Web Services mainly focuses on issues of software architecture and infrastructure deployment. However, these areas also involve algorithmic problems such as fault diagnosis and alarm correlation, testing, QoS evaluation, negotiation, and monitoring. The DistribCom team develops the foundations supporting such algorithms. Our algorithms are model-based. Our research topics are therefore structured as follows:

1. *Fundamentals of distributed observation and supervision of concurrent systems*: this provides the foundations for deriving models and algorithms for the above mentioned tasks.
2. *Self-modeling*: for obvious reasons of complexity, our models cannot be built by hand. We thus address the new topic of self-modeling, i.e., the automatic construction of models, both structural and behavioral.
3. *Algorithms for distributed management of telecommunications systems and services*.
4. *Web Services orchestrations, functional and QoS aspects*.
5. *Active XML peers for Web scale data and workflow management*.

Our main industrial ties are with Alcatel-Lucent, on the topic of networks and service management.

2.2. Highlights

- In a joint work with Pierre Peloso from Alcatel-Lucent Bell Labs France (ALBLF), Eric Fabre has developed a novel decentralized and adaptive scheme for the optimal tuning of amplification gains in photonic networks. Two patents have been derived from this work, jointly registered by INRIA and ALBLF. This work is part of the Research Action on *High Manageability* in the framework of the Joint ALU-Bell Labs / INRIA laboratory.
- In a joint work with Nathalie Bertrand (Vertecs team) and Hugo Gimbert, Blaise Genest has published a paper entitled *Qualitative Determinacy and Decidability of Stochastic Games with Signals* in the conference Logic In Computer Science (LICS), 2009 [18]. This work was performed in the framework of the DST associated team.

3. Scientific Foundations

3.1. Overview of the needed paradigms

Management of telecommunications networks and services, and Web services, involves the following algorithmic tasks:

Observing, monitoring, and testing large distributed systems: Alarm or message correlation is one of the five basic tasks in network and service management. It consists in causally relating the various alarms collected throughout the considered infrastructure—be it a network or a service sitting on top of a transport infrastructure. Fault management requires in particular reconstructing the set of all state histories that can explain a given log of observations. Testing amounts to understanding and analyzing the responses of a network or service to a given set of stimuli; stimuli are generally selected according to given test purposes. All these are variants of the general problem of *observing* a network or service. Networks and services are large distributed systems, and we aim at observing them in a distributed way as well, namely: logs are collected in a distributed way and observation is performed by a distributed set of supervising peers.

Quality of Service (QoS) evaluation, negotiation, and monitoring: QoS issues are a well established topic for single domain networks or services, for various protocols — e.g., Diffserv for IP. Performance evaluation techniques are used that follow a “closed world” point of view: the modeling involves the overall traffic, and resource characteristics are assumed known. These approaches extend to some telecommunication services as well, e.g., when considering (G)MPLS over an IP network layer.

However, for higher level applications, including composite Web services (also called *orchestrations*, this approach to QoS is no longer valid. For instance, an orchestration using other Web services has no knowledge of how many users are calling the same Web services. In addition, it has no knowledge of the transport resources it is using. Therefore, the well developed “closed world” approach can no longer be used. *Contract* based approaches are considered instead, in which a given orchestration offers promises to its users on the basis of promises it has from its subcontracting services. In this context, contract composition becomes a central issue. Monitoring is needed to check for possible breaching of the contract. Countermeasures would consist in reconfiguring the orchestration by replacing the failed subcontracted services by alternative ones.

The DistribCom team focuses on the algorithms supporting the above tasks. Therefore models providing an adequate framework are fundamental. We focus on models of discrete systems, not models of streams or fluid types of models. And we address the distributed and asynchronous nature of the underlying systems by using models involving only local, not global, states, and local, not global, time. These models are reviewed in section 3.2. We use these mathematical models to support our algorithms and we use them also to study and develop formalisms of Web services orchestrations and workflow management in a more general setting.

3.2. Models of concurrency: nets, scenarios, event structures, and their variants

For Finite State Machines (FSM), a large body of theory has been developed to address problems such as: observation (the inference of hidden state trajectories from incomplete observations), control, diagnosis, and learning. These are difficult problems, even for simple models such as FSM's. One of the research tracks of DistribCom consists in extending such theories to distributed systems involving concurrency, i.e., systems in which both time and states are local, not global. For such systems, even very basic concepts such as "trajectories" or "executions" need to be deeply revisited. Computer scientists have for a long time recognized this topic of concurrent and distributed systems as a central one. In this section, we briefly introduce the reader to the models of scenarios, event structures, nets, languages of scenarios, graph grammars, and their variants.

3.2.1. Scenarios.

The simplest concept related to concurrency is that of a finite execution of a distributed machine. To this end, scenarios have been informally used by telecom engineers for a long time. In scenarios, so-called "instances" exchange asynchronous messages, thus creating events that are totally ordered on a given instance, and only partially ordered by causality on different instances (emission and reception of a message are causally related). The formalization of scenarios was introduced by the work done in the framework of ITU and OMG on High-level Message Sequence Charts and on UML Sequence Diagrams in the last ten years, see [44], [48]. This allowed in particular to formally define infinite scenarios, and to enhance them with variables, guards, etc [53], [50], [51]. Today, scenarios are routinely offered by UML and related software modeling tools.

3.2.2. Event structures.

The next step is to model sets of finite executions of a distributed machine. *Event structures* were invented by Glynn Winskel and co-authors in 1980 [47], [54]. Executions are sets of events that are partially ordered by a *causality* relation. Event structures collect all the executions by superimposing shared prefixes. Events not belonging to a same execution are said in *conflict*. Events that are neither causally related nor in conflict are called *concurrent*. Concurrent processes model the "parallel progress" of components.

Categories of event structures have been defined, with associated morphisms, products, and co-products, see [55]. Products and co-products formalize the concepts of parallel composition and "union" of event structures, respectively. This provides the needed apparatus for composing and projecting (or abstracting) systems. Event structures have been mostly used to give the semantics of various formalisms or languages, such as Petri nets, CCS, CSP, etc [47], [54]. We in DistribCom make a nonstandard use of these, e.g., we use them as a structure to compute and express the solutions of observation or diagnosis problems, for concurrent systems.

3.2.3. Nets and languages of scenarios.

The next step is to have finite representations of systems having possibly infinite executions. In DistribCom, we use two such formalisms: *Petri nets* [49], [40] and *languages of scenarios* such as High-level Message Sequence Charts (HMSC) [44], [51]. Petri nets are well known, at least in their basic form, we do not introduce them here. We use so-called *safe* Petri Nets, in which markings are boolean (tokens can be either 0 or 1); and we use also variants, see below.

3.2.4. Extensions and variants.

Two extensions of the basic concepts of nets or scenario languages are useful for us. Nets or scenario languages enriched with variables, actions, and guards, are useful to model general concurrent and distributed dynamical systems in which a certain discrete abstraction of the control is represented by means of a net or a scenario language. Manipulating such *symbolic nets* requires using abstraction techniques. Time Petri nets and network of timed automata are particular cases of symbolic nets. Probabilistic Nets or event structures: Whereas a huge literature exists on stochastic Petri nets or stochastic process algebras (in computer science), randomizing *concurrent models*, i.e., with ω 's being concurrent trajectories, not sequential ones, has been addressed only since the 21st century. We have contributed to this new area of research.

3.2.5. Handling dynamic changes in the systems.

The last and perhaps most important issue, for our applications, is the handling of dynamic changes in the systems model. This is motivated by the constant use of dynamic reconfigurations in management systems. Extensions of net models have been proposed to capture this, for example the *dynamic nets* of Vladimiro Sassone [39] and *net systems* [41]. For the moment, such models lack a suitable theory of unfoldings.

4. Software

4.1. SOFAT : a scenario toolbox

Participants: Loïc Hélouët, Deepak Bhatia, Gunjan Aggarwal.

The SOFAT toolbox is a scenario manipulation toolbox. Its aim is to implement all known formal manipulations on scenarios. The toolbox implements several formal models such as partial orders, graph grammars, graphs, and algorithm dedicated to these models (Tarjan, cycle detection for graphs, Caucal's normalization for graph grammars, etc.). The SOFAT toolbox is permanently updated to integrate new algorithms. It is freely available from Distribcom's website : <http://www.irisa.fr/distribcom/Prototypes/SOFAT/index.html>.

SOFAT is a demonstrator and a support for all our proposals in standardization committees at ITU. This involvement in standardization is also the occasion for numerous contacts with MSC users (France Telecom, Nokia, Motorola), but also with CASE tool designers at IBM.

Last year, the V3 of SOFAT had been released. Ever since, SOFAT has been extended with new functionalities such as scenario based diagnosis, and model checking of globally cooperative HMSCs. Version 4 of SOFAT will be published as soon as these new functionalities are documented.

Time analysis of scenarios was developed this summer by G. Aggarwal, and is currently under integration in the tool. The principle of the analysis consists in unfolding an annotated HMSC, transform it into a colored stochastic Petri Net, and then run a simulation to obtain performance indicators.

5. New Results

5.1. Fundamentals results and algorithms: distributed unfoldings

Participants: Eric Fabre, Blaise Genest.

This is a joint work with our former postdoc Agnes Madalinski, now assistant professor at the University of Santiago in Chile.

Distributed systems can be modeled as networks of interacting components, for example networks of Petri nets, or networks of automata. The simplest way to combine components into larger systems is to take their product. It is well known that the unfolding of a product Petri net can be expressed as the product (in a specific sense) of the unfoldings of the components. The factorized form of unfoldings is generally more compact than the expanded product, because each factor only represents its local conflicts and does not have to display the choices that are made by other components. This property was the basis of our previous results about distributed diagnosis.

With Agnes Madalinski, we explored the construction of finite complete prefixes (FCP) for distributed systems [45], [14]. More precisely, our goal was to obtain a FCP in factorized form, when the underlying system is expressed as a product of components. FCP represent in a compact manner all possible behaviors, and states, that a concurrent system can reach, which makes them a central tool for model checking applications for example. It is likely that factorized forms will be more compact, and will thus open the way to new distributed model checking techniques. A trivial but impractical solution consists in computing a FCP of the global system, and then deriving from it its factorized form. We explored a fully decentralized method to obtain directly an FCP of each component, such that the product of all these local FCP would yield an FCP of the global system. A solution was proposed in some limited situations, through the notion of extended stopping point. This is a first step to solving problems like distributed reachability analysis, distributed planning, etc., as they are now explored in the FAST cooperation, and in the DISC European project.

This year, we also developed another technique, similar in its goal to unfolding but very different in its methodology. Instead of considering a true concurrency model which does not exhibit redundancy by construction, we consider a priori several interleavings resulting in the same distributed execution. Then, we perform on the fly partial order reduction, such that the redundancy is kept as small as possible [10]. We showed that performing a breadth first search, no redundant path needs to be considered, whereas in depth first search (DFS), some limited redundancy needs to be used. It is quite similar to unfolding. The experiment we performed showed that the technique is simple yet very efficient, and that the amount of redundancy needed in DFS is very limited in practice.

5.2. Fundamentals results and algorithms: distributed planning

Participants: Eric Fabre, Loig Jézéquel.

A planning problem consists in organizing some actions in order to reach an objective. Formally, this is equivalent to finding a path from an initial state to a goal state in a huge automaton. The latter is specified by a collection of resources, that may be available or not, and actions that consume and produce resources. In the case of optimal planning, actions have a cost, and the objective is to find a path of minimal cost to the goal.

Our interest in this problem is threefold. First, it is naturally an instance of a concurrent system, given that actions have local effects on resources. Secondly, it is a weak form of an optimal control problem for a concurrent/distributed system. Finally, we are interested in distributed solutions to such problems, which is a very hot topic in the planning community under the name of “factored planning.”

Our contribution to this topic is the first optimal factored planning algorithm [23]. It is based on the observation that a planning problem can be translated into a network of components, modeled as weighted automata in our case. We have then designed a message passing procedure on this network, based on weighted automata calculus, where each component determines its part of the best global action plan using only local information: its local model, and messages received from neighbors about shared actions. This distributed solution resolves both a constraint solving problem, and an optimization problem. The optimal plan is given as a tuple of partially synchronized local plans, therefore as a partial order of actions. We are currently experimenting this approach on benchmarks of planning problems.

5.3. Fundamentals and algorithms: communication with messages and Scenarios

Participants: Loïc Hérouët, Blaise Genest.

In this paragraph, we collect our fundamental results regarding the models and algorithms we use for communicating systems, and in particular, scenarios.

A major challenge with models communicating with messages (e.g.: scenarios) is to *exhibit good classes of models* allowing users to *specify easily complex distributed systems* while *preserving the decidability* of some key problems, such as diagnosis, equality and intersection. Furthermore, when these problems are decidable for the designed models, the second challenge is to design algorithms to keep the *complexity low enough* to allow *implementation in real cases*.

For the last three years, we have developed a new model of scenarios, namely Causal HMSCs [13], in order to specify complex telecommunication protocols, such as sliding windows protocols. The main novelty of the approach is to allow for an independence alphabet on each process, instead of the rigid total order of HMSCs. Interestingly, many problems on this model remain decidable without requiring (existential) bounds on the message queues. The decidable questions we have considered are diagnosis and comparison (equality, intersection) with other Causal HMSCs, and we gave the optimal associated algorithms. However, when comparing Causal HMSCs with other models (logics, communicating automata, or even Causal HMSCs built with different independence alphabets), the problems turn out to be undecidable, unless there is an (existential) bounds on the number of messages present in any channel at a given time. We thus consider the problem to know whether a given system is existentially bounded, modeled as communicating automata [11]. We proved that this problem is undecidable. However, we give algorithms to solve the problem in a non trivial subclass.

Our last work in the topic of algorithms for Scenarios was to consider basic test on languages of scenarios, depicted as graphs, to be included in our tool SOFAT. The problem is that a graph of scenarios might exhibit counterintuitive behaviors, as an event appearing much further than a given event in the graph can happen at the same time or even before this event. We call such a case disorder [12], we give quantification of this disorder, and give an efficient algorithm to compute the worst disorder, that is to display its “most counterintuitive” behavior.

We have also considered an extension of coregions in HMSCs. A coregion is a part of a process description in which the ordering of events is relaxed. Usually, the ordering of events on a single process is a total order. The Z.120 standard [44] also allows for general orderings, that is a replacement of the total ordering imposed on a process by a partial ordering. However, coregions are limited to a finite set of events. We have extended the original formalism to allow for infinite coregions containing partial orderings. Within this context, we also have provided algorithms to detect discrepancies between the visual ordering of events and their actual ordering imposed by the semantics (this notion is usually called “race condition”). This work was done during a collaboration with Masaryk University (Czech Republic).

5.4. Fundamentals and algorithms: timed models

Participants: Claude Jard, Anne Bouillard, Bartosz Grabiec.

Our work on timed models was focused on the study and use of two different techniques: unfoldings of network of timed automata (and Time Petri nets) and the network calculus. The goals are supervision with time and performance evaluation.

In the context of the PhD of B. Grabiec, we are trying to extend the model of networks of timed automata by modifying the time semantics to take into account more realistic time models, allowing for instance some drift between clocks of different components. This has some consequences on the way of representing generic time zones and unfoldings. The focus this year was to introduce parameters in models in order to facilitate the modelling phase. The interest is to handle symbolically both the timings and the parameters. Unfoldings are an efficient method to determine the causal relations between the events in a system. Given a partial observation, as a list of actions, we propose to use our method to determine the causal relation between events in the model that explain the observation [24]. We can also synthesise parametric constraints associated with these explanations. In collaboration with the IRCCyN’s group in the DOTS project, this method was implemented in the ROMEO tool.

Network Calculus is a quite recent theory developed to compute deterministic worst-case bounds in queuing networks. Computing such bounds is necessary when dealing with real-time and critical systems (that can be found for example in embedded systems of airplanes or cars). The Network Calculus is based on the (min,plus) algebra. It models constraints on arrival and output processes by means of arrival and service curves. Our work has focused on three main aspects:

1. We first studied the algorithmic aspects of the Network Calculus operators, namely the (min,plus) convolution, the (min,plus)-deconvolution and the sub-additive closure. We have exhibited a stable class of functions regarding those operators and have given efficient algorithms to compute them [38]. A small software COINC has been written to implement those algorithms and a first version is now available (<http://www.istia.univ-angers.fr/~lagrange/spip.php?article21>).
2. We then studied the composition of network elements in presence of cross-traffic. Our contribution concerns two kinds of scenarios:
 - One flow in a network in presence of independent cross-traffic, that has to be transmitted from a given source to a given destination. What is the best path concerning the worst-case delay/backlog for that flow?
 - One flow on a fixed path interfering with dependent cross-traffic. Can we compute a service curve for the effective traffic for the flow?

To answer the first question, we derived some polynomial-time algorithms that can be seen as shortest-path algorithms with functional weights on the arcs, instead of constant weights as in the classical case ([36]). To answer the second question, we introduce a new operator, the multidimensional convolution. It appears that we are not able to compute the multidimensional convolution (hence the effective service curve) in polynomial time, but, thanks to linear programming, we are able to derive bounds on the delay and backlog in polynomial time([37]).

3. Finally some work has been initiated thanks to the associated team CASDS. It concerns the study of multi-mode Network Calculus: in classical Network calculus, constraints are static. Here, they evolve with time, and this evolution is modeled with a finite-state automaton. A first study concerns the block-writing server (a server either serves packets with a given constraint on the guaranteed service or serves nothing). Corresponding paper [21] has been accepted at RTAS'09.

5.5. Fundamentals and algorithms: probabilistic models

Participants: Albert Benveniste, Anne Bouillard, Sidney Rosario.

Our work on *true concurrency probabilistic models* is joint work with our former PhD student Samy Abbes, now Maître de Conférences in Mathematics at Paris VI, PPS Laboratory. We have established the foundations for probabilistic models of distributed and concurrent systems, in which traces, not interleaving sequences, are randomized. We have been able to extend our Markov net model beyond nets with finite confusion (those which yield locally finite event structures). Markov nets are now defined in most general cases. Extensions required a fine understanding of probabilistic fairness. This result was published in FOSSACS'2009 [17]. Regarding true concurrency *compositional* probabilistic models, we are currently working at extending our previous results concerning the fully probabilistic synchronous product of two Markov chains—this differs from products of Probabilistic Automata, which yield mixed probabilistic/nondeterministic models.

On another direction, we have studied probabilistic path criticality for stochastic Petri nets. Targeted applications are workflow nets such as used in the modeling of orchestrations. In concurrent real-time processes, the speed of individual components has a double impact: on the one hand, the overall latency of a compound process is affected by the latency of its components. But, if the composition has race conditions, the very outcome of the process will also depend on the latency of component processes. Using stochastic Petri nets, we investigate the probability of a transition occurrence being critical for the entire process, i.e. such that a small increase or decrease of the duration of the occurrence entails an increase or decrease of the total duration of the process. The first stage of the analysis focuses on occurrence nets, as obtained by partial order unfoldings, to determine criticality of events; we then lift to workflow nets to investigate criticality of transitions inside a workflow. This is joint work with Stefan Haar. It has been published in [20].

5.6. Specific studies: Web services orchestrations

Participants: Sidney Rosario, Ajay Kattapur, Albert Benveniste, Claude Jard, Anne Bouillard.

Web services *orchestrations* and *choreographies* refer to the composition of several Web services to perform a co-ordinated, typically more complex task. We decided to base our study on a simple and clean formalism for WS orchestrations, namely the ORC formalism proposed by Jayadev Misra and William Cook [46].

Main challenges related to Web services QoS (Quality of Service) include: 1/ To model and quantify the QoS of a service. 2/ To establish a relation between the QoS of queried Web services and that of the orchestration (contract composition); 3/ To monitor and detect the breaching of a QoS contract, possibly leading to a reconfiguration of the orchestration. Typically, the QoS of a service is modeled by a *contract* (or Service Level Agreement, SLA) between the provider and consumer of a given service. To account for variability. In 2007, we proposed soft probabilistic contracts specified as probabilistic distributions involving the different QoS parameters. We studied *contract composition* for such contracts, see 2007 Activity Report. Since then, the following directions have been developed.

First, we observed that the current practice regarding SLA seems to ignore the fact that *orchestrations may not be monotonic*, meaning that it is possible that if a called service improves its performance, then the overall orchestration performance nevertheless *decreases*. This artifact does not occur in SLA for networks; it can, however, occur for Web services because the latter involve complex interactions between control, data, and time (through the use of timers). In [22] we have established conditions (necessary and sufficient ones) ensuring monotonicity of orchestrations.

In a contract-based paradigm, one important duty of the orchestration is *QoS contract monitoring*. We have developed statistical techniques for soft probabilistic QoS contract monitoring. This work is part of the invited paper [52] of the ICWS07 special issue of *IEEE Transactions on Services Computing* and was presented at IM 09, special track on new results [26]. This monitoring service has been added to the *TOrQuE (Tool for Orchestration simulation and Quality of service Evaluation)* tool developed since 2007 by Sidney Rosario regarding soft probabilistic QoS contracts for Web services orchestrations.

Last, we have further extended our approach by soft probabilistic contracts to general QoS parameter, i.e., beyond response time. In particular, we now encompass composite parameters, which are thus only partially, not totally, ordered. We have developed a general algebra to capture how QoS parameters are transformed while traversing the orchestration and we have extended our study of monotonicity. Finally, we have developed corresponding contract composition procedures, which are iterative and must involve negotiation, unlike the simple case that we studied previously. This was presented at ICWS 09 [27] and this paper was selected for the special issue of *International Journal of Web Services Research* [15].

5.7. Specific studies: Active XML documents

Participants: Albert Benveniste, Loïc H elou et, Blaise Genest, Debmalya Biswas [PhD].

The language *Active XML* or *AXML* is an extension of XML which allows to enrich documents with *service calls* or *sc's* for short. These *sc's* point to web services that, when triggered, access other documents; this materialization of *sc's* produces in turn *AXML* code that is included in the calling document. One therefore speaks of dynamic or intentional documents; note in particular that materialization can be *total* (inserting data in XML format) or *partial* (inserting *AXML* code containing further *sc's*). *AXML* has been developed by the GEMO team at INRIA Futurs, headed by Serge Abiteboul; it allows to set up P2P systems around repositories of *AXML* documents (one repository per peer).

We are cooperating with the GEMO team (Serge Abiteboul) and the LABRI laboratory in Bordeaux (Anca Muscholl) to explore the behavioral semantics of *AXML* in the framework of the former ASAX INRIA-ARC (see the 2006 activity report), and to analyze such systems in the framework of the Docflow and Activedoc projects, see 6.1, 6.2 below.

AXML allows for complex and evolving systems. The challenge is thus to be able to check whether an *AXML* system (as a mail order system) satisfies some basic properties (e.g.: is it possible that a product is mailed after the customer canceled his order?), even for systems beyond the so-called “positive” class [35] (where nothing can get deleted). This problem is difficult due to unboundedness of the number of products (infinite states), unboundedness of the names of the product (infinite alphabet), distribution and asynchrony. We succeeded into defining a class of models subsuming positive systems, allowing deletion and unbounded paths, but not too permissive (e.g. the alphabet is supposed finite, the depth of the document is bounded, and guards on the system cannot test for the non existence of a pattern), such that the basic problem we want to solve stay decidable. This year, we have addressed the problem of distribution and composition of *AXML* systems [32]. In the usual *AXML* framework, the locality of active documents is not considered, and guards can be evaluated over documents that are located at any place. However, in a real implementation, it is not possible to evaluate a guard over a part of a document owned by a distant site without some additional communication. We have proposed a distributed version of *Active XML*, in which guards can only be evaluated locally to a given site, and agents of the system interact using calls to external services. From a site, the interactions with the environment are only seen as interfaces, that depict the relation between the parameters sent to an implementation of an interface, and the expected results. Intuitively, interfaces describe the needed functionalities of a site

in terms of legal inputs/outputs. Then, a service implements an interface if it accepts all allowed parameters, and returns only results described by the interface. This implementation relation builds on the well-known notion of query containment, and can rapidly become undecidable if the considered data is infinite. We have shown sufficient conditions for decidability. Distributed AXML compose, just by providing in one system some implementations needed by the other one. They can also be abstracted (a part of the system is replaced by interfaces for it) to allow modular model checking.

Another main activity of the AXML project is to propose algorithms to bring transactional features to AXML services, and more generally to composition of web services, without breaking the confidentiality of peers and by using distributed techniques to keep high performance. It is the central theme of the PhD thesis of Debmalya Biswas, who defended in January 2009 [8]. The transactional feature considered was atomicity, which is ensured through a compensating (rollback) mechanism. For this, actions performed are logged. This year, we did experiments on different implementations of rollback (deferred rollback through commit, or in place) and compared them in [19].

5.8. Specific studies: security

Participants: Loïc H elou et, Blaise Genest.

We have performed some work on security issues in the context of the DOTS project, and within a collaboration with the VERTECS team. In DOTS, we are involved in a working group on non-interference. This year, we have mainly focused on two topics, anomaly detection, and covert channels discovery using information theory.

Anomaly detection with diagnosis techniques: We have extended the work initiated last year on anomaly detection. The proposed technique uses partial order diagnosis techniques to discover unusual behaviors, that can be seen as potential intrusions. The approach consists in comparing executions of a running system with a partial order model that describes the “normal” behaviors and interactions of a group of users with the system. When the observation does not correspond to an explanation in the model, (i.e. when diagnosis does not provide a solution for an observation) an alarm is raised. This work relies on the results of [42]. The work initiated last year [43] has been extended to deal with online diagnosis. Being able to perform online security monitoring is an important issue, as post mortem analysis can only take place after an intrusion occurred. We have proposed optimized online algorithms for intrusion detection, and shown under which conditions this online monitoring could be performed with finite memory. This extended work should be submitted to a journal next year.

Covert channels discovery using information theory: In the DOTS project, we have studied covert channels with the help of information theory. Roughly speaking, a covert channel is an obfuscated use of a system to create hidden communication between agents that are not allowed to exchange information. We have adapted work on channel capacities to discover covert flows of information. Namely, if we represent a distributed system with agents u_1, \dots, u_n as a transition system, a covert channel can be seen as a way to increase average mutual information between what an agent u_i does and what another agent u_j observes from the system.

This year, in the context of the DST associated team, we started to consider control in order to ensure security of distributed systems. We modeled the distributed system as a system with partial information (each process does not know precisely the state of the other process), where 2 processes have antagonist goal (a zero sum turn based game), with stochastic transition functions. In such cases, stochastic controllers are strictly more powerful than deterministic ones. We give a fix point algorithm to determine from which state there exists a controller for one process that ensures almost surely (resp. with positive probability) to stay in a safe zone. We provide constructively the associated controller, and characterize how much memory it needs. In the case where such a controller does not exist from a given state, we provide a controller for the other process ensuring with positive probability (resp. almost surely) that an insecure state can be reached. Again, we characterize precisely how much memory is needed.

5.9. Specific studies: Distributed optimization in network management

Participant: Eric Fabre.

This work represents part of our activities within the research group “High Manageability,” supported by the common lab of Alcatel-Lucent Bell Labs (ALBLF) and INRIA. It concerns a key issue for the autonomic management of photonic networks, i.e. optically routed networks. The problem concerns the fine tuning of wavelength reamplification gains at the input of each fiber, in order to optimize the optical signal to noise ratio (OSNR) at egress of the connection. The tuning of these gains directly influences the reach of a connection, that is the distance over which the signal can be transported optically, without necessity of an electronic regeneration. This in turn has a direct financial impact since less equipment is needed.

The problem is made complex by several phenomena. First of all, the total amount of power allowed in a fiber is limited (or equivalently each optical cross-connect has a bounded power budget). This implies that each node must select which wavelengths will be reamplified, and by how much. Secondly, the per-wavelength amplification gains are themselves limited, so an important loss on some connection in a link may have to be compensated for by several consecutive reamplifications along this connection. These two phenomena, combined with other non-linear effects, make this optimal tuning of gains a huge and complex network-scale optimization problem under constraints. The objective function is of course to maximize the OSNR of all connections in the network, and at the same time equalize these OSNR, so that long-range connections have the same quality as short range ones. For the moment, all these gains are manually adjusted, one by one, which is extremely difficult and suboptimal.

We have designed and tested an iterative and distributed solution to solve this problem: Each optical cross-connect in the network tunes its own reamplification gains, based on information provided by its neighbors. No global topology information is necessary, and convergence is guaranteed. The algorithm is adaptive to network changes: it redistributes optimally the power left by closed connections, and symmetrically pumps power in the less important connections to feed a newly created one.

Two patents have been derived from this work, jointly registered by ALBLF and INRIA. This delayed the publication work, that will take place in 2009. A demo of the prototype solution was performed at the Alcatel-Lucent Innovation Days, June 2009.

6. Other Grants and Activities

6.1. ANR Docflow

Participants: Albert Benveniste, Eric Fabre, Loïc H elou et, Blaise Genest, Debmalya Biswas, Beno t Masson.

Contract INRIA ANR-06-MDCA-005 January 2007 - April 2010

Docflow (<http://www.labri.fr/perso/anca/docflow/main.html>) is a national research project where Distribcom cooperates with INRIA’s GEMO team, and the LABRI/Bordeaux. It started in January 2007 and is scheduled to end in April 2010. It is a follow-up of the ARC Asax (see below). The aim of the docflow project is to model, analyze and monitor real life composite services, as tour operators (Opodo) or supply chains (DELL). It builds on the understanding between the Database community (data centric views) and the Discrete Event community (control centric), brought by the past ASAX meetings. The main tool is Active XML, see URL <http://activexml.net> on Active XML and Web services. So far, only a fragment of AXML was considered. It is called “positive AXML”, and have simplistic control (no move or deletion of data, only copy of nodes are possible at some given nodes, and every copy is possible in parallel). We try to develop a model where control can simulate workflow, and structured data (XML) can be used in the same formalism. This starting point will allow us to develop algorithms to analyze, monitor and optimize workflows with rich data.

6.2. CREATE ActiveDoc

Participants: Albert Benveniste, Eric Fabre, Loïc H elou et, Blaise Genest, Debmalya Biswas, Sidney Rosario, Beno t Masson.

Contract INRIA CREATE February 2007 - August 2011

Activedoc is funded by Région Bretagne, supporting the ANR Docflow project. It started in February 2007, for 18 months, and can be extended twice for 18 months. In addition to the Docflow program, it grants funding to study composite web services in a quantitative way. The fundamental models proposed in Docflow will be a starting point. For instance, developing methods to compose the Quality of Service of different web services is a difficult problem if one wants realistic values which are not too imprecise. Methods to elaborate and use contracts between heterogeneous services would thus be simplified.

6.3. ANR Dots

Participants: Claude Jard, Loïc Hélouët, Blaise Genest, Bartosz Grabiec.

Contract INRIA ANR-06-SETI January 2007 - December 2010

Dots (<http://www.lsv.ens-cachan.fr/anr-dots/>) is a national research project where Distribcom cooperates with the LSV/ENS Cachan, the LABRI/Bordeaux, the LAMSADE/Paris Dauphine and the IRCCyN/Nantes. It started in January 2007 and is scheduled to end in December 2010. The ambitious goal of the project is to consider open systems (that is interacting with other undefined systems) which are distributed and require timing information, in order to analyze concrete systems without abstracting one of these aspects. For instance, the interference between several systems require a combination of opened, distributed and timed information. Distribcom is in charge of the interaction of distributed systems with timing aspect (as timed Petri nets) or openness (as distributed controllers and distributed games).

6.4. Associated team DST

Participants: Loïc Hélouët, Blaise Genest, Thomas Gazagnaire, Anne Bouillard.

DST : Distributed systems, Supervision and Time.

Associated Team INRIA-NUS-Chennai — 2009-2011

This associated team is a collaboration with the National University of Singapore, Chennai Mathematical Institute and Institute for Mathematical Science in Chennai, and also involves members of the S4 team. The main research theme is to study supervision and time issues in distributed systems with the help of concurrency models, which follows and extend the work done in the former associated team CASDS. Two application areas are targeted: Real-time embedded systems and telecommunications systems and services. Although very different in nature, both areas make fundamental use of models of concurrency. Several types of formal models are considered: scenario languages, communicating automata and Petri-nets. More specifically, we work together on the following problems:

- Distributed Control of Concurrent Systems and the problem of synthesizing small controllers.
- Quantitative aspects of timed distributed systems.
- Qualitative Verification of timed constraints concurrent models.

Two long versions of paper we wrote two and three years ago have been accepted then published to top journals this year [13], [11], one considering the minimal control. On quantitative aspects of time, [21] have been published in a conference. Loïc Helouet and Philippe Darondeau visited in Chennai in January for 10 days, spending time at a Indo French Workshop. Blaise Genest visited Singapore for 2 weeks in February, spending time at SinFra (Singapore-French) conference. Loïc Helouet and Blaise Genest are going to visit NUS in early december. We received the visit of several PhD students from Chennai, S. Akshay in july for a week and Paul Soumya in November for 2 weeks. We also hired an intern from India in the summer working on implementing a tool to compute mean throughput of probabilistic and timed distributed system through simulation.

6.5. PHC FAST

Participants: Eric Fabre, Loig Jézéquel.

Funded by the French Ministry of Foreign Affairs. January 2008 - December 2009

This small exchange program involves two research teams of INRIA Rennes, DistribCom and S4 (Sophie Pinchinat), and a research team of NICTA (National Information and Communication Technologies institute of Australia) in Canberra, Australia. For the DistribCom part, the aim is to design modular planning algorithms. The stay of L. Jézéquel in November was funded by this program.

6.6. DISC

Participants: Eric Fabre, Loig Jézéquel.

European STREP project - Call FP7-ICT September 2008 - September 2011

Distributed Supervisory Control of Large and Complex Plants. This project involves as well team S4 (Ph. Darondeau), and a starting collaboration with Serge Haddad (LSV, ENS Cachan) will also be hosted by DISC. The main collaborations of DistribCom will be with the LSV, the University of Cagliari (Italy), the CWI (Amsterdam, NL), Ghent University (Belgium), the Czech Academy of Sciences (Czech Republic), and with Canadian and US partners that will soon be attached to DISC. Distribcom is in charge of three main workpackages related to 1/ the distributed optimal control of coupled MDP (Markov Decision Processes), 2/ distributed planning algorithms, and in particular distributed reachability tests, and 3/ the existence of distributed observers for a distributed system.

6.7. Joint Bell Labs INRIA Laboratory

Participants: Eric Fabre, Albert Benveniste.

Started January 2008.

The *Joint Bell Labs INRIA Laboratory* is the ongoing framework for the overall research cooperation between Alcatel-Lucent Bell Labs and INRIA. This joint Laboratory was launched in January 2008. It is a virtual lab, meaning that researchers remain hosted by their home institutions. The lab has the general area of *self-organizing networks* in its central focus. It is organized into three *Actions de Recherche (ADR)*:

- SelfNets (Self-Organizing networks), headed by Olivier Marcé (BellLabs) and Bruno Gaujal (INRIA);
- Semantic Networking, headed by Ludovic Noirie (BellLabs) and Pascale Vicat-Blanc (INRIA);
- High Manageability, HiMa, headed by Pierre Peloso (BellLabs) and Eric Fabre (INRIA, DistribCom).

Overall, the joint lab involves about 50 people. It is jointly headed by Olivier Audouin (BellLabs, president), and Albert Benveniste (INRIA, president of the Scientific Committee). The lab organizes bi-yearly seminars with progress reports and keynotes by key engineers from Alcatel-Lucent — the first one was about LTE (Long Term Evolution) by Denis Rouffet and the second one was about optical networks, by Paolo Fogliata, both from business divisions.

6.8. ADR HiMa

Participants: Eric Fabre, Albert Benveniste, Claude Jard, Anne Bouillard, Carole Hounkonnou.

Research Action "High Manageability", hosted by the common research laboratory of Alcatel-Lucent-Bell Labs and INRIA. June 2008 - June 2011.

This research group involves three INRIA teams, DistribCom, Madynes (O. Festor, INRIA Lorraine), and Mascotte (J.-C. Bermond) who joined the group recently in 2009. On the Alcatel-Lucent side, 5 persons of the PTI group (Packet Transport Infrastructure) are involved. It is jointly headed by P. Peloso (ALU, in replacement of M. Vigoureux) and E. Fabre (INRIA). The objective is to contribute to the autonomic networking trend, that is to design telecommunication networks that would be programmed by objectives, with minimal human operations, and that would then adapt themselves in order to reach these objectives. More specifically, this covers both the architectural and the algorithmic aspects of self-management methodologies. The activity is organised around several case-studies and working groups. In 2009, the team brought to maturity an activity about the optimal power allocation to wavelengths in photonic (meaning optically routed) networks, and worked on the problem of organization tasks for network maintenance with minimal service disruption, which corresponds to the starting PhD of Carole Hounkonnou.

The activities of HiMa also cover security issues for VOIP (studied by Madynes), and network defragmentation issues (studied by Mascotte), that is the joint optimization of connection paths for both the IP and the photonic layers. This represents a total of four PhD theses. In 2009, this group actively contributed to the preparation of "UniverSelf," an FP7 IP proposal led by ALU-Bell Labs, that aims at gathering the most significant teams in Europe dealing with autonomic networking.

6.9. ANR Pegase

Participant: Anne Bouillard.

Contract INRIA ANR-09-SEGI-009 October 2009 - October 2012

Pegase (Performances garanties dans les systèmes embarqués) is a national research project where DistribCom interacts both with academical partners (ENS Lyon and INRIA Rhône-Alpes) and with industrial partners (Thalès, ONERA and RT@W). It started in October 2009 and is scheduled to end in October 2012. The aim of Pegase is to develop the theory of Network Calculus and study the applicability to embedded networks (SpaceWire, AFDX). A prototype is planned to be developed.

7. Dissemination

7.1. Partner-Countries

Chili, United Kingdom, Singapore, Macao, Germany, India, Netherlands, Germany, Italy, Israel

7.2. Scientific animation

A. Benveniste is associated editor at large (AEAL) for the journal *IEEE Trans. on Automatic Control*. He is member of the Strategic Advisory Council of the Institute for Systems Research, Univ. of Maryland, College Park, USA. This year, he has been in the program committee of the conference WS-FM. A. Benveniste is president of the Scientific Committee of the *Joint Bell Labs INRIA Laboratory*. The lab organizes bi-yearly seminars with progress reports and keynotes by key engineers from Alcatel-Lucent. A. Benveniste has also co-organized, jointly with Philippe Jacquet, the Second joint BellLabs-INRIA workshop on the Fundamentals of Communication Networks and sErvices.¹

E. Fabre is associate editor (AE) for the journal *IEEE Trans. on Automatic Control*. He gave a plenary address at the first joint INRIA-Bell Labs workshop on the "Fundamentals of Communications and Networks," Murray Hill, June 1-2, 2009. Together with Pierre Peloso (Alcatel-Lucent), he presented a demo of the DOBLIN software prototype (Distributed Optimal Balancing of LIght in photonic Networks) during the Alcatel-Lucent Bell Labs Innovation Days, June 1-5, 2009. He has been member of the Program Committee of DX'09 (Int. Workshop on the Principles of Diagnosis).

¹Only facts related to the activities of DistribCom team are mentioned. Other roles or duties concern the S4 or Ssystem teams, to which A. Benveniste also belongs.

B. Genest is an elected member of the Comité National de la Recherche Scientifique for 2008-2012.

C. Jard has been in 2009 member of the Program Committee of the following international conferences: FORMATS, NOTERE, FMOODS/FORTE. He is also member of the editorial board of the *Annales des Télécommunications* and the steering committee of MSR series of conferences. C. Jard supervises a CNRS national transverse program on formal approaches for embedded systems (AFSEC). C. Jard is the director of the research of its Brittany extension (director of the pluridisciplinary institute called the Hubert Curien Research College). He is member of the scientific council of the European University of Brittany. He is expert of the AERES, the national evaluation agency and expert for the French ministry of research, he has also served as an expert in several programs of the ANR. In 2009, C. Jard was reviewer of the PhD thesis of T. Q. Tran (University of Bordeaux) and L.-M. Traonouez (University of Nantes) and member of the PhD Committees of A. Degorre (University of Grenoble), N. Sznajder (ENS Cachan) and S. Rosario (University of Rennes).

Loïc Hérouët is co-reporter at ITU for the question 17 on MSC language. Loïc is also the co-organizer (together with S. Pinchinat (S4), D. Cachera (Celtique) and N. Bertrand (Vertecs)) of the 68NQRT, a weekly seminar of IRISA on software, theory of computing, discrete mathematics in relation to computer science and artificial intelligence. He is the coordinator for the DST associated team between Rennes, the National University of Singapore, and two computer science institutes in Chennai. He is a member of GROLO, a working group on the evaluation of softwares at INRIA. He is also a member of the working group for international relations in the scientific orientation council of INRIA.

7.3. Teaching

L. Hérouët gave a lecture on non-interference in February, during an winter school on security organized at IRISA.

E. Fabre teaches "information theory and coding" at École Normale Supérieure de Cachan, Ker Lann campus, in the computer science and telecommunications Master program. He also teaches "numerical and combinatorial optimization," and "distributed algorithms and systems" in the computer science Master program at the University of Rennes 1.

C. Jard is a full-time professor at the ENS Cachan and teaches mainly at the Master level, in Computer Science and Telecom, and in Maths. He supervises the third year of the cursus (the research master's degree). He is also in charge of the competitive examination for the entry of new students in computer science in the French ENS schools.

A. Bouillard is an Assistant Professor at the ENS Cachan and teaches at the last year of Bachelor and at Master level in computer science. She is also the responsible for the computer science option of the Agrégation of Mathematics (highest competitive examination for teachers in France), where she is involved in the training of the candidates.

7.4. Visits and invitations

C. Jard gave an invited talk on symbolic concurrent semantics of safe Petri nets in the Dagstuhl Seminar on Design and Validation of Concurrent Systems, August 30, 2009, Germany.

B. Genest spent two weeks in February 2009 in Singapore (NUS) to work with P.S. Thiagarajan.

A. Benveniste, C. Jard and S. Rosario spent one week in February 2009 in Austin to work with J. Misra and W. Cook about some QoS aspects in ORC and its partial order instrumented semantics.

L. Jézéquel spent the first two weeks of November at the NICTA (National Information and Communication Technologies institute of Australia, Canberra, Australia) to perform experiments on distributed planning with P. Haslum and S. Thiebaut.

8. Bibliography

Major publications by the team in recent years

- [1] S. ABBES, A. BENVENISTE. *True-concurrency probabilistic models: Markov nets and a law of large numbers*, in "Theoretical Computer Science", vol. 390, n^o 2-3, 2008, p. 129-170.
- [2] A. BENVENISTE, E. FABRE, C. JARD, S. HAAR. *Diagnosis of asynchronous discrete event systems, a net unfolding approach*, in "IEEE Transactions on Automatic Control", vol. 48, n^o 5, May 2003, p. 714–727.
- [3] A. BOUILLARD, E. THIERRY. *An Algorithmic Toolbox for Network Calculus*, in "Journal of Discrete Event Dynamic Systems (JDEDS)", vol. 18, n^o 1, 2008, p. 3-49.
- [4] T. CHATAIN, C. JARD. *Complete finite prefixes of symbolic unfoldings of safe time Petri nets*, in "Proc. of ICATPN", LNCS, n^o 4024, June 2006, p. 125–145.
- [5] E. FABRE, A. BENVENISTE. *Partial Order Techniques for Distributed Discrete Event Systems: why you can't avoid using them*, in "Journal of Discrete Event Dynamic Systems (JDEDS)", vol. 17, 2007, p. 355-403.
- [6] E. FABRE. *Trellis Processes: a Compact Representation for Runs of Concurrent Systems*, in "Journal of Discrete Event Dynamic Systems (JDEDS)", vol. 17, n^o 3, 2007, p. 267–306.
- [7] T. GAZAGNAIRE, B. GENEST, L. HÉLOÛET, P. THIAGARAJAN, S. YANG. *Causal Message Sequence Charts*, in "Theoretical Computer Science (TCS)", vol. 410, n^o 41, 2009, p. 4094-4110.

Year Publications

Doctoral Dissertations and Habilitation Theses

- [8] D. BISWAS. *Visibility in Hierarchical Systems*, Université de Rennes 1, 2009, Ph. D. Thesis FR .
- [9] S. ROSARIO. *Quality of Service issues in compositions of Web services*, Université de Rennes 1, 2009, Ph. D. Thesis FR .

Articles in International Peer-Reviewed Journal

- [10] D. BOSNACKI, E. ELKIND, B. GENEST, D. PELED. *On Commutativity Based Edge Lean Search*, in "Special issue of BISFAI 2007, Annals of Mathematic and Artificial Intelligence (AMAI)", vol. in press, 2009 FR SG IL NL GB .
- [11] P. DARONDEAU, B. GENEST, P.S. THIAGARAJAN, S. YANG. *Quasi-Static Scheduling of Communicating Tasks*, in "Special issue of CONCUR 2008 accepted to Information and Computation", vol. in press, 2009 FR SG MO .
- [12] E. ELKIND, B. GENEST, D. PELED, P. SPOLETINI. *Quantifying the Discord: Order Discrepancies in Message Sequence Charts*, in "Special Issue of ATVA'07, International Journal of Foundations of Computer Science (IJFCS)", vol. in press, 2009 FR SG IL IT BG .

- [13] T. GAZAGNAIRE, B. GENEST, L. HÉLOUËT, P.S. THIAGARAJAN, S. YANG. *Causal Message Sequence Charts*, in "Theoretical Computer Science (TCS)", vol. 410, n^o 41, 2009, p. 4094-4110 FR SG MO GB .
- [14] A. MADALINSKI, E. FABRE. *Modular Construction of Finite and Complete Prefixes of Petri Net Unfoldings*, in "Fundamenta Informaticae", n^o 95, 2009 FR CL .
- [15] S. ROSARIO, A. BENVENISTE, C. JARD. *Flexible Probabilistic QoS Management of Orchestrations*, in "International Journal of Web Services Research", n^o 2, 2010, to appear FR .

Invited Conferences

- [16] A. BOUILLARD, B. COTTENCEAU, B. GAUJAL, L. HARDOUIN, S. LAGRANGE, M. LHOMMEAU, E. THIERRY. *COINC Library: a toolbox for the Network Calculus*, in "Fourth International Conference on Performance Evaluation Methodologies and Tools (ValueTools)", ICST, 2009 FR .

International Peer-Reviewed Conference/Proceedings

- [17] S. ABBES, A. BENVENISTE. *Concurrency, Sigma-Algebras, and Probabilistic Fairness*, in "FOSSACS", 2009, p. 380-394 FR .
- [18] N. BERTRAND, B. GENEST, H. GIMBERT. *Qualitative Determinacy and Decidability of Stochastic Games with Signals*, in "Twenty-Fourth Annual IEEE Symposium on Logic In Computer Science (LICS)", vol. IEEE 978-0-7695-3746-7, 2009, p. 319-328 FR .
- [19] D. BISWAS, A. JIWANE, B. GENEST. *Atomicity for XML Databases*, in "Sixth International XML Database Symposium (XSym)", vol. LNCS 5679, 2009, p. 160-167 DE IN FR .
- [20] A. BOUILLARD, S. HAAR, S. ROSARIO. *Critical Paths in the Partial Order Unfolding of a Stochastic Petri Net*, in "FORMATS", J. OUAKNINE, F. W. VAANDRAGER (editors), Lecture Notes in Computer Science, vol. 5813, Springer, 2009 FR .
- [21] A. BOUILLARD, L. T.X. PHAN, S. CHAKRABORTY. *Lightweight Modeling of Complex State Dependencies in Stream-Processing Systems*, in "15th IEEE Real-Time and Embedded Technology and Applications Symposium (RTAS)", 2009 FR DE .
- [22] A. BOUILLARD, S. ROSARIO, A. BENVENISTE, S. HAAR. *Monotonicity in Service Orchestrations*, in "Petri Nets", Lecture Notes in Computer Science, vol. 5606, Springer, 2009, p. 263-282 FR .
- [23] E. FABRE, L. JÉZÉQUEL. *Distributed Optimal Planning: an Approach by Weighted Automata Calculus*, in "Conference on Detection and Control (CDC)", 2009 FR .
- [24] B. GRABIEC, C. JARD. *Unfolding of Networks of Automata and their Application in Supervision*, in "9th annual international conference on New Technologies of Distributed Systems (NOTERE)", july 2009 FR .
- [25] F. KERSCHBAUM, D. DAHLMEIER, A. SCHRÖPFER, D. BISWAS. *On the practical importance of communication complexity for secure multi-party computation protocols*, in "Proceedings of the 2009 ACM Symposium on Applied Computing (SAC), Honolulu, Hawaii, USA, March 9-12, 2009", 2009, p. 2008-2015 FR GE .

[26] S. ROSARIO, A. BENVENISTE, C. JARD. *Monitoring Probabilistic SLAs in Web Service Orchestrations*, in "IFIP/IEEE Intern. Symposium on Integrated Network Management, Mini-conference", IEEE, june 2009 FR .

[27] S. ROSARIO, A. BENVENISTE, C. JARD. *Probabilistic QoS Management of Transaction based Web services Orchestrations*, in "IEEE 7th International Conference on Web Services (ICWS 2009)", IEEE, july 2009 FR .

National Peer-Reviewed Conference/Proceedings

[28] M. BOYER, L. JOUHET, A. BOUILLARD. *Notations pour le calcul réseau*, in "Modélisation des systèmes réactifs (MSR)", JESA, 2009 FR .

Books or Proceedings Editing

[29] C. JARD, O.-H. ROUX (editors). *Communicating Embedded Systems – Software and Design*, ISTE Publishing / John Wiley, October 2009, ISBN:978-1-8482-1143-8 FR .

Research Reports

[30] A. BOUILLARD, S. HAAR, S. ROSARIO. *Critical paths in the Partial Order Unfolding of a Stochastic Petri Net*, n^o 7003, INRIA, 2009, Technical report FR .

[31] A. BOUILLARD, L. JOUHET, E. THIERRY. *Tight performance bounds in the worst-case analysis of feed-forward networks*, n^o 7012, INRIA, 2009, Technical report FR .

[32] L. HÉLOUËT, A. BENVENISTE. *Distributed Active XML and Service Interfaces*, n^o 7082, INRIA, 2009, Research report FR .

[33] L. HÉLOUËT, A. ROUMY. *Canaux cachés et théorie de l'information*, Projet ANR DOTS, 2009, Research report, DOTS deliverable 2.4.b FR .

[34] V. REHÁK, P. SLOVÁK, J. STREJČEK, L. HÉLOUËT. *Decidable Race Condition for HMSC*, n^o FIMU-RS-2009-10, Faculty of Informatics, Masaryk Univ., 2009, Technical report FR PO .

References in notes

[35] S. ABITEBOUL, Z. ABRAMS, S. HAAR, T. MILO. *Diagnosis of asynchronous discrete event systems: datalog to the rescue!*, in "PODS '05: Proceedings of the twenty-fourth ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems, New York, NY, USA", ACM Press, 2005, p. 358–367, <http://doi.acm.org/10.1145/1065167.1065214>.

[36] A. BOUILLARD, B. GAUJAL, S. LAGRANGE, E. THIERRY. *Optimal routing for end-to-end guarantees using Network Calculus*, in "Performance Evaluation", vol. 65, n^o 11-12, 2008, p. 883-906.

[37] A. BOUILLARD, L. JOUHET, E. THIERRY. *Computation of a (min,+) multi-dimensional convolution for end-to-end performance analysis*, in "3rd international conference on Performance Evaluation Methodologies and Tools (ValueTools 2008), Athens, Greece", October 2008.

[38] A. BOUILLARD, E. THIERRY. *An Algorithmic Toolbox for Network Calculus*, in "Journal of Discrete Event Dynamic Systems (JDEDS)", vol. 18, n^o 1, 2008, p. 3-49.

-
- [39] M. G. BUSCEMI, V. SASSONE. *High-Level Petri Nets as Type Theories in the Join Calculus*, in "FoSSaCS", Lecture Notes in Computer Science, vol. 2030, Springer, 2001, p. 104-120.
- [40] C. CASSANDRAS, S. LAFORTUNE. *Introduction to discrete event systems*, Kluwer Academic Publishers, 1999.
- [41] R. DEVILLERS, H. KLAUDEL. *Solving Petri net recursions through finite representation*, in "IASTED International Conference on Advances in Computer Science and Technology, ACST'2004", ACTA Press, 2004, p. 145-150, ISBN 0-88986-497-3.
- [42] L. HÉLOUËT, T. GAZAGNAIRE, B. GENEST. *Diagnosis from Scenarios*, in "WODES'06", 2006, p. 307-312.
- [43] L. HÉLOUËT, T. GAZAGNAIRE, H. MARCHAND. *Anomaly Detection with diagnosis technique*, Projet Politess, 2008, Technical report.
- [44] ITU-TS. *ITU-TS Recommendation Z.120: Message Sequence Chart (MSC)*, ITU-TS, Geneva, September 1999.
- [45] A. MADALINSKI, E. FABRE. *Modular construction of finite and complete prefixes of Petri net unfoldings*, in "ACSD'08, Applications of Concurrency to System Design, Xian, China", July 2008.
- [46] J. MISRA. *A Programming Model for the Orchestration of Web Services*, in "SEFM", IEEE Computer Society, 2004, p. 2-11.
- [47] M. NIELSEN, G. PLOTKIN, G. WINSKEL. *Petri nets, event structures and domains, part 1*, in "Theoretical Computer Science", 1981.
- [48] OBJECT MANAGEMENT GROUP. *Unified Modeling Language Specification version 2.0: Superstructure*, n^o pct/03-08-02, OMG, 2003, Technical report.
- [49] W. REISIG. *Petri nets*, Springer Verlag, 1985.
- [50] M. RENIERS, S. MAUW. *High-level Message Sequence Charts*, in "SDL97: Time for Testing - SDL, MSC and Trends, Evry, France", A. CAVALLI, A. SARMA (editors), Proceedings of the Eighth SDL Forum, September 1997, p. 291-306.
- [51] M. RENIERS. *Message Sequence Charts: Syntax and Semantics*, Eindhoven University of Technology, 1998, Ph. D. Thesis.
- [52] S. ROSARIO, A. BENVENISTE, S. HAAR, C. JARD. *Probabilistic QoS and soft contracts for transaction based Web services orchestrations*, in "IEEE Trans. on Service Computing", vol. 1, n^o 14, Oct-Dec 2008, p. 1-14.
- [53] E. RUDOLPH, P. GRAUBMAN, J. GRABOWSKI. *Tutorial On Message Sequence Charts*, in "Computer Networks and ISDN Systems", vol. 28, 1996, p. 1629-1641.
- [54] G. WINSKEL. *Event structures semantics in CCS and related languages*, in "Lecture Notes in Computer Science", vol. 140, Springer Verlag, 1982.

- [55] G. WINSKEL. *Categories of Models for Concurrency*, in "Seminar on Concurrency", Lecture Notes in Computer Science, vol. 197, Springer, 1985, p. 246-267.