INRIA

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 COM

Table of contents

1. Team

Research Scientist

Albert Benveniste [DR INRIA, part–time, HdR]
Éric Fabre [CR INRIA , HdR]
Stefan Haar [CR INRIA, on sabbatical at University of Ottawa, Moved to Saclay On september 1st 2008, HdR]
Loïc Hélouët [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]

PhD Student

Thomas Gazagnaire [ENS Cachan grant, until May 1st]
Debmalya Biswas [INRIA grant until december 1st]
Sidney Rosario [INRIA grant]
Bartosz Grabiec [INRIA grant]

Post-Doctoral Fellow

Shaofa Yang [Until July 1st]

Administrative Assistant

Laurence Dinh

2. Overall Objectives

2.1. Objectives of the team

Keywords: *Orchestrations*, *Quality of Service*, *Web services*, *distributed algorithms*, *distributed testing*, *fault management*, *self-management*, *telecommunications*.

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

- The Joint ALU-Bell Labs / INRIA laboratory has been launched this year and our team is involved in a Joint Research Action called *High Manageability.*
- Our participation to ITU regarding the standardisation of scenario formalisms (Message Sequence Charts and their variants) has achieved the publication of a new appendix on applications of MSCs, that will be distributed with standard Z.120.
- Albert Benveniste has obtained the *2008 Grand Prix France-Telecom* of the french Academy of Sciences.

3. Scientific Foundations

3.1. Overview of the needed paradigms

Keywords: *asynchronous system*, *distributed system*, *monitoring*, *quality of service*, *testing*.

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. Coutermeasures would consist in reconfigurating 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, graph grammars, and their variants

Keywords: *Models of concurrency*, *event structures*, *graph grammars*, *nets*, *scenarios*.

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 [50], [53]. This allowed in particular to formally define infinite scenarios, and to enhance them with variables, guards, etc [58], [55], [56]. 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 [52], [59]. 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 [60]. 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 [52], [59]. 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* [54], [46] and *languages of scenarios* such as High-level Message Sequence Charts (HMSC) [50], [56]. 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: graph grammars.

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 [45] and *net systems* [47]. For the moment, such models lack a suitable theory of unfoldings. A relevant alternative is the class of *graph grammars* [43], [44]. Graph grammars transform graphs by means of a finite set of rules.

4. Software

4.1. SOFAT : a scenario toolbox

Keywords: *partial orders*, *scenarios*.

Participants: Loïc Hélouët, Deepak Bhatia.

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.

This year, a complete redevelopment and improvement has been performed. The graphical interface has been completely redevelopped, existing functionnalities enhanced and corrected, and new functionalities have been integrated to the tool (e.g. computation of existential bounds in MSCs). A new distribution (SOFAT V3) has been released. This was the occasion for the redaction of a complete user manual, and of a licence. This software has been presented at the International Telecommunication Union during a standardization meeting in september. This demonstration was done as an illustration of the importance of detection of MSC subclasses, after the presentation of the new appendix to Z.120 that deals with this problem [39]. We also have established contacts with a tool vendor called PRAGMADEV, and we have connected SOFAT to one of their tool that allows tracing of MSCs during a simulation. We hope that this first approach will bring opportunities of industrial transfer.

5. New Results

5.1. Fundamentals results and algorithms: distributed unfoldings

Keywords: *category theory*, *concurrent systems*, *distributed algorithms*, *event structures*, *probabilistic event structures*, *unfoldings*.

Participant: Eric Fabre.

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 [33]. 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 factored 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 factored 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.

5.2. Fundamentals and algorithms: communication with messages and Scenarios

Keywords: *concurrent systems*, *distributed algorithms*, *probabilistic event structures*, *scenarios*.

Participants: Loïc Hélouët, Blaise Genest, Thomas Gazagnaire.

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 class of models* allowing users to *specify easily complex distributed systems* while *preserving the decidability* of some key problems, 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 developped in the context of Thomas Gazagnaire's thesis [10] a new model of scenarios, namely Causal HMSCs, in order to specify complex telecomunication protocols, as sliding windows protocols. The main novelty of the approach is to allow for an independance 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 build with different independence alphabets), the problems turn out 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 a product of HMSCs with synchronizations [28] and as communicating automata [29]. We proved that this problem is undecidable for these two classes of systems in general. However, we give algorithms to solve the problem in non trivial subclasses. In particular, the subclass of products of HMSCs that we exhibited in [28] (synchronizing is allowed on a single process) is compatible with Causal HMSCs, which means that we can hope to extend the result to Causal HMSCs, using the same techniques.

A new paradigm was designed this year in [18] in order to perform diagnosis of scenarios with uncertain models. In order to specify what we know about the model, the specification is given in the form of a partial order logical formula, and the scenarios are extanded to take into account more behaviors. That is, instead of giving all the possible runs of the system, there is only a conjunction of known facts about the system. Hence, instead of giving (a priori infinite) sets of behavior explaining some observation, precise questions can be asked about what happened and is not directly observable through the observation (hidden causes). Unfortunately, if there is no bound on the message queue, the diagnosis question is undecidable [18]. Hence, we assume that a bound on this message queue is given a priori on the system (for instance, we could use the algorithm in [29] to provide such a bound).

At last, we considered hierarchical models in this context of scenario, but also in other contexts [24]). Hierarchical models provide an elegant mechanism to analyze system functionality at different level of abstraction, and allow for conciseness through the reuse of components. Decidability is easily ensured since hierarchical systems can always be unnested to obtain flat systems, but at an exponential cost. Nevertheless, we prove that this unnesting is both not necessary and negative complexitywise. Indeed, we design in [24] componentwise algorithms, that lead to a modular and exponentially faster time complexity.

5.3. Fundamentals and algorithms: timed models

Keywords: *network calculus* , *unfoldings*.

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 the network calculus. The goals are supervision with time and performance evaluation.

Networks of timed automata have not been studied as such until now. The usual approach is to compute a product, which gives an equivalent single timed automaton. This forces to compute all the interleavings of actions (too space consuming for large models) and destroy the causality and concurrency information, especially interesting for the supervision task. This is why, based on our past experience in unfolding time Petri nets (the thesis of T. Chatain, two years ago), we are designing algorithms to compute finite complete prefixes of unfoldings of network automata. We mix graphical techniques inherited from contextual Petri nets and symbolic techniques to handle time constraints on the firing dates of the transitions. This is a collaborative work in the context of the DOTS project, with the IRCCyN and LSV labs. We lead the implementation works.

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.

Network Calculus is a quite recent theory developed to compute deterministic worst-case bounds in queuing networks. Computing such bound 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 and models constraints on arrival processes and on the guaranteed with arrival and service curves. Our work has been 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 given efficient algorithms to compute them [15]. 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 ([14]). 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([27]).

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). A paper about that has been submitted to RTAS'09.

5.4. Fundamentals and algorithms: probabilistic models

Keywords: *concurrent systems*, *probabilistic event structures*, *unfoldings*.

Participant: Albert Benveniste.

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. The work of this year has consisted in finalizing papers and getting one new result. We review our progresses and refer the reader to the 2004-2006 activity reports for the motivations and previous results. Regarding ongoing publications, our long paper on Markov nets, their renewal theory and associated Law of Large Numbers has been accepted for publication in the special issue of *Theoretical Computer Science* devoted to selected papers of FOSSACS'05 [12].

Our understanding has significantly improved this year. Firstly, 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. A paper has been submitted to a conference.

Regarding true concurency *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.

We now understand that the above two types of models—Markov nets and synchronous products of Markov chains—deeply differ.

5.5. Specific studies: Web services orchestrations

Keywords: *Monitoring*, *Orchestrations*, *P2P Systems*, *QoS*, *SLA*, *Web Services*.

Participants: Sidney Rosario, Albert Benveniste, Claude Jard, Stefan Haar, 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 [51]. In 2007, jointly with the inventors of ORC, we developed a semantics of ORC into partial order models (Asymmetric Event Structures). Based on this work, Claude Jard developed a small tool written in Prolog which implements the original sequential semantics of ORC, as well as its partial order semantics. Most studies of WS service compositions concentrate on the functional aspects.

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.

This year, we have further progressed in two directions. 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 [37] 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 of the ICWS07 special issue of *IEEE Transactions on Services Computing* that is currently under review, cf. our ICWS07 publication [57]. This monitoring service has been added to the *TOrQuE* (***T**ool for **Or**chestration simulation and **Qu**ality of service **E**valuation*) tool developed since 2007 by Sidney Rosario regarding soft probabilistic QoS contracts for Web services orchestrations.

5.6. Specific studies: Asynchronous Testing

Keywords: *Distributed systems*, *Partial order I/O automata*, *conformance testing*.

Participants: Claude Jard, Stefan Haar.

Cooperation on asynchronous conformance testing has started during Stefan Haar's sabbatical leave at the School of Information Technology and Engineering (SITE), University of Ottawa, Canada, in 2007, with Claude Jard of DistribCom and G.V. Jourdan and G. Bochmann of SITE. The conformance testing problem can be resumed as follows: Suppose given a specification model $\mathcal{M}$ and an implementation $\mathcal{M}'$ that supposedly conforms to $\mathcal{M}$; both are assumed to have input and output channels, so that the behaviour is partially controled by the input streams received, and an output stream allows to derive partial information about the behaviour performed. The relations of input to state changes and to output are supposed entirely known - via knowledge of $\mathcal{M}$ - concerning the specification; the model $\mathcal{M}'$ is unknown.

Our work strives to transcend the limitations imposed by the synchronization constraints brought about by the automaton model; we have been pushing the synchronization restrictions back into that direction:

- In [48], we use a model of I/O automata with *bipartite* partial orders, composed of one layer of concurrent inputs and one of concurrent outputs that are related to some of the inputs by precedence relations. We obtain checking sequences for conformity of such models that are structurally similar to those obtained in the classical case. Concerning the complexity of our checking approach, one might have expected that the new model were just a concise way of specifying the same behavior, and thus that testing would be the same in both cases (that is, that all combinations of concurrent inputs would have to be tested anyways). It turns out not to be the case; in fact, the number of tests required for concurrent input in our model is polynomial in the number of inputs.
- In [26], we have presented this year an extension to a more general model in which the partial order on a transition need not be bipartite as above ; the testing procedure for identifying individual transitions is an adaptation of the above. In [26], we adopt a more relaxed set of conformance requirements allowing to develop dedicated tests for specific types of deviation from the specification.

In future work, we plan to remove the remaining limits to asynchronicity that still prevail in the automaton models, by developping adequate Petri net models with I/O capabilities, and testing procedures adapted to their asynchronous semantics.

5.7. Specific studies: Active XML documents

Keywords: *Active XML*, *P2P Systems*, *XML.*

Participants: Stefan Haar, Albert Benveniste, Loïc Hélouët, 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 frameword 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 [2] (where nothing can get deleted). This problem is difficult due to undboundedness 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 [30]. However, we do not know yet how to handle some key properties (as confluence or distribution), or key data (as the names in the infinite alphabet). It will be the main activity of the AXML project next year, by proposing an interface based semantics to AXML, in order to model the distribution, and by extending the result to classes that do not impose finite alphabets.

Another main activity of the AXML project is to propose algorithms to bring transactional features to AXML services [34], and more generaly 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 will defend in January 2009. The transactional feature considered was replication and recovery of peers [22], and more importantly atomicity, which is ensured through a compensating (rollback) mechanism. For this, actions performed are logged. However, in order to ensure security and performance, all actions are not logged. On the other hand, computing an (absolute) minimal set of action to log allowing to perform compensation is untractable in general. Hence, we proposed to work on hierachical systems, designing a modular algorithm using the hierarchy to obtain good performance [24]. Ususally, obtaining the absolute minimal set of actions to log is not necessary, and using a slightly bigger set is good enough. Hence, we also designed algorithms over approximating this set [23], bringing a trade off between speed performance and size of the log. Concerning security, a visibility model has been proposed in [25], allowing to talk about what should be secret and what should be shared in a P2P networks, while allowing for fast updating algorithm in this highly dynamic network. Performance has also been considered with respect to the number of messages being sent (this is a usual bottleneck in communicating protocols) in [32].

5.8. Specific studies: security

Keywords: *concurrent systems*, *covert channels*, *diagnosis*, *distributed games*, *security*, *test*.

Participant: Loïc Hélouët.

We have performed some work on security issues in the context of the DOTS project, and in the context of the RNRT POLITESS project with the VERTECS team. In DOTS, we are involved in a working group on non-interference. The POLITESS RNRT project, aims at providing techniques to test security issues, such as the presence of covert channels in software, the portential violation of a security policy, the presence of interference, etc. This year, we have mainly focused on two topics, anomaly detection, and covert channels testing.

Anomaly detection with diagnosis techniques: We have proposed an anomaly detection technique based on partial order diagnosis techniques. 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 [49], and has led to an internal report of the project [38], that should be submitted to a conference or a journal next year.

Testing covert channels: the objective of this work is to detect in a specification the presence of covert information flows from a description of a system given by means of scenarios. When such flow may exist, we transpose the use of this channel in the system to a distributed game problem. If a strategy exists for the derived game, then the corrupted users may have a mean to transfer covert information if the real system implements the mechanisms used by the covert flow. We then derive distributed tests to check whether a potential covert channel detected on a model exists in an implementation of a system. A test is succesful if the implementation reacts as described in the strategy during a large enough test campain. Note that it is more difficult to implement a distributed strategy than to test it, as a centralized tester has more inforamtion on the state of the system than each corrupted user of a covert flow which can only observe their side of the system. Hence, when an implementation does not pass the covert channel test, then we can consider that it does not implement a covert flow. This work has led to the production of an internal report of the POLITESS project [41], and should be submitted to a conference next year.

6. Other Grants and Activities

6.1. ANR Docflow

Participants: Albert Benveniste, Eric Fabre, Loïc Hélouët, Blaise Genest, Hélia Pouyllau, Debmalya Biswas.

Contract INRIA *ANR-06-MDCA-005 January 2007 - December 2009*

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 December 2009. 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 worflow, and structured data (XML) can be used in the same formalism. This starting point will allow us to develop algorithm to analyse, monitor and optimize worflows with rich data.

6.2. CREATE ActiveDoc

Participants: Albert Benveniste, Eric Fabre, Loïc Hélouët, Blaise Genest, Hélia Pouyllau, Il-Gon Kim, Debmalya Biswas, Sidney Rosario.

Contract INRIA *CREATE February 2007 - August 2011*

Activedoc is funded by Region 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.

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. RNRT Politess

Participant: Loïc Hélouët.

RNRT POLITESS : POLItiques de sécurité pour des systèmes dinformation en réseau : modélisation, déploiement, TESt et Surveillance.

— september 2005 - september 2008

The purpose of the Politess RNRT project (http://www.rnrt-politess.info/) is to study metholodologies that ensure conformance between high-level security policies, their deployment, their test and monitoring. The partners involved in this RNRT are : GET (Teams in Rennes and Évry), INPG/IMAG - GrenobleCentre de recherche INRIA Rennes - Bretagne Atlantique (Vertecs and Distribcom teams) France Télécom - R&D Caen, Leirios Technologies, SAP Research, Silicomp-AQL, LSR/IMAG - INPG (Vasco team), VERIMAG (DCS team).

6.5. Associated team CASDS

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

CASDS : Control, Analysis and Synthesis of Distributed Systems

Associated Team INRIA-NUS — 2006-2008

This associated team is a collaboration with the National University of Singapore, and also involves members of the S4 team. The main research theme is the control and diagnosis of distributed communicating systems. 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:

- An extension of scenario models for distributed systems diagnosis.
- Distributed control synthesis, with applications to the quasi-static scheduling problem.

A joint paper on scenarios has been published in a top conference last year, and a long version has been submitted to a journal. The works on quasi static scheduling has been accepted in CONCUR 2008 [29].The associated team has been extended in 2008, both geographically (to Chennai, India, team of Madhavan Mukind) and in subject (process calculi). Anne Bouillard has spent two weeks in NUS in july, and Blaise Genest has spent one week in NUS. Visitors from Chennai came in may to start new projects.

6.6. PHC FAST

Participant: Eric Fabre.

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 E. Fabre at NICTA (April) was partly funded by this program. It was followed by a return visit of Sylvie Thiebaux in November.

6.7. DISC

Participant: Eric Fabre.

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.

6.8. ADR HiMa

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

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 two INRIA teams, DistribCom and Madynes (O. Festor, INRIA Lorraine), and 5 persons of the PTI group (Packet Transport Infrastructure) on the Alcatel-Lucent side. It is jointly headed by M. Vigoureux (ALU) and E. Fabre (INR). 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 2008, where the team was created, the activity focused on distributed algorithms for the optimal power allocation to wavelengths in transparent (=optically routed) networks, and to the problem of organization tasks for network maintenance with minimal service disruption.

7. Dissemination

7.1. 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. He is in charge of assisting François Baccelli in the management of the INRIA side of the Alcatel external Research Programme (ARP).[1]

E. Fabre is associated editor (AE) for the journal *IEEE Trans. on Automatic Control*. He participated to the selection commission of ENS Cachan, for the hiring of new associate professors. He served as an evaluator for the ANR, the national funding agency of the French ministry of research. He has been member of the Program Committee of DX'08 (Int. Workshop on the Principles of Diagnosis).

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

B. Genest is an elected member of the Comité National de la Rercherche Scientifique for 2008-2012; He also served as Program Committee Member for the ICE'08 (Interaction and Concurrency Experience) International Workshop, Satellite event of Icalp'08, Reykjavik, Island, July 2008.

C. Jard has been in 2008 member of the Program Committee of the following international conferences: ICAPTN, ROGICS, FORTE, NOTERE, MOVEP, FORMATS, A-MOST. 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 a member of the board of directors of the ENS Cachan and 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 2008, C. Jard was member of the PhD Committees of T. Gazagnaire and T. Turpin (University of Rennes 1), and reviewer of the habilitation of S. Graf (University of Grenoble).

Stefan Haar is a member of the working group for evaluation of international activities with the COST committee of INRIA; upon his transfer to INRIA Saclay, he has been appointed correspondent for international activities of the Saclay Center. His mandate as associate editor of *IEEE Transactions on Automatic Control* has been extended to the end of 2009. Stefan supervised the thesis of Helia Pouyllau, successfully defended in December 2007.

Loïc Hélouët is co-rapporteur at ITU for the question 17 on MSC language. He has proposed a document [39] that summarizes all undecidability results in the application of Message Sequence Charts published this last decade. This proposal was accepted during the plennary session of the standardization meeting of September 2008, and is now an appendix of standard Z.120. Loïc was member of the PhD jury of Cedric Meuters at the Université Libre de Bruxelles. Loïc also animates (with S. Pinchinat (S4), Th. Genet (Lande) and N. Bertrand (Vertecs)) the 68NQRT, a weekly seminar of IRISA on software, theory of computing, discrete mathematics in relation to computer science and artificial intelligence. He was the coordinator for the CASDS associated team between Rennes, the National University of Singapore, and two computer science institutes in Chennai. At the occasion of the visit of indian partners at IRISA in may, he has organised a one day seminar. He is a member of GROLO, a working group on the evaluation of softwares at INRIA.

7.2. Teaching

L .Helouët 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 at Bachelor and 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.

7.3. Visits and invitations

A. Benveniste gave an invited talk entitled "Composing Web Services in an open world: QoS issues" in the 5th International Conference on Quantitative Evaluation of SysTems (QEST'08) in September 2008.

C. Jard gave an invited talk on runs of distributed systems and partial orders in the 1st International Conference on Relations, Orders and Graphs (ROGICS08) in May 2008.

A. Bouillard spent two weeks in July 2008 in Singapore (NUS) to work with P.S Thagarajan, S. Chakraborty and Phan Thi Xuan Linh, about some generalization of the Network Calculus framework.

B. Genest spent one week in September 2008 in Singapore (NUS) to work with P.S. Thiagarajan and Shaofa Yang.

S. Rosario spent two weeks in August and September 2008 in Austin to work with J. Misra and W. Cook about some QoS aspects in ORC.

É. Fabre spent the month of April 2008 at the NICTA (National Information and Communication Technologies institute of Australia, Canberra, Australia) to work with S. Thiebaux, P. Haslum and J. Rintanen on algorithms for distributed planning.

M. Mukund (CMI), N. Kumar (CMI) and R. Ramanujam (IMSc), the Indian partners of our associated team CASDS came from Chennai, India and visited our project in May 2008. A mini-seminar was organised and some works have been initiated.

V. Rehak and J. Sterjchek from Masaryk university (Brno) came for a one week visit in Rennes in November 2008. A new collaboration on Message Sequence Charts and race problems was initiated.

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] 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.

[3] 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.

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

[5] 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.

[6] 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.

[7] 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.

[8] T. GAZAGNAIRE, B. GENEST, L. HÉLOUËT, P. THIAGARAJAN, S. YANG. *Causal Message Sequence Charts*, in "Proceedings of CONCUR 2007", 2007.

[9] B. GENEST, A. MUSCHOLL, D. PELED. *Message Sequence Charts.*, in "Lectures on Concurrency and Petri Nets", vol. LNCS 3098, 2003, p. 537-558.

Year Publications

Doctoral Dissertations and Habilitation Theses

[10] T. GAZAGNAIRE. *Scenario Languages: Using Partial Orders to Model, Verify and Supervize Distributed and Concurrent Systems*, Ph. D. Thesis, Université de Rennes 1, 2008.

[11] S. HAAR. *Law and Partial Order; Nonsequential Behaviour and Probability in Asynchronous Systems*, Ph. D. Thesis, Université de Rennes 1, 2008.

Articles in International Peer-Reviewed Journal

[12] 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.

[13] A. BOUILLARD, B. GAUJAL. *Backward Coupling in Bounded Free-Choice Nets Under Markovian and Non-Markovian Assumptions*, in "Journal of Discrete Event Dynamic Systems (JDEDS)", vol. 18, n^{o} 4, 2008, p. 473-498.

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

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

[16] B. GENEST, A. MUSCHOLL. *Pattern Matching and Membership for Hierarchical Message Sequence Charts*, in "Theory of Computing Systems (TOCS)", vol. 42, n^{o} 4, 2008, p. 536-567.

[17] H. POUYLLAU, S. HAAR. *Distributed Busacker-Gowen algorithm for end-to-end QoS pipe negotiation in X-domain networks*, in "Annals of Telecommunications", DOI 10.1007/s12243-008-0055-0, October 2008, -.

[18] S. YANG, L. HÉLOUËT, T. GAZAGNAIRE. *Logic-based diagnosis for distributed systems*, in "Theoretical Computer Science, Special Issue for P.S. Thiagarajan's Festchrift", to appear, 2008.

Invited Conferences

[19] A. BENVENISTE. *Composing Web Services in an Open World: Issues of Quality of Service*, in "FORMATS", 2008, 14.

[20] C. JARD. *Runs of Distributed Systems and Partial Orders*, in "ROGICS'2008, 1st International Conference on Relations, Orders and Graphs: Interaction with Computer Science, Mahdia, Tunisia", may 2008.

International Peer-Reviewed Conference/Proceedings

[21] P. BALDAN, T. CHATAIN, S. HAAR, B. KÖNIG. *Unfolding-Based Diagnosis of Systems with an Evolving Topology*, in "CONCUR 2008", 2008, p. 203-217.

[22] D. BISWAS. *Active XML Replication and Recovery*, in "2nd IEEE International Conference on Complex, Intelligent and Software Intensive Systems (CISIS-2008), Barcelona, Spain", March 2008, p. 263-269.

[23] D. BISWAS, T. GAZAGNAIRE, B. GENEST. *Small Logs for Transactional Services*, in "11th IEEE High Assurance Systems Engineering Symposium (HASE'08), Nanjing, China", To appear, December 2008.

[24] D. BISWAS, B. GENEST. *Minimal Observability for Transactional Hierarchical Services*, in "20th International Conference on Software Engineering and Knowledge Engineering (SEKE'08), San Francisco, USA", vol. ISBN 1-891706-22-5, July 2008, p. 531-536.

[25] D. BISWAS, K. VIDYASANKAR. *A Highly Flexible Data Structure for Multi-level Visibility of P2P Communities*, in "9th International Conference on Distributed Computing and Networking (ICDCN 2008), Kolkata, India", vol. LNCS 4904, January 2008, p. 363-374.

[26] G. BOCHMANN, S. HAAR, C. JARD, G.-V. JOURDAN. *Testing Systems Specified as Partial Order Input/Output Automata*, in "Testcom/Fates 08, 20th IFIP International Conference on Testing of Communicating Systems and 8th International Workshop on Formal Approaches to Testing of Software, Tokyo", vol. LNCS 5047, june 2008, p. 169-183.

[27] A. BOUILLARD, L. JOUHET, É. 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.

[28] P. DARONDEAU, B. GENEST, L. HÉLOUËT. *Products of Message Sequence Charts*, in "11th International Conference on Foundations of Software Science And Computation Structures (FOSSACS'08), Budapest, Hungary", vol. LNCS 4962, April 2008, p. 459-474.

[29] P. DARONDEAU, B. GENEST, P. THIAGARAJAN, S. YANG. *Quasi-Static Scheduling of Communicating Tasks*, in "19th International Conference on Concurrency Theory (CONCUR'08), Toronto, Canada", vol. LNCS 5201, August 2008, p. 310-324.

[30] B. GENEST, A. MUSCHOLL, O. SERRE, M. ZEITOUN. *Tree Pattern Rewrite Systems*, in "6th International Symposium on Automated Technology for Verification and Analysis (ATVA'08), Seoul, Korea", To appear, October 2008.

[31] C. JARD. *Teaching Distributed Algorithms using Spin*, in "Formal Methods in Computer Science Education, ETAPS 2008, Budapest, Hungary", march 2008, -.

[32] F. KERSCHBAUM, D. BISWAS, D. DAHLMEIER, A. SCHROPFER. *An Experimental Study on the Practical Importance of Communication Complexity for Secure Multi-Party Computation Protocols*, in "24th Annual ACM Symposium on Applied Computing (SAC 2009), Hawaii, USA", To appear, March 2009.

[33] 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.

Scientific Books (or Scientific Book chapters)

[34] D. BISWAS, I.-G. KIM. *Active XML Transactions, chapter of Service and Business Computing Solutions with XML: Applications for Quality Management and Best Processes*, IGI Global, 2008.

[35] F. CASSEZ, C. JARD. *The 6th International Conference on Formal Modelling and Analysis of Timed Systems, FORMATS 2008*, Springer-Verlag, LNCS 5215, 2008.

[36] C. JARD, O.-H. ROUX. *Systèmes embarqués communicants : Approches formelles*, Hermès, Traité IC2, 2008.

Research Reports

[37] A. BOUILLARD, S. ROSARIO, A. BENVENISTE, S. HAAR. *Monotonicity in Service Orchestrations*, INRIA Research Report, n^{o} RR-6528, INRIA, 2008, http://hal.inria.fr/inria-00276627/fr/.

[38] L. HÉLOUËT, T. GAZAGNAIRE, H. MARCHAND. *Anomaly Detection with diagnosis technique*, Technical report, Projet Politess, 2008.

[39] L. HÉLOUËT. *Appendix 1 to Z.120: Applications of Messages Sequence Charts*, Technical report, ITU-T, 2008.

[40] L. HÉLOUËT. *SOFAT V3 : Suer Manual*, August 2008.

[41] L. HÉLOUËT, H. MARCHAND, T. JÉRON. *Testing Covert Channels*, Technical report, Projet Politess, 2008.

[42] A. MADALINSKI, E. FABRE. *Modular construction of finite and complete prefixes of Petri net unfoldings*, Technical report, n^{o} 6412, IRISA, jan. 2008.

References in notes

[43] H. EHRIG, H.-J. KREOWSKI, U. MONTANARI, G. ROZENBERG (editors). *Handbook of graph grammars and computing by graph transformations: vol. 3, concurrency, parallelism, and distribution*, World Scientific, 1999.

[44] P. BALDAN, A. CORRADINI, U. MONTANARI. *Unfolding and event structure semantics for graph grammars*, in "Proc. of FOSSACS 1999", Lecture Notes in Computer Science, vol. 1578, Springer Verlag, 1999, p. 73–89.

[45] M. BUSCEMI, V. SASSONE. *High-Level Petri Nets as Type Theories in the Join Calculus*, in "FOSSACS'2001", 2001.

[46] C. CASSANDRAS, S. LAFORTUNE. *Introduction to discrete event systems*, Kluwer Academic Publishers, 1999.

[47] R. DEVILLERS, H. KLAUDEL. *Solving Petri net recursions through finite representation*, in "IASTED International Conference on Advances in Computer Science and Technology, ACST'2004", ISBN 0-88986-497-3, ACTA Press, 2004, p. 145-150.

[48] S. HAAR, C. JARD, G.-V. JOURDAN. *Testing Input/Output Partial Order Automata*, in "TestCom/FATES", 2007, p. 171-185.

[49] L. HÉLOUËT, T. GAZAGNAIRE, B. GENEST. *Diagnosis from Scenarios*, in "WODES'06", 2006, p. 307-312.

[50] ITU-TS. *ITU-TS Recommendation Z.120: Message Sequence Chart (MSC)*, ITU-TS, Geneva, September 1999.

[51] J. MISRA. *A Programming Model for the Orchestration of Web Services.*, in "SEFM", 2004, p. 2-11.

[52] M. NIELSEN, G. PLOTKIN, G. WINSKEL. *Petri nets, event structures and domains, part 1*, in "Theoretical Computer Science", 1981.

[53] OBJECT MANAGEMENT GROUP. *Unified Modeling Language Specification version 2.0: Superstructure*, Technical report, n^{o} pct/03-08-02, OMG, 2003.

[54] W. REISIG. *Petri nets*, Springer Verlag, 1985.

[55] 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.

[56] M. RENIERS. *Message Sequence Charts: Syntax and Semantics*, Ph. D. Thesis, Eindhoven University of Technology, 1998.

[57] S. ROSARIO, A. BENVENISTE, S. HAAR, C. JARD. *Probabilistic QoS and soft contracts for transaction based Web services*, in "International Conference on Web Services", 2007, p. 126-133.

[58] E. RUDOLPH, P. GRAUBMAN, J. GRABOWSKI. *Tutorial On Message Sequence Charts*, in "Computer Networks and ISDN Systems", vol. 28, 1996, p. 1629-1641.

[59] G. WINSKEL. *Event structures semantics in CCS and related languages*, in "Lecture Notes in Computer Science", vol. 140, Springer Verlag, 1982.

[60] G. WINSKEL. *Categories of models for concurrency*, in "Seminar on Concurrency, Carnegie-Mellon Univ. (July 1984),Lecture Notes in Computer Science 197", 1985, p. 246-267.