IN PARTNERSHIP WITH:
CNRS

Université Rennes 1

Ecole normale supérieure de Cachan

Activity Report 2011

Project-Team DISTRIBCOM

Distributed and Iterative Algorithms for the Management of Telecommunications Systems

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

RESEARCH CENTER
Rennes - Bretagne-Atlantique

THEME
Networks and Telecommunications

Table of contents

Project-Team DISTRIBCOM

Keywords: Quality Of Service, Distributed Algorithms, Monitoring, Optical Networks, Service Orchestration, Workflow

1. Members

Research Scientists

Albert Benveniste [Team Leader, DR INRIA, part–time, HdR]
Eric Fabre [DR INRIA , HdR]
Loïc Hélouët [CR INRIA]
Blaise Genest [CR CNRS 15% with DistribCom, 85% with UMI IPAL, Singapore]
Axel Legay [CR INRIA]

Faculty Members

Claude Jard [PROF. ENS CACHAN , HdR]
Guillaume Aucher [CHAIRE MAÎTRE DE CONF. UNIV. RENNES I – INRIA, since May 1st]
François Schwarzentruber [MAÎTRE DE CONF. ENS CACHAN, since October 1st]

PhD Students

Ajay Kattepur [INRIA grant]
Loig Jézéquel [ENS Cachan grant]
Rouwaida Abdallah [ENS Cachan grant]
Cyrille Jégourel [Univ. Rennes I grant]
Carole Hounkonnou [INRIA grant]
Aurore Junier [INRIA grant]

Post-Doctoral Fellows

Sean Sedwards
Benoît Masson [until July 31st]
Uli Fahrenberg
Paolo Ballarini [Jan.-Aug.]
Akshay Sundararaman [ANR Impro ENS Cachan, since November 1st]

Administrative Assistant

Laurence Dinh

2. Overall Objectives

2.1. Objectives of the team

The DistribCom team is jointly headed by Albert Benveniste (official head for Inria) and Claude Jard. It 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, and France-Telecom, on the topic of networks and service management.

This year, Inria, Centre of Rennes-Bretagne-Atlantique, decided that Axel Legay and his group of post-docs and PhDs would get hosted by DistribCom. It was decided that, for a first period (until august 2011), this hosting would have no effect on the topics of both Axel's group and DistribCom in its previous geometry. In august 2011, we decided to make the merge more effective, by having Axel's group joining the group seminars with full participation. It is expected that some move in the activities of both bodies may result. The activities of Axel's group are specifically reported in Sections 3.4, 6.5, and 6.6.

Guillaume Aucher joined the team in May 2011.

Also, François Schwarzentruber joined the team in September 2011, so his work is not reported here.

2.2. Highlights

The Laboratoire d'Excellence (Labex) CominLabs was selected as the only Labex in France in the area of software for the first Labex competition held in 2010 (Labex are Excellence Centers within the framework of *Investissements d'Avenir*; 100 have been selected in all disciplines). CominLabs gathers ten labs from Bretagne and Nantes in the sector of telecommunications and information systems, with an overall amount of 500 researchers ("'équivalents chercheurs") and a funding of 14M Euros for a duration of 10 years; 1.4M Euros was provided for the first year. Albert Benveniste is the Scientific Director of CominLabs.

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. 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, 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 [67], [73]. This allowed in particular to formally define infinite scenarios, and to enhance them with variables, guards, etc [77], [75], [76]. 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 [72], [78]. 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 [79]. 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 [72], [78]. 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* [74], [59] and *languages of scenarios* such as High-level Message Sequence Charts (HMSC) [67], [76]. 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 [58] and *net systems* [60]. For the moment, such models lack a suitable theory of unfoldings.

3.3. Modal logics for distributed systems

Modal logics are a family of logics that were developped originally to reason about different modalities occurring in natural language, such as for example the modality of knowledge (epistemic logic), the modalities of obligation and permission (deontic logic) and the modality of time (temporal logic). Temporal logics (CTL, LTL, μ-calculus...) are the most prominent (modal) logics used in computer science nowadays, especially in the field of verification.

3.3.1. Epistemic logic and distributed systems.

In the 1980's, epistemic logic was propounded by computer scientists such as Fagin, Halpern, Moses and Vardi to address problems in distributed systems, resulting in the TARK conference series (Theoretical Aspects of Rationality and Knowledge) and the books [63], [70]. This interest in epistemic logic was due to their observation that the notion of knowledge plays a central role in the informal reasoning used in the design of distributed protocols. This lead these authors to "hope that a theory of knowledge, communication and action will prove rich enough to provide general foundations for a unified theoretical treatment of distributed systems" [66]. The research pursued in DistribCom follows this line of thought, although we also strive to feed and confront our theoretical developments with actual problems stemming from diverse areas of application of distributed systems.

In [63], the behaviour of a distributed system is represented by a set of *runs*, each run being a possible execution of the distributed system, determined by a given protocol. Processors are called *agents* and their partial observation of the system is represented at any point in the run by indistinguishability relations between local states of different runs (the local state of a processor represents the state of this processor at a moment of time). This model was used to show for example that the specific notion of common knowledge of epistemic logic is necessary to reach agreement and to coordinate actions [66]. Dynamic Epistemic Logic (DEL) is another logical framework that can be used to represent and reason about distributed systems (connections between these two logical frameworks were made in [80]). DEL deals with the representation of global states of synchronous distributed systems. The global state of the system at a moment in time is represented directly by means of an *epistemic model*. Events occuring in this distributed system are represented by means of *event models* and their effects on the local states of agents (processors) are represented by means of a *product update*.

The contributions in this sub-module are described in Section 6.4.

3.3.2. Deontic logic and privacy in distributed systems.

We also use deontic logic in combination with epistemic logic for the formalization of privacy regulations. We intend to use this formalization to reason about privacy in the composition of web-services. The combination of these two modal logics can be used to express statements such as "it is *forbidden* for agent 1 to *know* that agent 2 sent message m" or "if agent 1 is an administrator of the system, then it is *permitted* for him to *know* information i". This provides a formal language very close to the natural language used in actual privacy regulations by law legislators. In the long run, we expect this formal language to be used at the level of interfaces of the web-service in order to:

1. check that the privacy policy declared by the web-service on its interface is indeed compliant (coherent) with respect to the privacy regulations expressed by law makers;
2. check that the web-service does enforce and apply the privacy policy it has declared on its interface.

The contributions in this sub-module are described in Section 6.9.

3.4. Statistical Model Checking

Complex systems pose two particular challenges to formal verification: (i) the non-determinism caused by concurrency and unpredictable environmental conditions and (ii) the size of the state space. Our interest is probabilistic model checking, that can verify intricate details of a system's dynamical behaviour and where non-determinism is handled by assigning probabilistic distributions to unknowns and quantifying results with a probability. Exact probabilistic model checking quantifies these probabilities to the limit of numerical precision by an exhaustive exploration of the state space, but is restricted by what can be conveniently stored in memory. Our focus is therefore statistical model checking (SMC), that avoids an explicit representation of the state space by building a statistical model of the executions of a system and giving results within confidence bounds. The key challenges of this approach are to reduce the length (simulation steps and cpu time) and number of simulation traces necessary to achieve a result with given confidence. Rare properties pose a particular problem in this respect, since they are not only difficult to observe but their probability is difficult to bound. A further goal is to make a tool where the choice of modelling language and logic are flexible.

4. Application Domains

4.1. Telecommunication network management

The management of telecommunication networks is traditionally a human performed activity that covers the five FCAPS functions: Fault management, network Configuration, Accounting, Performances and Security. This simple classification has exploded in the last decade, under the pressure of several phenomena. The first one concerns the growth in size and complexity of networks, with the emergence of new (possibly virtual)

operators, the multiplication of vendors, new core and (wireless) access technologies, the variety of terminal devices, the convergence of phone/computer/radio/TV networks, the multiplication of services over the top, the necessity to provide QoS for a wide variety of traffic demands, etc. As a consequence, the management task is reaching the limits of human operators and demands automation. It is estimated that telecommunication companies spend over 50% of their manpower on management tasks. They naturally want to reduce it and dedicate their effort to the design and offer of innovative services, where the added value is more important (as witnessed by the success of some over-the-top companies). The result of these trends is that network management now covers a much wider variety of problems, for which automatic solutions are requested. This takes the name of self-management, or autonomic management: one wishes to manage networks by high-level objectives, and networks should be able to adapt themselves automatically to fulfill these objectives.

DistribCom is contributing to this field with its background on the modelling of distributed/concurrent systems, and its expertise in distributed algorithms. Networks are perfect examples of large distributed and concurrent systems, with specific features like the dynamicity (their structure evolves) and a hierarchical structure (multiple layers, multiple description granularities). We have proposed model-based distributed algorithms to solve problems like failure diagnosis, negotiation of QoS (quality of service) parameters, parameter optimization, graceful shutdown of OSPF routers for maintenance operations... The present activities in this domain are related to the joint diagnosis for access network + core network + services, within the European IP UniverSelf. The challenges cover self-modelling methods (how to obtain the network model that is used by the management algorithms), active diagnosis methods that both adapt the scope of their network model and perform tests to explain a fault situation, and self-healing methods.

4.2. Web services and active structured documents

Keywords: Active documents, Web services, choreographies, orchestrations, QoS.

Web services architectures are usually composed of distant services, assembled in a composite framework. This raises several practical issues: one of them is how to choose services, assemble them, and coordinate their executions in a composite framework. Another issue is to guarantee good properties of a composite framework (safety but also QoS properties). All this has to be done in a context where a distant service provided by a subcontractor is only perceived as an interface, specifiying legal inputs and outputs, and possibly a quality contract.

The standard in industry for Web-services is now BPEL [57], but most of the problems listed above are untractable for this language. Composition of services can also be performed using choreography languages such as ORC [69]. The implementation of orchestration and choreography description languages raises a number of difficulties related to efficiency, clean semantics, and reproducibility of executions, issues of composite QoS associated with orchestrations. We develop studies in these areas, with the aim of proposing service composition frameworks equiped with tools to specify, but also to monitor and analyse the specified architectures.

Another issue is the convergence between data and workflows. Web Services architectures are frequently considered exclusively as workflows, or as information systems. Many approaches to Web Service orchestration and choreography abstract data away. Symmetrically, modern approaches to Web data management typically based on XML and Xqueries rely on too simplistic forms of control. We develop a line of research on Active documents. Active documents are structured data embedding references to services, which allow for the definitions of complex workflows involving data aspects. The original model was proposed by S. Abiteboul [56], but the concept of active document goes beyond AXML, and offers a document oriented alternative to Web services orchestrations and choreographies. This approach is in particular well adapted to the modeling of E-business processes, or information processing in organizations, etc. Our aim is to extend and promote the concept of active document. This means developing verification and composition tools for document-based architectures, considerd not only as theoretical models but also as effectively running systems. To this extend, we develop an active document plateform.

5. Software

5.1. SOFAT

Participants: Loïc Hélouët [correspondant], Rouwaida Abdallah.

SOFAT is the acronym for Scenario Oracle and Formal Analysis Toolbox. As this name suggests it is a formal analysis toolbox for scenarios. Scenarios are informal descriptions of behaviors of distributed systems. SOFAT allows the edition and analysis of distributed systems specifications described using Message Sequence Charts, a scenario language standardized by the ITU [Z.120]. The main functionalities proposed by SOFAT are the textual edition of Message Sequence Charts, their graphical visualization, the analysis of their formal properties, and their simulation. The analysis of the formal properties of a Message Sequence Chart specification determines if a description is regular, local choice, or globally cooperative. Satisfaction of these properties allow respectively for model-checking of logical formulae in temporal logic, implementation, or comparison of specifications. All these applications are either undecidable problems or unfeasible if the Message Sequence Chart description does not satisfy the corresponding property. The SOFAT toolbox implements most of the theoretical results obtained on Message Sequence Charts this last decade. It is regularly updated and re-distributed. The purpose of this software is twofold:

- Provide a scenario based specification tool for developers of distributed applications
- Serve as a platform for theoretical results on scenarios and partial orders

SOFAT provides several functionalities, that are: syntactical analysis of scenario descriptions, Formal analysis of scenario properties, Interactive Simulation of scenarios when possible, and diagnosis. This year, SOFAT was extended with code synthesis functionalities, allowing to generate communicating automata, promela code, or rest based web services from HMSCs. A new release of the software is expected before the end of the year.

See also the web page http://www.irisa.fr/distribcom/Prototypes/SOFAT/index.html.

- AMS: Order; lattices; ordered algebraic structures
- APP: IDDN.FR.001.080027.000.S.P.2003.00.10600
- Programming language: Java

6. New Results

6.1. Fundamental 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/marked state in a huge automaton. The latter is specified by a collection of resources, that may be available or not (which defines a state), and actions that consume and produce resources (which defines a transition). 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 an active topic in the planning community under the name of “factored planning.”

Our previous contributions to the domain was the first optimal factored planning algorithm [61]. The main idea is to represent a planning problem as a network of interacting weighted automata, the objective being to jointly drive all of them to a target state, while minimizing the cost of their joint trajectory. We have developed and tested [68] a distributed algorithm to solve this problem, based on a weighted automata calculus, and that takes the shape of a message passing procedure. Components perform local computations, exchange messages with their neighbors, in an asynchronous manner, and the procedure converges to the path that each component should follow. The optimal global plan is thus given as a tuple of (compatible) local plans, i.e. a partial order of actions.

In 2011, we have extended this framework in two directions. In terms of modelling, first. In most planning problems, some actions consume/produce resources, but also are enabled by the presence of other resources, that they only read but not consume. We have proposed to model this feature under the form of networks of automata with read arcs. Interactions then take the form of synchronous actions, as previously, but also the form of readings: a component may only be allowed to fire some local transition if another component is in a specific state. Our distributed planning approach has then been extended to this new model of distributed systems [38].

The second improvement is algorithmic. So far, our distributed optimal planning algorithm computes all possible distributed plans, in a factored form. This contrasts with the philosophy of planning algorithms, that look for one plan only, and organize the computations to quickly reach the best plan. In other words, most planning algorithms are based on a common ground known as the A-star algorithm, a depth-first search procedure in a graph, guided by a heuristic function that estimates the remaining cost to reach the goal. We have developed a truly distributed version of this algorithm, to perform a search on a product graph. Each component runs an A-star procedure to find a path to its goal, taking into account the costs of its neighbors in order to guarantee that all components converge to local plans that are compatible and jointly optimal.

6.2. Fundamental results and algorithms: communication with messages and senarios

Participants: Loïc Hélouët, Rouwaida Abdallah, Claude Jard, 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*.

This year, we have considered analysis for a timed extension of scenarios called Time-constrained MSCs and implementation techniques that take scenarios as an input model and output an equivalent distributed implementation.

The first part of our work is the study of Time-Constrained MSC graphs (TC-MSGS for short). Time-constrained MSCs (TC-MSCs) are simply MSCs decorated with constraints on the respective occurrence dates of events. The semantics of a TC-MSC T is a dated MSC, that is a MSC where events are associated with an occurrence date. For a given TC-MSC, there can be an infinite set $L(T)$ of dated MSCs satisfying its constraints. Note however that some time-constraints in a TC-MSC may not be satisfiable, and hence $L(T)$ can simply be empty. TC-MSCs can be extended by composition mechanisms such as TC-MSC graphs. TC-MSC graphs are simply automata labeled by TC-MSC. Each path ρ of a TC-MSC G is associated with a TC-MSC T_ρ obtained by concatenation of TC-MSC along ρ. The language $L(G) = \bigcup_{\rho \text{ path of } G} L(T_\rho)$ of a TC-MSC Graph is then the union of all dated MSCs associated to paths of G. Because of inconsistent timing constraints, some path may have no possible realization (i.e $L(T_\rho = \varnothing)$). One can even design a MSC Graph G such that $L(G) = \varnothing$ - such TC-MSC graph is clearly inconsistent-. It has been shown [64] that checking whether $L(G) = \varnothing$ is an undecidable problem in general, but can be decided for the restricted

subclass of regular TC-MSC graphs (that have the expressive power of event-count timed automata). We have proposed two restrictions allowing for the decision of emptiness. The first one is K-drift boundedness, which imposes for a fixed integer K that for every T_ρ there exists one dated realization such that for every pair of events e, f appearing in the same transition of G, the dates of e and f differ by at most K. We have shown that K-drift boundedness is decidable in a symbolic and efficient way, and that for $K-$drift bounded TC-MSC graphs, emptiness is decidable [52]. This extends decidability results beyond regular specifications. The second restriction is K-non-zenoness, which imposes that for a fixed K, for every path ρ of G, there exists one realization such that at every date d, at most K events occur between date d and $d+1$. When a TC-MSC graph is A-drift-bounded and B-non-zeno, then $L(G)$ has a regular set of representants, which opens the way for more involved model-checking applications [51].

The second part of our work is the study of realistic implementation of scenarios. The main idea is to propose distributed implementation (communicating state machines) of High-level MSCs that do not contain deadlocks, and behave exactly as the original specification. It is well known that a simple projection of a HMSC on each of its process to obtain communicating finite state machines results in an implementation with more behaviors than the original specification. An implementation of a HMSC H is considered as consistent if and only if is exhibits the same prefix closed set of behaviors as H. We have studied how such projection with additional local controlers allows the distributed synthesized behavior to remain consistent with the original specification. This work has been implemented in our scenario prototype (see the Software section). As usually for scenarios, the synthesis algorithm works for a particular syntactic class of scenarios, namely the class of local HMSCs. Roughly speaking, in local HMSC, a decision to behave according to a scenario or another is always taken by a single participant. The deciding process need not be the same at each choice. This class is a sensible restriction of HMSCs, as distributed choices can not be implemented without additional synchronization among processes [53].

Last, we have extended existing results on diagnosis from scenarios [15]. We have shown that when a distributed implementation is instrumented with software probes that publish their observations while the system is running, and when the system is modeled as a High-level MSC, then diagnosis can be expressed as a new HMSC the executions of which are all explanations of the observation. The construction of diagnosis can be performed offline or online, and we have considered the conditions under which online diagnosis can run with finite memory.

6.3. Fundamental results and algorithms: timed models

Participants: Claude Jard, Aurore Junier, Akshay Sundararaman.

Our works in that subject concern Time Petri Nets (TPNs) and Network Calculus. With TPNs, we are particularly interested in symbolic unfoldings (extended with parameters). Possible applications are supervision of distributed timed systems [8] and testing of concurrent systems (work done in collaboration with Stefan Haar, INRIA-LSV in Cachan).

The article [26] was made during the internship of the master degree of Aurore Junier under the supervision of Anne Bouillard (ENS Paris). It uses a (min, plus)-algebra to define a worst-case delay bound for networks where flows have fixed priorities.

After that, we studied a well-known problem: detection of congestion and failure in networks. The idea was to find an efficient and deterministic method that is very reactive and takes little memory space. Such a method does not exists for now and is an important issue for Alcatel-Lucent. We achieved a solution that solves this problem based on the analysis of flows behaviour. This work is part of the work done within the Alcatel-Lucent-Inria joint lab and a patent is being established.

We are also studying the way buffers of routers can increase. The objective is to find a method that can detect if sizes of buffers can dangerously increase on a defined topology. We start by looking at Link State Advertisements (LSA) in the OSPF protocol. We represent the topology and a part of the protocol by a Time Petri Net and try to infer parameters ensuring stability.

6.4. Fundamental results and algorithms: dynamic epistemic logic

Participants: Guillaume Aucher, François Schwarzentruber.

Dynamic Epistemic Logic (DEL) deals with the representation and the study of knowledge and belief change in a multi-agent setting. The core representative task of this logical framework can be split up in three parts: 1/ the inital global state of the distributed system, 2/ an event occurring in this system, 3/ a product update taking as argument these two representations and yielding a new representation of the new global state of the distributed system. Therefore, we can express uniformly within the DEL framework epistemic statements about:

(i) what is true about an initial state

(ii) what is true about an event occurring in this initial state

(iii) what is true about the resulting state after the event has occurred.

We axiomatized within the DEL framework what we can infer about (iii) given (i) and (ii), what we can infer about (ii) given (i) and (iii), and what we can infer about (i) given (ii) and (iii). Given three logical formulas ϕ, ϕ' and ϕ'' describing respectively (i), (ii) and (iii), we also showed how to build three formulas that capture respectively all the information which can be inferred about (iii) from ϕ and ϕ', all the information which can be inferred about (ii) from ϕ and ϕ'', and all the information which can be inferred about (i) from ϕ' and ϕ''. We showed how our results extend to other modal logics than the minimal modal logic K. These results are to appear in [9] and [10]. In [19], we also provided a tableau method deciding whether such inferences are valid. We implemented it in LOTRECscheme and showed that this decision problem is NEXPTIME-complete. This work contributes to the proof theory and the study of the computational complexity of DEL which have rather been neglected so far.

Application to fault localization in IMS network (see the UNIVERSELF project) has started. The various agents involved in an IMS network (clients, assistance, administrators...) have a partial view of the network and so need to communicate their partial knowledge of the network to each other in order to localize the fault in the network (each communication having possibly a different cost). One of the main problems is to determine which communication should occur and which agent should be queried so that the fault is eventually localized. This problem can naturally be expressed in the DEL framework. We have shown how the initial state of an IMS network representing the knowledge of each agent can be represented by a particular kind of epistemic model (i) and how the desired state where the fault is localized can be expressed by a logical formula (iii). The problem amounts to determining which communication or sequence of communications should occur (ii) so that one passes from the initial epistemic model (i) to another epistemic model where the fault is localized (iii), and also to determine if such a communication or sequence of communications is possible. We have focused so far on the case of a single communication, but we plan to extend it to a sequence of communications. Further theoretical work still needs to be done to address the issue of sequential communication.

In parallel to this work, we also axiomatized different notions of knowledge and belief which are defined by means of a 'sphere' semantics. This work is the result of an invited contribution and is to appear in [48].

6.5. Fundamental results and algorithms: statistical model checking

Participants: Sean Sedwards, Cyrille Jégourel, Axel Legay.

Our work on statistical model checking (SMC) avoids an explicit representation of the state space by building a statistical model of the executions of a system and giving results within confidence bounds. The key challenges of this approach are to reduce the length (simulation steps and cpu time) and number of simulation traces necessary to achieve a result with given confidence. Rare properties pose a particular problem in this respect, since they are not only difficult to observe but their probability is difficult to bound. A further goal is to make a tool where the choice of modelling language and logic are flexible.

We have developed the prototype of a compact, modular and efficient SMC platform which we have named *PLASMA* (PLatform for Statistical Model checking Algorithms). PLASMA incorporates an efficient discrete event simulation algorithm and features an importance sampling engine that can reduce the necessary number of simulation runs when properties are rare. We have found that PLASMA performs significantly better than PRISM (the de facto reference probabilistic model checker) when used in a similar mode: PLASMA's simulation algorithm scales with a lower order and can handle much larger models. When using importance sampling, PLASMA's performance with rare properties is even better.

6.6. Fundamental results and algorithms: quantitative model checking and quantitative specification Theories

Participants: Uli Fahrenberg, Axel Legay.

Model checking of systems deals with the question whether a given model of a computer system satisfies the properties one might want to require of it. This is a well-established and successful approach to formal verification of safety-critical computer systems.

When the models of the systems contain quantitative information, the model checking problem becomes complicated by the fact that in most cases, quantitative properties of the systems do not need to be satisfied exactly. Indeed, the model or the properties might be subject to measurement error, or probabilistic information might only be an approximation. In this case, it is of little use to know whether or not a model satisfies a specification precisely; what is needed instead is a notion of *satisfaction distance*: a measure which can assess to which extent a quantitative model satisfies a quantitative specification.

In other words, what is needed is a notion of satisfaction which is robust in the sense that small deviations in the model or the specification only lead to small changes in the outcome of the model checking question. We have published work on such distances in the papers [37], [34].

For more elaborate reasoning about distributed systems or systems-of-systems, an important role is played by specification theories. Such systems are often far too complex to reason about, or model-check, as a whole, and additionally they might be composed of a large number of components which are implemented by different vendors. Hence one needs methods for compositional reasoning, which allow to infer properties of a system from properties of its components, and for incremental design, which allow to synthesize and refine specifications in a step-wise manner.

Such specification theories are by now well-established e.g. in the incarnations of interface theories and modal transition systems. Additionally to defining a formalism for describing and model-checking specifications, they provide notions of refinement of specifications, logical conjunction of specifications, and structural composition and quotient.

When the models and specifications contain quantitative information, all the above notions need to be made robust. One needs to introduce a quantitative version of refinement, and the operations on specifications need to be continuous with respect to refinement distance: compositions of specifications with small refinement distance need themselves to have small refinement distance. We have published work on these issues in the papers [21], [35]; additionally, two other papers within this research area are currently under submission.

6.7. Specific studies: Web services orchestrations

Participants: Ajay Kattepur, Albert Benveniste, Claude Jard.

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 [71].

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 previous years, we proposed soft probabilistic contracts specified as probabilistic distributions involving the different QoS parameters; we studied *contract composition* for such contracts; we developed probabilistic QoS contract monitoring; and we studied the *monotonicity* of orchestrations; an orchestration is monotonic if a called service improves its performance, then so does the overall orchestration.

This year, in the framework of the Associated Team FOSSA with the University of Texas at Austin (John Thywissen (PhD), Jayadev Misra and William Cook), we have extended our approach to general QoS parameters, i.e., beyond response time. 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. John Thywissen (from UT Austin) and Ajay Kattepur have started extending the Orc language and execution engine to support QoS according to our theory. This extension mainly consists in 1/ providing a rich type system to declare QoS domains and related algebra, and 2/ providing a new operator for Orc that allows for selecting competing returns from different sites on the basis of their QoS. A journal paper is under revision.

A key task in extending Orc for QoS was to extend the Orc engine so that causalities between the different site calls are made explicit at run time while execution progresses. This benefits from our previous work on Orc semantics, but a new set of rules has been proposed to generate causalities in an efficient way, by covering new features of the language. This is joint work of Claude Jard, Ajay Kattepur and John Thywissen from Austin. A publication is in preparation.

Besides this main line of work, other topics have been addressed by Ajay Kattepur as part of his thesis.

- In [41], we study variability of composite services. We model variability as a feature diagram (FD) that captures all valid configurations of its orchestration. Then, we apply pair-wise testing to sample the set of all possible configurations to obtain a concise subset. Finally, we test the composite service for selected pairwise configurations for a variety of QoS metrics such as response time, data quality, and availability. Using two case studies, Car crash crisis management and e-Health management, we demonstrate that pairwise generation effectively samples the full range of QoS variations in a dynamic orchestration. The pairwise sampling technique eliminates over 99% redundancy in configurations, while still calling all atomic services at least once.
- Web services orchestrations conventionally employ exhaustive comparison of runtime quality of service (QoS) metrics for decision making. The ability to incorporate more complex mathematical packages is needed, especially in case of workflows for resource allocation and queuing systems. By modeling such optimization routines as service calls within orchestration specifications, techniques such as linear programming can be conveniently invoked by non-specialist workflow designers. Leveraging on previously developed QoS theory, we propose the use of a high-level flexible query procedure for embedding optimizations in languages such as Orc. The Optima site provides an extension to the sorting and pruning operations currently employed in Orc. Further, the lack of an objective technique for consolidating QoS metrics is a problem in identifying suitable cost functions. We use the *analytical hierarchy process* (AHP) to generate a total ordering of QoS metrics across various domains. With constructs for ensuring consistency over subjective judgements, the AHP provides a suitable technique for producing objective cost functions. Using the Dell Supply Chain example, we demonstrate the feasibility of decision making through optimization routines, specially when the control flow is QoS dependent. This work was published in [39].
- With web services quality of service (QoS) modeled as random variables, the accuracy of sampled values for precise service level agreements (SLAs) come into question. Samples with lower spread are more accurate for calculating contractual obligations, which is typically not the case for

web services QoS. Moreover, the extreme values in case of heavy-tailed distributions (eg. 99.99 percentile) are seldom observed through limited sampling schemes. To improve the accuracy of contracts, we propose the use of variance reduction techniques such as importance sampling. We demonstrate this for contracts involving demand and refuel operations within the Dell supply chain example. Using measured values, efficient forecasting of future deviation of contracts may also be performed. A consequence of this is a more precise definition of sampling, measurement and variance tolerance in SLA declarations. This work was published in [40].

6.8. Specific studies: active documents and web services

Participants: Albert Benveniste, Loïc Hélouët, Benoît Masson.

Active Documents have been introduced by the GEMO team at INRIA Futurs, headed by Serge Abiteboul, mainly through the language *Active XML* (or *AXML* for short). 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. In the past years, we have collaborated with the GEMO team to study a distributed version of their language.

This year, we have addressed the problem of distributed documents from a different point of view. Starting from our knowledge of distributed active XML (DAXML), we have first proposed a Petri Net semantics for a subset of DAXML [43], and then considered compositionality issues [54]. Compositionality in services can be addressed in several ways: first one have to ensure that modules that provide services an modules that use them agree on the data that they exchange. This notion is called composability of modules. However, composability does not ensure that a service always terminates (i.e. it returns a result to the caller) when it is invoked with appropriate data. Composability plus termination of services is called compatibility. We have shown that under some restrictions on the recursion in active documents, on the data, and upon the assumption that services use positive guards, composability is decidable. This work has also helped us isolate the core idea behind active documents, and propose a model for them called *Docnets*. Docnets are dynamic Petri nets which places are typed, which transitions are guarded computable type transformations, and which can receive new tokens from their environment. Docnet modules compose well, and if their closure by type transformation is finite, their compatibility is decidable. This work led to a publication [43].

Within the context of the DST associated team, we have proposed a new model, that combines arbitrary numbers of finite workflows, hence allowing for the definition of sessions. Sessions is a central paradigm in web-based systems. As messages exchange between two sites need not follow the same route over the net, a site can not rely on the identity of machines to uniquely define a transaction. This unique identification is essential, as commercial site, for instance, need to manage several interactions at a given time. The current trend, as in BPEL, is to associate a unique identifier to each session. Modeling realistic session hence often forces to include session counters, and hence render most of properties undecidable. The session formalism studied in 2011 can be seen as a mix of BPEL and ORC elements, but was designed to keep several properties decidable. The strength of this formalism is to allow designing systems that use sessions without the obligation to provide identifiers. The formalism has the expressive power of reset Petri nets for which coverability is decidable. This is sufficient to decide whether a set of agent can be found in some bad configuration during the lifetime of a system. This joint work with Ph. Darondeau from the S4 Team, and with M. Mukund from the Chennai Mathematical Institute led to a publication in the ATVA conference [28].

Our last work on Web-services was the development of an experimental platform. During his post-doc, Benoît Masson has designed a distributed Active XML engine, which can be distributed over a network. We have built a lightweight experimentation platform, made of four linux machines, that run DAXML services and communicate with one another. Simultaneously, R. Abdallah has designed a synthesis tool to generate REST services from High-level Message Sequence Charts. These services were successfully tested on the platform.

6.9. Specific studies: security and privacy

Participants: Guillaume Aucher, Blaise Genest.

We have worked on three parallel lines of research related to security and privacy. The first line deals with problems of deleguation and revocation in distributed systems. The second line deals with problems of compliance of a system with respect to a privacy regulation expressed in a language combining epistemic, deontic and dynamic modalities. The third line tackles the minimal information needed at runtime to e.g. break in a (stochastic) system.

6.9.1. Delegation and revocation in distributed systems

Together with Steve Barker from King's College London, Guido Boella from the University of Torino, Valerio Genovese and Leon van der Torre from the University of Luxembourg, we defined a (sound and complete) propositional dynamic logic to specify and reason about delegation and revocation schemes in distributed systems. This logic describes formally a family of delegation and revocation models that are based on the work of [65]. We extended our logic to accommodate an epistemic interpretation of trust. What emerges from this work is a rich framework of formally well-defined delegation and revocation schemes that accommodates an important trust component. In particular, we showed how to automatically reason about whether an agent is authorized to do an operation on an object and about the authorization policy resulting from the execution of a sequence of actions. We used our logical framework to give a formal account of eight different types of revocation schemes informally introduced in previous literature. This work is published in [18].

6.9.2. Privacy policy with modal logic: the dynamic turn

As explained in Section 6.4, we want to define a logical language to specify privacy policies which is close to the natural language. In general, privacy policies can be defined either in terms of permitted and forbidden *knowledge*, or in terms of permitted and forbidden *actions*. For example, it may be forbidden to know the medical data of a person, or it may be forbidden to disclose these data. Implementing a privacy policy based on permitted and forbidden *actions* is relatively easy, since we can add a filter on the system checking the outgoing messages. Such a filter is an example of a security monitor. If the system attempts to send a forbidden message, then the security monitor blocks the sending of that message. However, the price to pay for this relatively straightforward implementation is that it is difficult to decide which actions are permitted or forbidden so that a piece of information is not disclose. We are therefore interested in privacy policies expressed in terms of permitted and forbidden knowledge. Expressing a privacy policy in terms of permitted and forbidden knowledge is relatively easy, since it lists the situations, where, typically, it may not be permitted to know some sensitive information. Implementing a privacy policy based on permitted and forbidden knowledge is quite difficult, since the system has to reason about the relation between permitted knowledge and actions. The challenge is that the exchange of messages changes the knowledge, and the security monitor therefore needs to reason about these changes. This inference problem is already non trivial with a static privacy policy, and becomes challenging when privacy policies can change over time. Together with Guido Boella and Leon van der Torre, we therefore introduced a dynamic modal logic that permits not only to reason about permitted and forbidden knowledge to derive the permitted actions, but also to represent explicitly the declarative privacy policies together with their dynamics. The logic can be used to check both regulatory and behavioral compliance, respectively by checking that the permissions and obligations set up by the security monitor of an organization are not in conflict with the privacy policies, and by checking that these obligations are indeed enforced. We also showed that the complexity of the model checking problem is quadratic in the size of the model and the formula and provided the corresponding model-checking algorithms. This work is published in [11].

6.9.3. Minimal information needed

Together with Nathalie Bertrand from Vertecs, we tackle the problem of the *minimal information* a user needs *at runtime* to achieve a simple goal, modeled as reaching an objective with probability one [25]. The natural question is then to minimize the additional information the user needs to fulfill her objective. This optimization question gives rise to two different problems, whether we consider to minimize the *worst case cost*, or the *average cost*. On the one hand, concerning the worst case cost, we show that efficient techniques from the model checking community can be adapted to compute the optimal worst case cost and give optimal strategies for the users. On the other hand, we show that the optimal average price (a question typically considered in

the AI community) cannot be computed in general, nor can it be approximated in polynomial time even up to a large approximation factor. Following this negative results, we investigate with P.S. Thiagarajan's group at NUS, Singapore basic algorithms of the AI community to infer the exact probability in (compact) stochastic systems. We proposed in [45] a simple parametrized extension of the usual Factored Frontier algorithm in order to choose the desired accuracy of the algorithm, at the cost of additional but managable computations. We showed its benefit when dealing with biological pathways.

6.10. Specific studies: network maintenance

Participants: Eric Fabre, Carole Hounkonnou.

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 methodology for the graceful shut down and restart of routers in OSPF networks, one of the core protocols of IP networks. A methodology has been proposed to safely switch off the software layer of a router while still maintaining this router in the forwarding plane: the router still forwards packets, but is not able to adapt its routing table to changes in network conditions or topology. Nevertheless, it is possible to check whether this frozen router is harmless or can cause packet losses, through a centralized or distributed algorithm. And if ever it puts the network at risk, minimal patches can be set up temporarily until the router comes back to normal a activity. This avoids running twice a global OSPF update at all nodes (once for shutdown of the equipment, one for restart). There is a patent project on this activity, that we don't detail more here.

6.11. Specific studies: network and service diagnosis

Participants: Eric Fabre, Carole Hounkonnou.

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 is also supported by the UniverSelf EU integrated project, and conducted in relation with Orange Labs.

The objective is to develop a framework for the joint diagnosis of networks and of the supported services. We are aiming at a model-based approach, in order to taylor the methods to a given network instance and to follow its evolution. We also aim at active diagnosis methods, that collect and reason on alarms provided by the network, but that can also trigger tests or the collection of new observations in order to refine a current diagnosis.

In 2011, the main effort was dedicated to a key and difficult part of this approach: the definition of a methodology for self-modelling. This consists in automatically building a model of the monitored system, by instantiating generic network elements. There are several difficulties to address:

- The model must capture several layers, from the physical architecture up to the service architecture and its protocols. As a case-study, we have chosen VoIP services on an IMS network, deployed over a wired IP network.
- The model should be hierarchical, to allow for multiscale reasoning, and to reflect the intrinsic hierarchical nature of the managed network.
- The model should be generic, i.e. obtained by assembling component instances coming from a reduced set of patterns, just like a text is obtained by assembling words.
- The model should be adaptive, to capture the evolving part of the network (e.g. introduction of new elements) but also its intrinsically dynamic nature (e.g. opened/closed connections).
- The model should display the hierarchical dependency of resources, specifically the fact that lower-level resources are assembled to provide a support to a higher level resource or functionality.
- The model should allow progressive discovery and refinement: for a matter of size, it is not possible to first build a model of the complete network and then monitor it; one must adopt an approach where the model is build on-line, and where the construction is guided by the progress of the diagnosis algorithms.

The first elements of a methodology achieving these objectives have been designed in 2011. The next efforts will aim at refining the grammar of this model, for our specific case study, and at developing the dedicated diagnosis algorithms. For the latter, we envision a new setting of hierarchical and generic Bayesian networks, in order to capture the dependencies between network elements at different granularities.

7. Contracts and Grants with Industry

7.1. Contracts with Industry

7.1.1. Alcatel-Lucent

Participant: Albert Benveniste.

Title: Laboratoire Commun Alcatel-Lucent BellLabs / Inria

Type: Laboratoire Commun

Defi: The Network of the Future

Duration: January 2008 - January 2012

Coordinator: Albert Benveniste (Inria) and Olivier Audouin (Alcatel-Lucent BellLabs)

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, with a 5-year research program. 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 MarcÈ (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 researchers. It is jointly headed by Olivier Audouin (Bell Labs, president), and Albert Benveniste (INRIA, president of the Scientific Committee). The lab organizes two yearly seminars with progress reports and keynotes by key engineers from Alcatel-Lucent. So far, its production represents 80 papers (21 cosigned), 12 patents and 13 PhDs.

7.1.2. Laboratoire Commun Alcatel-Lucent Bell Labs / Inria: ADR HiMa

Participants: Eric Fabre, Albert Benveniste, Claude Jard, Carole Hounkonnou, Aurore Junier.

Title: High Manageability

Type: Joint Bell Labs INRIA Laboratory *Action de Recherche (ADR)*

Defi: The Network of the Future

Duration: 5 years

Coordinator: Eric Fabre (Inria) and Pierre Peloso (Alcatel-Lucent Bell Labs)

Others partners: involves also members of the Madynes team (Loria), and of the Mascotte team (Sophia)

On the Alcatel-Lucent side, this research group involves 5 persons of the PTI group (Packet Transport Infrastructure). The objective of the ADR 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. The current outcome is of 15 papers, 2 patents [62], 3 PhDs.

In the previous years, DistribCom cosigned two joint patents about the distributed optimization of power allocation in photonic networks. In 2011, the team was involved in three main activities.

- The design of graceful shutdown and restart algorithms for the OSPF protocol (patent project). Corresponds to part of Carole Hounkonnou's PhD, started in 2009.
- The study of network stability when protocol parameters are modified, using network calculus techniques (patent project). Corresponds to Aurore Junier's PhD, started in 2010.
- The definition of a methodology for joint network and service self-diagnosis, in IMS/IP architectures. This corresponds to part of Carole Hounkonnou's PhD.

A great part of the activity of this ADR (in particular the last item above) is now hosted within the EU IP UniverSelf (sept. 2010 - sept. 2013).

8. Partnerships and Cooperations

8.1. Regional Initiatives

8.1.1. Contrat CREATE ActivDoc

Title: ActivDoc

Type: CREATE

Defi: Telecom

Duration: February 2007 - August 2011

Coordinator: Albert Benveniste

Abstract: Activedoc is funded by Région Bretagne. It started in February 2007, for 18 months, and was extended twice for 18 months. This project ended in august 2011, and funded studies on composite web services in a quantitative and qualitative way. The fundamental models studied during this project are models for Quality of Service and models for active documents. We have developed composition techniques for Web systems based on the paradigm of active documents. In particular, Activdoc funded Benoît Masson's Post doctoral stay in Distribcom, which eventually led to the design of a distributed active document simulator.

8.1.2. Contrat CREATE Estase

Title: Estase

Type: CREATE

Defi: Lifting the applicability of formal methods to real life application.

Duration: Three years

Coordinator: Axel Legay

Others partners: None

Abstract: The main objective of the Estase project is to develop new statistical model checking algorithms. In addition, we shall study the concept of stochastic abstraction that allows to abstract the global behavior of a system by probability distribution. The results of Estase shall be implemented in the PLASMA toolset developed at INRIA Rennes.

8.2. National Initiatives

8.2.1. ANR DOTS

Participants: Claude Jard, Loïc Hélouët.

Contract INRIA *ANR-06-SETI January 2007 - December 2011*

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 was originally scheduled to end in December 2010. It was extended for one additional year. 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).

8.2.2. ANR IMPRO

Participants: Claude Jard, Loïc Hélouët, Rouwaida Abdallah, Akshay Sundararaman.

Contract ENS CACHAN *ANR-2010-BLAN-0317 March 2010 - February 2013.*

ImpRo (http://anr-impro.irccyn.ec-nantes.fr/) is an academic research project funded by the French national research agency, within its non-thematic ("Blanc") program. This project addresses the issues related to the practical implementation of formal models for the design of communicating embedded systems: such models abstract many complex features or limitations of the execution environment. The modeling of time, in particular, is usually ideal, with infinitely precise clocks, instantaneous tests or mode commutations, etc. Our objective is thus to study to what extent the practical implementation of these models preserves their good properties. We will first define a generic mathematical framework to reason about and measure implementability, and then study the possibility to integrate implementability constraints in the models. We will particularly focus on the combination of several sources of perturbation such as resource allocation, the distributed architecture of applications, etc. We will also study implementability through control and diagnostic techniques. We will finally apply the developed methods to a case study based on the AUTOSAR architecture, a standard of the automotive industry.

Distribcom cooperates with IRCCyN (Nantes), LIP6 (Paris), LSV (Cachan), LIAFA (Paris) and LIF (Marseille). The coordinator is Didier Lime from IRCCyN. It mainly adresses implementability of scenarios and Time Petri Nets, focusing on concurrency aspects.

8.3. European Initiatives

8.3.1. FP7 Projects

8.3.1.1. DISC

Title: Distributed supervisory control of large plants

Type: COOPERATION (ICT)

Defi: Networked embedded and control systems

Instrument: Specific Targeted Research Project (STREP)

Duration: September 2008 - December 2011

Coordinator: Univ. of Cagliari (Italy)

Others partners: Univ. of Cagliary (IT), CWI (NL), Univ. of Gent (B), Tech. Univ. Berlin (G), Univ. Zaragosa (S), Akhela (IT), CyBio (G)

See also: http://www.disc-project.eu/

Abstract: Supervisory control is a formal approach for the control of discrete event systems that aims to solve logical problems of safety, resource allocation, liveness, and fault diagnosis that can be encountered in all systems with a high degree of automation. It provides a conceptual framework for developing methods and tools for system design.

An open issue is the application of this methodology to those control problems that arise in networked embedded systems. These distributed plants are composed by several local agents that take concurrently decisions, based on information that may be local or received from neighbouring agents; they require scalable and self-organising platforms for advanced computing and control. An important feature of this type of processes is the possibility of studying them at an appropriate level of abstraction where the resulting model is a purely discrete event one. The evolution is guided by the occurrence of asynchronous events, as opposed to other real-time models where the event occurrence is time-triggered.

We plan to use several techniques to reduce the computational complexity that is one of the major obstacles to the technology transfer of supervisory control methodologies to distributed plants. These techniques are: modularity in the modelling and control design phases; coordinating control; fluidisation of some discrete-event dynamics to reduce state-space cardinality; modular state identification and modular fault detection based on the design of decentralized observers.

8.3.1.2. Univerself

Title: UniverSelf

Type: COOPERATION (ICT)

Defi: The Network of the Future

Instrument: Integrated Project (IP)

Duration: September 2010 - August 2013

Coordinator: Alcatel Lucent (France)

Others partners: Alcatel Lucent (F, Ir, G), NEC (G), Thales (F), Orange (F), Telecom Italia (It), Telefonica (E), Univ. College London (GB), Univ. of Surrey (GB), Univ. of Twente (NL), Univ. of Piraeus (G), Univ. of Athens (G), IBBT (B), VTT

See also: http://www.univerself-project.eu/

Abstract: UniverSelf unites 17 partners with the aim of overcoming the growing management complexity of future networking systems, and to reduce the barriers that complexity and ossification pose to further growth.

While there has been undeniable progress in the field of autonomics research over the past several years across the world and especially in Europe, widespread deployments of self-management techniques are still missing. At the same time the need for techniques enabling the transformation of operational models, the evolution of networks towards a flexible playground for operators, and more generally techniques participating to the increase of the return on investment, is becoming more and more evident. Further, most efforts and initiatives have been focussed on solving manageability bottlenecks in a given technological domain, while services extend anywhere, regardless of the technological boundaries (e.g., wireline/wireless). UniverSelf arises from this context and is thus driven by the need and objective to take self-management a leap further, and, in doing so, be both a federating and impactful project.

8.3.1.3. Danse

Type: COOPERATION (ICT)

Defi: Studying Systems of Systems (Dynamical Systems)

Instrument: Integrated Project (IP)

Duration: November 2011 - October 2014

Coordinator: Alcatel Lucent (France)

Others partners: OFFIS Institute for Information Technology (Germany), IBM Israel - Science and Technology LTD (ISRAEL), Israel Aerospace Industries (ISRAEL), Advanced Laboratory on Embedded Systems S.R.L (Italy), INRIA (France), Loughborough University (United Kingdom), EADS Innovation Works (United Kingdom), Selex Sistemi Integrati (Italy)

Abstract: Our objective is to build theory and practice for Systems of Systems (SoS). More precisely, we shall provide more important insights on SoS and understand why they must be treated differently to conventional systems.

8.3.1.4. Dali

Title: Devices for assisted Living

Type: Collaboration (ICT)

Defi: Building an automatic machine capable of assisting elderly people.

Instrument: STREP

Duration: November 2011 - October 2014

Coordinator: Trento (France)

Others partners: University of Trento (Italy), Visual Tools (Spain), Forth (Greece), Northumbria University (United Kingdom), University of Siena (Italy), INRIA (France), INDRA Software (Spain), Siemens AG $\sqrt{}$ñsterreich (Austria)

Abstract: The objective is to build a machine that can help an elderly personn to avoid obstacle. The role of INRIA is to design the algorithm that will run within the engine of the machine.

8.4. International Initiatives

8.4.1. INRIA Associate Teams

8.4.1.1. DST

Title: Distributed Supervision and time

INRIA principal investigator: LoïcHélouët

International Partner:

Institution: National University of Singapore (Singapore)

Laboratory: National University of Singapore

Researcher: Madhavan Mukund

International Partner:

Institution: Chennai Mathematical Institute (India)

Laboratory: Institute for Mathemetical Sciences

Researcher: P.S. Thiagarajan

International Partner:

Institution: Institute of Mathematical Science Chennai (India)

Laboratory: Theoretical Computer Science

Researcher: R. Ramanujam

Duration: 2009 - 2011

See also: http://www.irisa.fr/distribcom/DST09/
This associated team is a tripartite collaboration between two projects at INRIA Rennes (S4 & Distribcom), the National University of Singapore (NUS), and two institutes in Chennai (INDIA), the Chennai Mathematical Institute (CMI) and the Institute of Mathematical Sciences (IMS). The objective of the DST project is to study distributed systems, supervision and time issues with the help of concurrency models. The two main themes of the project are supervision, and quantitative/timed aspects of systems. The supervision theme focuses on distributed scheduling policies of distributed systems to ensure satisfaction of some properties (preservation of some bound on communication channels, for instance), diagnosis, and distributed control techniques. The second theme on time aspects of distributed systems focuses on the analysis of qualitative and quantitative properties of timed systems and models. The quantitative approaches rely on network calculus applied to multimode Real Time Calculus, and the timed models studied during the collaboration are time-constrained scenarios. A recent advance is DST is the elaboration of a model to describe and verify sessions in web-based systems.

8.4.1.2. FOSSA

Title: Formalizing Orchestration & Secure Services Analysis
INRIA principal investigator: Albert Benveniste
International Partner:

Institution: University of Texas Austin (United States)
Laboratory: Computer Science Department

Duration: 2010 - 2012
See also: http://www.irisa.fr/distribcom/FOSSA2010/index.htm
The widespread deployment of networked applications and adoption of the internet has fostered an environment in which many distributed services are available. There is great demand to automate business processes and workflows among organizations and individuals. Solutions to such problems require orchestration or choreography of concurrent and distributed services in the face of arbitrary delays and failures of components and communication. The Orc team, lead by Jayadev Misra at the University of Texas at Austin, has developed the Orc language to support orchestrations. The DistribCom team has developed studies regarding the Quality of Services of orchestrations and choreographies, with emphasis on Orc. The teams cooperate since 2006 and have decided to join their efforts in lauching the associated team FOSSA.

The above tracks have been developed to success in 2011:

- We have come up with a comprehensive theory of QoS for service orchestrations, and more generally composite services. We believe our contract-based approach for QoS is deeply novel and we have submitted a joint paper to the IEEE Transactions on Software, which is curently under revision.
- Causality analysis of Orc programs has been completed. An efficient implementation is under development by John Thywissen (Austin) and Ajay Kattepur, Claude Jard (DistribCom). A joint publication is planed.
- The combination of orchestration languages (such as Orc) and document based workow formalisms (such as Active XML) is of primary interest, as it offers a nice blending of declarative and functional/imperative styles of programming, for large applications. This topic has now started, under the leadership of Loïc Hélouët, with the ongoing deployment on top of Rest of a platform of servers implementing Distributed AXML.

Visits and Exchanges in 2011:

- February 20 – 25, 2011: Albert Benveniste, Claude Jard and Ajay Kattepur visited Austin.
- February 25 – March 4: Ajay Kattepur has extended his stay for one more week in Austin.
- June 27 – July 1st: John Thywissen visited Rennes. Minutes of his stay are available.

9. Dissemination

9.1. Animation of the scientific community

C. Jard has been in 2011 member of the Program Committee of the following international conferences: NOTERE and 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 the Brittany extension of the ENS Cachan (director of the pluridisciplinary institute called the Hubert Curien Research College). He is member of the scientific council of the European University of Brittany and member of the executive board of the CominLabs (french excellence initiative in ICST). 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. He was recently nominated to the National Council of Universities (CNU). In 2011, C. Jard was president of the PhD or Habilitation jurys of V. Gripon (Telecom Bretagne), L. Paulevé (Ecole Centrale Nantes), S. Brault (University of Rennes 2), P. Niebert (University of Marseille).

A. Benveniste is the Scientific Director of the CominLabs Excellence Center (Laboratoire d'Excellence, part of the program *Investissements d'Avenir* of the french government). He is member of the Strategic Advisory Council of the Institute for Systems Research, Univ. of Maryland, College Park, USA. He is president of the Scientific Committee of the *Common Bell Labs INRIA Laboratory*. He is member of the Scientific Council of France Telecom.[1] Albert Benveniste has been elected to the French Academy of Technologies.

Eric Fabre is associate editor (AE) for the journal *IEEE Trans. on Automatic Control*. He participated to the Alcatel Lucent Open Days in Lannion, where he presented the results of the joint lab on the distributed optimal power allocation in photonic networks. Eric Fabre was reviewer of Jingxian Lu's PhD thesis (Univ. of Bordeaux 1).

L. Hélouët is the co-organizer of a weekly Seminar (68NQRT) at IRISA-INRIA Rennes on theoretical aspects of computer science. In 2011, L. Hélouët was member of the jury of Gilles Benattar at IrccYn Nante, and of Mario S. Alvim at Ecole Polytechnique.

G. Aucher was an organizer of the workshop Gipsy which took place in Rennes from the 25th to the 27th of October 2011 (http://www.irisa.fr/prive/pinchina/GIPSy/gipsy11.html). In 2011, he served as a reviewer in the following international journals: *Journal of Logic and Computation, Studia Logica, Mathematical Social Science, Journal of Applied Non-classical Logic, Logique et analyse*. He was a program committee member of the international conferences IJCAI 2011 (International Joint Conference in Artificial Intelligence) and LORI 2011 (Logic, Rationality and Interaction) and served as external reviewer for the international conference TARK 2011 (Theoretical Aspects of Rationality and Knowledge). He was invited to give talks at the Beth Foundation Symposium at the 14th Congress of Logic, Methodology and Philiosophy of Science in Nancy on the 20th of July 2011 and at the Institute of Philosophy of Language in Lisbon on the 8th of November 2011.

9.2. Teaching

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. He teaches the following courses: L. Hélouët teaches algorithms (graphs, sorting algorithms, classical problems) to maths students preparing the aggregation at ENS Cachan, antenne de Bretagne. He is also involved in the INS project, which aim is to give the basic knowledge in computer science to teachers from other disciplines at the baccalaureat level. In this context, he supervises a development project by 5 teachers from Rennes academy. He is the co-supervisor (with C. Jard) of Rouwaida Abdallah's Thesis. Eric Fabre teaches Numerical optimization methods, and Distributed algorithms and systems in the Master of computer

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

science (M2), University Rennes 1. He teaches also Information theory and coding at ENS Cachan, magister Math, Info., Telecom, 2nd year.

Guillaume Aucher teaches Algorithms for graphs (L3 MIAGE, 20h TD) and Imperative programming in Java (L1 Computer Science, 20hTD+20hTP), both at the Univ. of Rennes 1.

François Schwarzentruber is a lecturer at ENS Cachan - Britany extension. He teaches an introductory course about algorithms to first year students of ENS (L3) and a course about conception and verification of programs to second year students (M1). He is also involved in the preparation of the computer science option of the competitive examination 'agregation de maths'.

PhD & HdR

PhD : Bartosz Grabiec, Supervision de systèmes répartis utilisant des dépliages avec contraintes de modèle temporisés, ENS Cachan, 4 Octobre 2011, C. Jard.

PhD in progress : Ajay Kattepur, QoS in Web services, september 2009, A. Benveniste, C. Jard.

PhD in progress : Rouwaida Adballah, Implémentation de scénarios, march 2010, L. Hélouët, C. Jard.

PhD in progress : Cyrille Jégourel, Statistical Model Checking, January 2011, Axel Legay.

PhD in progress : Valérie Murat, Tree automata, January 2011, Axel Legay.

PhD in progress : Loig Jézéquel, Distributed optimal planning methods, Oct. 2009, Eric Fabre.

PhD in progress : Carole Hounkonnou, A methodology for joint network and service self-diagnosis, Oct. 2009, Eric Fabre.

PhD in progress : Aurore Junier, Network calculus applied to network stability analysis, sept. 2010, C. Jard, A. Bouillard.

PhD in progress : Bastien Maubert, Logical foundations of games with imperfect information, sept. 2010, S. Pinchinat, G. Aucher.

10. 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", 2008, vol. 390, n^{o} 2-3, 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", May 2003, vol. 48, n^{o} 5, p. 714–727.

[3] T. CHATAIN, C. JARD. *Complete finite prefixes of symbolic unfoldings of safe time Petri nets*, in "Proc. of ICATPN", LNCS, June 2006, n^{o} 4024, p. 125–145.

[4] 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)", 2007, vol. 17, p. 355-403.

[5] E. FABRE. *Trellis Processes: a Compact Representation for Runs of Concurrent Systems*, in "Journal of Discrete Event Dynamic Systems (JDEDS)", 2007, vol. 17, n^{o} 3, p. 267–306.

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

[7] T. GAZAGNAIRE, B. GENEST, L. HÉLOUËT, P. S. THIAGARAJAN, S. YANG. *Causal Message Sequence Charts*, in "Theoretical Computer Science (TCS)", 2009, vol. 410, n° 41, p. 4094-4110.

Publications of the year

Doctoral Dissertations and Habilitation Theses

[8] B. GRABIEC. *Distributed system supervision using timed constrained unfoldings*, ENS Cachan, 2011.

Articles in International Peer-Reviewed Journal

[9] G. AUCHER. *DEL-Sequents for progression*, in "Journal of Applied Non-Classical Logics", 2011, to appear.

[10] G. AUCHER. *DEL-Sequents for regression and epistemic planning*, in "Journal of Applied Non-Classical Logics", 2011, to appear.

[11] G. AUCHER, G. BOELLA, L. VAN DER TORRE. *A dynamic logic for privacy compliance*, in "Journal of artificial intelligence and law", 2011, vol. 19, n° 2–3.

[12] L. BOZZELLI, A. LEGAY, S. PINCHINAT. *Hardness of preorder checking for basic formalisms*, in "Theor. Comput. Sci.", 2011, vol. 412, n° 49, p. 6795-6808.

[13] B. CAILLAUD, B. DELAHAYE, K. G. LARSEN, A. LEGAY, M. L. PEDERSEN, A. WASOWSKI. *Constraint Markov Chains*, in "Theor. Comput. Sci.", 2011, vol. 412, n° 34, p. 4373-4404.

[14] B. DELAHAYE, B. CAILLAUD, A. LEGAY. *Probabilistic contracts: a compositional reasoning methodology for the design of systems with stochastic and/or non-deterministic aspects*, in "Formal Methods in System Design", 2011, vol. 38, n° 1, p. 1-32.

[15] T. GAZAGNAIRE, B. GENEST, L. HÉLOUËT, H. MARCHAND. *Diagnosis from Scenarios and Applications*, in "Journal of Discrete Events and Dynamic Systems", 2011, to appear.

[16] L. JEZEQUEL, E. FABRE. *On the construction of probabilistic diagnosers for modular systems*, in "Journal of Discrete Events Dynamical Systems", 2012, to appear.

[17] J.-B. RACLET, E. BADOUEL, A. BENVENISTE, B. CAILLAUD, A. LEGAY, R. PASSERONE. *A Modal Interface Theory for Component-based Design*, in "Fundam. Inform.", 2011, vol. 108, n° 1-2, p. 119-149.

International Conferences with Proceedings

[18] G. AUCHER, S. BARKER, G. BOELLA, V. GENOVESE, L. VAN DER TORRE. *Dynamics in Delegation and Revocation Schemes: A Logical Approach*, in "25th IFIP WG 11.3 International Conference on Data and Applications Security and Privacy, DBSec 2011", Y. LI (editor), 2011, p. 90–105.

[19] G. AUCHER, B. MAUBERT, F. SCHWARZENTRUBER. *Tableau Method and NEXPTIME-completeness of DEL-sequents*, in "Proceedings of Methods for Modalities (M4M 2011)", ENTCS, 2011.

[20] S. S. BAUER, U. FAHRENBERG, L. JUHL, K. G. LARSEN, A. LEGAY, C. R. THRANE. *Quantitative Refinement for Weighted Modal Transition Systems*, in "MFCS", 2011, p. 60-71.

[21] S. S. BAUER, U. FAHRENBERG, L. JUHL, K. G. LARSEN, A. LEGAY, C. THRANE. *Quantitative Refinement for Weighted Modal Transition Systems*, in "MFCS, Mathematical Foundations of Computer Science 2011 - 36th International Symposium, MFCS 2011", Warsaw, Poland, August 22-26 2011, p. 60-71.

[22] S. S. BAUER, P. MAYER, A. LEGAY. *MIO Workbench: A Tool for Compositional Design with Modal Input/Output Interfaces*, in "ATVA", 2011, p. 418-421.

[23] S. BENSALEM, L. DE SILVA, A. GRIESMAYER, F. INGRAND, A. LEGAY, R. YAN. *A Formal Approach for Incremental Construction with an Application to Autonomous Robotic Systems*, in "Software Composition", 2011, p. 116-132.

[24] S. BENSALEM, A. GRIESMAYER, A. LEGAY, T.-H. NGUYEN, J. SIFAKIS, R. YAN. *D-Finder 2: Towards Efficient Correctness of Incremental Design*, in "NASA Formal Methods", 2011, p. 453-458.

[25] N. BERTRAND, B. GENEST. *Minimal Disclosure in Partially Observable Markov Decision Processes*, in "31th IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS)", 2011, vol. LIPIcs, to appear.

[26] A. BOUILLARD, A. JUNIER. *Worst-case delay bounds with fixed priorities using network calculus*, in "Proc. of Valuetools'2011", 2011.

[27] A. CLASSEN, P. HEYMANS, P.-Y. SCHOBBENS, A. LEGAY. *Symbolic model checking of software product lines*, in "ICSE", 2011, p. 321-330.

[28] P. DARONDEAU, L. HÉLOUËT, M. MUKUND. *Assembling Sessions*, in "ATVA", Lecture Notes in Computer Science, Springer, 2011, vol. 6996, p. 259-274.

[29] A. DAVID, K. G. LARSEN, A. LEGAY, M. MIKUCIONIS, D. B. POULSEN, J. VAN VLIET, Z. WANG. *Statistical Model Checking for Networks of Priced Timed Automata*, in "FORMATS", 2011, p. 80-96.

[30] A. DAVID, K. G. LARSEN, A. LEGAY, M. MIKUCIONIS, Z. WANG. *Time for Statistical Model Checking of Real-Time Systems*, in "CAV", 2011, p. 349-355.

[31] B. DELAHAYE, J.-P. KATOEN, K. G. LARSEN, A. LEGAY, M. L. PEDERSEN, F. SHER, A. WASOWSKI. *Abstract Probabilistic Automata*, in "VMCAI", Lecture Notes in Computer Science, Springer, 2011, vol. 6538, p. 324-339.

[32] B. DELAHAYE, K. G. LARSEN, A. LEGAY, M. L. PEDERSEN, A. WASOWSKI. *APAC: A Tool for Reasoning about Abstract Probabilistic Automata*, in "QEST", 2011, p. 151-152.

[33] B. DELAHAYE, K. G. LARSEN, A. LEGAY, M. L. PEDERSEN, A. WASOWSKI. *Decision Problems for Interval Markov Chains*, in "LATA", 2011, p. 274-285.

[34] U. FAHRENBERG, A. LEGAY, C. THRANE. *The Quantitative Linear-Time–Branching-Time Spectrum*, in "FSTTCS", LIPIcs, Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2011, To appear.

[35] U. FAHRENBERG, A. LEGAY, A. WASOWSKI. *Vision Paper: Make a Difference! (Semantically)*, in "Model Driven Engineering Languages and Systems, 14th International Conference, MODELS 2011", Wellington, New Zealand, October 16-21 2011, p. 490-500.

[36] U. FAHRENBERG, A. LEGAY, A. WASOWSKI. *Vision Paper: Make a Difference! (Semantically)*, in "MoDELS", 2011, p. 490-500.

[37] U. FAHRENBERG, C. THRANE, K. G. LARSEN. *Distances for Weighted Transition Systems: Games and properties*, in "QAPL", Electronic Proceedings in Theoretical Computer Science, 2011, vol. 57, p. 134-147.

[38] L. JEZEQUEL, E. FABRE. *Networks of automata with read arcs: a tool for distributed planning*, in "IFAC World Congress", Sept. 2011.

[39] A. KATTEPUR, A. BENVENISTE, C. JARD. *Optimizing Decisions in Web Services Orchestrations*, in "9th International Conference on Service-Oriented Computing (ICSOC)", Springer, 2011, p. 77-91.

[40] A. KATTEPUR. *Importance Sampling of Probabilistic Contracts in Web Services*, in "9th International Conference on Service-Oriented Computing (ICSOC)", Springer, 2011, p. 557-565.

[41] A. KATTEPUR, S. SEN, B. BAUDRY, A. BENVENISTE, C. JARD. *Pairwise testing of dynamic composite services*, in "6th international symposium on Software engineering for adaptive and self-managing systems (SEAMS)", New York, NY, USA, SEAMS '11, ACM, 2011, p. 138–147, http://doi.acm.org/10.1145/1988008.1988028.

[42] K. G. LARSEN, A. LEGAY, L.-M. TRAONOUEZ, A. WASOWSKI. *Robust Specification of Real Time Components*, in "FORMATS", 2011, p. 129-144.

[43] B. MASSON, L. HÉLOUËT, A. BENVENISTE. *Compatibility of Data-Centric Web Services*, in "WS-FM, 8th International Workshop on Web Services and Formal Methods", Lecture Notes in Computer Science, Springer, 2011, to appear.

[44] J. J. ORTIZ, A. LEGAY, P.-Y. SCHOBBENS. *Distributed Event Clock Automata - Extended Abstract*, in "CIAA", 2011, p. 250-263.

[45] S. PALANIAPPAN, S. AKSHAY, B. GENEST, P. THIAGARAJAN. *A Hybrid Factored Frontier Algorithm for Dynamic Bayesian Networks*, in "9th International Conference on Computational Methods in Systems Biology (CMSB)", 2011, vol. ACM 978-1-4503-0817-5, p. 35-44.

National Conferences with Proceeding

[46] R. ABDALLAH, C. JARD. *An experiment in automatic generation of protocols from HMSCs*, in "Notere", May 2011.

Conferences without Proceedings

[47] G. AUCHER, C. BARREAU-SALIOU, G. BOELLA, A. BLANDIN-OBERNESSER, S. GAMBS, G. PIOLLE, L. VAN DER TORRE. *The Coprelobri project : the logical approach to privacy*, in "2e Atelier Protection de la Vie Privée (APVP 2011)", Sorèze, France, June 2011, 6 pages.

Scientific Books (or Scientific Book chapters)

[48] G. AUCHER. *From belief to knowledge with a 'sphere' semantics: axiomatizations*, in "Dialogue, Rationality, Formalism.", G. HEINZMANN, M. MUSIOL, M. REBUSCHI, A. TROGNON (editors), Logic, Epistemology and the Unity of Science, Springer, 2011, to appear.

[49] E. FABRE. , C. SEATZU, M. SILVA, J. H. VAN SCHUPPEN (editors)*Distributed Control of Large Plants; Chapter 5, Observers and automata*, Springer, 2012, to appear.

[50] S. HAAR, E. FABRE. , C. SEATZU, M. SILVA, J. H. VAN SCHUPPEN (editors)*Distributed Control of Large Plants; Chapter 13, Diagnosis with Petri nets unfoldings*, Springer, 2012, to appear.

Research Reports

[51] S. AKSHAY, B. GENEST, L. HÉLOUËT, S. YANG. *Regular Set of Representatives for Time-Constrained MSC Graphs*, 2011, http://perso.crans.org/~genest/AGHY11b.pdf.

[52] S. AKSHAY, B. GENEST, L. HÉLOUËT, S. YANG. *Symbolically Bounding the Drift in Time-Constrained MSC Graphs*, 2011, http://perso.crans.org/~genest/AGHY12.pdf.

[53] C. JARD, R. ABDALLAH, L. HÉLOUËT. *Realistic Implementation of Message Sequence Charts*, INRIA, April 2011, n^{o} RR-7597, http://hal.inria.fr/inria-00584530/en/.

[54] B. MASSON, L. HÉLOUËT, A. BENVENISTE. *Compatibility between DAXML Schemas*, INRIA, March 2011, n^{o} RR-7559, http://hal.inria.fr/inria-00573774/en/.

Other Publications

[55] A. DAVID, K. G. LARSEN, A. LEGAY, M. MIKUCIONIS, D. B. POULSEN, J. VAN VLIET, Z. WANG. *Stochastic Semantics and Statistical Model Checking for Networks of Priced Timed Automata*, 2011, CoRR.

References in notes

[56] S. ABITEBOUL, O. BENJELLOUN, I. MANOLESCU, T. MILO, R. WEBER. *Active XML: A Data-Centric Perspective on Web Services*, in "BDA'02", 2002.

[57] T. ANDREWS, F. CURBERA, H. DHOLAKIA, Y. GOLAND, J. KLEIN, F. LEYMANN, K. LIU, D. ROLLER, D. SMITH, S. THATTE, I. TRICKOVIC, S. WEERAWARANA. *Business Process Execution Language for Web Services (BPEL4WS). Version 1.1*, 2003, http://xml.coverpages.org/BPELv11-May052003Final.pdf.

[58] M. G. BUSCEMI, V. SASSONE. *High-Level Petri Nets as Type Theories in the Join Calculus*, in "FoSSaCS", Lecture Notes in Computer Science, Springer, 2001, vol. 2030, p. 104-120.

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

[60] 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.

[61] E. FABRE, L. JÉZÉQUEL. *Distributed Optimal Planning: an Approach by Weighted Automata Calculus*, in "Conference on Detection and Control (CDC)", 2009.

[62] E. FABRE, P. PELOSO, P. PECCI. *Method and equipment for adjusting power amplification in an optical network*, in "European patents EP 09290408 and EP 09290409", 2009.

[63] R. FAGIN, J. HALPERN, Y. MOSES, M. VARDI. *Reasoning about knowledge*, MIT Press, 1995.

[64] P. GASTIN, K. NARAYAN. KUMAR, M. MUKUND. *Reachability and boundedness in time-constrained MSC graphs*, in "Perspectives in Concurrency Theory, festchrift for P.S. Thiagarajan", 2008.

[65] Å. HAGSTRÖM, S. JAJODIA, F. PARISI-PRESICCE, D. WIJESEKERA. *Revocations-A Classification*, in "Procs. of CSFW-14", 2001, p. 44-58.

[66] J. HALPERN, Y. MOSES. *Knowledge and common knowledge in a distributed environment*, in "Journal of the ACM", 1990, vol. 37, n^{o} 3, p. 549–587.

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

[68] L. JEZEQUEL, E. FABRE, P. HASLUM, S. THIEBAUX. *Cost-Optimal Factored Planning: Promises and Pitfalls*, in "ICAPS, Int. Conf. on Applications of Planning and Scheduling", May 2010.

[69] D. KITCHIN, W.R. COOK, J. MISRA. *A Language for Task Orchestration and Its Semantic Properties*, in "CONCUR'06", 2006, p. 477-491.

[70] J.-J. C. MEYER, W. VAN DER HOEK. *Epistemic Logic for AI and Computer Science*, Cambridge University Press, Cambridge, 1995.

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

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

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

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

[75] 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.

[76] M. RENIERS. *Message Sequence Charts: Syntax and Semantics*, Eindhoven University of Technology, 1998.

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

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

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

[80] J. VAN BENTHEM, J. GERBRANDY, E. PACUIT. *Merging Frameworks for Interaction: DEL and ETL*, in "Theoretical Aspect of Rationality and Knowledge (TARK XI)", Brussels, D. SAMET (editor), June 2007, p. 72–82.