Inria international program

Associate Team proposal (2013-2015)

Submission form

Please name this file “acronymeoftheassociateteam Proposal 2013.pdf”
Online submission on https://international-programs.inria.fr
Deadline: September 30, 2012

Associate Team acronym: DISTOL

(Distributed systems, stochastic models and logics)

Principal investigator (Inria): Loïc Hélouët (Distribcom, INRIA Rennes)

Principal investigator (partner): Madhavan Mukund (Chennai Mathematical Institute(CMI), Chennai, India)

Principal investigator (partner): P.S. Thiagarajan (National University of Singapore)

Other participants: S4 (INRIA Rennes), Vertecs (INRIA Rennes)
Institute of Mathematical Sciences (IMSC), Chennai, India

1. Partnership

1.1 Detailed list of participants

Distribcom, INRIA Rennes: (d1) Loïc Hélouët (CR1 INRIA, senior researcher) http://people.rennes.inria.fr/Loic.Helouet, (d2) Guillaume Aucher (Chaire INRIA/univ. Rennes 1) http://www.irisa.fr/prive/gaucher/ (d3) Francois Schwarzentruber (ENS Cachan-Ker Lann) http://www.irisa.fr/prive/fschwarz/ (d4) S. Akshay (Post Doc) http://people.irisa.fr/Akshay.Sundararaman/
S4, INRIA Rennes : (s1) Philippe Darondeau (DR Inria, senior researcher) http://www.irisa.fr/s4/wg22/phd/ (s2) Sophie Pinchinat (Professor, Univ. Rennes 1) http://people.irisa.fr/Sophie.Pinchinat/ (s3) Bastien Maubert (PhD Student) http://www.irisa.fr/prive/bmaubert/
Vertecs, INRIA Rennes : (v1) Nathalie Bertrand http://www.irisa.fr/prive/nbertran/ (v2) Paulin Fournier (PhD Student)
CMI, Chennai, India : (c1) Madhavan Mukund (Professor and Dean of Studies) http://www.cmi.ac.in/~madhavan/ (c2) Narayan. K. Kumar (Professor) http://www.cmi.ac.in/~kumar/, (c3) S.P. Suresh (Associate Professor), http://www.cmi.ac.in/~spsuresh/ (c4) Prakash Saivasan (PhD Student) (c5) Prateek Karandikar (PhD Student) (c6) Gautham Shenoy R (PhD Student)
IMSC, Chennai, India : (i1) R. Ramamnujam http://www.imsc.res.in/~jam/ (i2) K. Lodaya http://www.imsc.res.in/~kamal/ (i3) Anup Basil Mathew (PhD Student) (i4) Ramchandra Phawade (PhD student)
NUS, Singapore : (n1) P.S. Thiagarajan (Professor,) http://www.comp.nus.edu.sg/~thiagu/, (n2) Blaise Genest (NUS Adjunct A/Prof. and CR1 CNRS) http://perso.crans.org/~genest/ (n3) Bruno Karelovic (PhD Student) (n4) Sucheendra Palaniappan (PhD Student)

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

- **R1** : Robustness and time issues in distributed systems models (Distribcom +S4 : competences in robustness, models for distributed systems – CMI : competences in models for timed and distributed systems)
- **R2** : Applications of formal models & techniques to Web Services (Distribcom +S4 : competences on modeling of Web Services – CMI : competences in modeling of Web services and verification of distributed systems)
- **R3:** Quantitative verification for distributed systems (Distribcom + Vertecs: competences in probabilities, markovian models +NUS : competences in inference in Bayesian networks)
- **R4** : Unification of Control Theory of Distributed Systems (S4+Distribcom : competences in logics – IMSC : competences in logics, control theory, games)

1.2 Nature and history of the collaboration

History: All members of the project have collaborations for a long time. P.S. Thiagarajan and P. Darondeau collaborated in European project in the 80's. In 2006, Distribcom and NUS launched an associated team (CASDS, 2007-2009), joined by CMI and IMSC two years later. This new consortium then applied successfully for a renewed associated team (DST, 2009-2011). These two teams led to many exchanges (an average of 7 bilateral visits each year in the lifetime of DST). In addition to the research visits, several researchers moved from one institute to another. B.Genest moved from Distribcom to Singapore in 2009. S.Akshay came for a post-doc in Distribcom after completing a PhD under the supervision of M. Mukund, and a post-doc in NUS with P.S. Thiagarajan. S. Yang, completed a PhD in NUS before joining Distribcom as a post-doc (2008-2010).

Thanks to these lively collaborations, researchers from Rennes were invited to join the CNRS international associated laboratory (LIA) INFORMEL, which involves several french research institutes, the CMI, IMSC, and the Indian Institute of Science. The goal of INFORMEL is to strengthen and extend the scientific collaborations between India and France in the domain of formal methods and verification of complex systems.

During the lifetime of CASDS and DST, many common publications in renowned conferences and journals were published (see the attached publication list). After the end of the DST team, the collaborations stayed alive, and led to new joint publications and common work. M. Mukund and N.K. Kumar came to Rennes, B. Maubert (PhD) did an internship at IMSC in June 2012. Last, M. Mukund and P.S. Thiagarajan will be members of L. Hélouët's HDR jury planned this year.

Several moves are also planned in the next coming months: S. Akshay is likely to move to India, but will stay connected to the project via IMSC or CMI. B. Genest will come back in Distribcom, but will keep strong connections with Singapore.

Complementarity :

All members of the team have been working on close topics in the general context of formal methods, distributed systems, stochastic models and logics. However, there are complementarities in all topics.

R1 : L.Hélouët and S. Akshay (Rennes) are involved in research on robustness of timed models (via the IMPRO ANR project). B.Genest also contributes to the study of timed models for distributed systems. At CMI, N.K. Kumar was involved in project on verification of time-constrained scenarios, and started working on robustness issues in Time Petri nets with L.Hélouët and S. Akshay during his last visit in Rennes (2012).

R2 : M. Mukund ,P. Darondeau, and L.Hélouët launched jointly research activities on formal modeling and verification of Web Services. Web Services is a research topic in Distribcom since 2009, but the collaboration

with M.Mukund was the occasion to concretize ongoing discussion in both groups. We share a common background on true concurrency models (communicating automata, scenarios, Petri nets,…). M. Mukund has strong competences in implementation models such as communicating finite state machines, while researchers in Rennes often consider models at a more global level of abstraction (Petri nets, process algebras,…).

R3: P.S.Thiagarajan & B. Genest have gained expertise in the last years on the computation of probabilities in large distributed stochastic networks (DBN…). This is used to perform parameter estimation in biological pathways. On the other hand, N. Bertrand is recognized as an expert of probabilistic models and quantitative aspects of verification. These three researchers are involved in the STIC Asie proposal SQALE (see description later) on scalable quantitative analysis of large distributed systems

R4: R. Ramanujam, K. Lodaya, S. Pinchinat, have a strong background in temporal logic. Moreover, R. Ramanujam is also an expert in epistemic logic, game theory and their applications to distributed systems. This aspect will be strengthened by the competences in dynamic epistemic logic of G. Aucher and the competences of F. Schwarzentruber in game theory and modal logic. Besides, S. Pinchinat and R. Ramanujam have already strong competences in control theory and the application of game theory to computer science. To sum up, they all share an interest in (non-classical) logic, game theory and control theory.

2. Scientific program

2.1 Context

The context of this project is formal modeling, and analysis of behaviors of distributed systems. We want to address verification and supervision of distributed systems through formal modeling and automated reasoning on models. By distributed systems, we mean software architectures made of several independent communicating entities. In the 90's the kind of system addressed was mainly telecommunication protocols. Nowadays, distributed systems are frequently web-based systems such as Web Services, but several aspects of distributed systems can be found in biological applications. Within this context, there are several key challenges to bring formal tools with applications to real systems.

The first one is to ensure that models considered are **realistic**. A model can be an abstraction of a real system, but one should ensure that such abstraction does not affect important properties. It is then important to focus on actual distributed systems (such as Web Services), how to represent them. We plan to develop **implementable but yet tractable models for Web services**. The simple nature of Web services (workflows with structured dat) calls for the use of formal tools such as Well-Structured Transition systems [O-FS01]. A lot of efforts is also devoted to model services in the pi-calculus community [O-HYC08]. A second aspect to consider is robustness of formal models w.r.t. verification of properties, i.e. the question of whether a property checked on a model still hold for an implementation of the model.In particular, we want to consider timed robustness: most models have an idealized representation of time (global clock, infinite precision,…) , that a real system cannot implement. Puri [O-Puri00] showed that a slight change in the semantics of time could alter properties of a model, and Bouyer [O-BMS11] showed how to check robustness of ω-regular properties for timed automata. We plan to study **robustness issues for true concurrency models**, with distinct and imprecise clocks.

A second challenge is to find **quantitative** rather than Boolean answers to formal properties: knowing the probability of some set of paths and the most probable path is usually much more informative than knowing the set of possible path satisfying some property. We focus on qualitative (« almost surely, a call to a service is successful") or quantitative ("the average failure rate is lower than 0.01").One possibility to obtain the probability is to compute its exact value. Such questions have answers for markovian models (Markov Chains, Markov Decision Process subsuming finite state systems) and some quantitative logic (PCTL mainly) [O-HJ94]. However, when dealing with very distributed systems (that can be parametric) and/or other logics [O-KVAK10], computing exactly the probability may be computationaly infeasible. Dynamic Bayesian Networks for instance allows representing compactly very distributed systems. Approximated inference of the probabilities is a pragmatic solution in that cases [O-BK98,O-MW01].

The third challenge addressed by this project is **unification of control** for distributed systems. The theory of supervisory control aims at synthesizing supervisors, whose role is to control the behavior of a discrete event system so as to produce a specified behavior [O-RW89,O-CL08], and is addressed under various assumptions (partial control or observation, decentralization,…). Recent approaches consider distributed control with communication. However, the formalisms tend to be quite complex and it is difficult to derive ***automatically*** algorithms to solve the problems of the theory of supervisory control. This part of the project is in line with Halpern and Moses who hoped in their seminal paper [O-HM90] that "a theory of knowledge, communication, and action will prove rich enough to provide general foundations for a unified theoretical treatment of distributed systems". We want to consider connections between problems from other fields (like distributed artificial intelligence) and theory of supervisory control, to benefit from the different approaches brought in different communities. Several logics [O-AHK98, O-HW02], are the result of interactions between logicians and game-theorists. Such cross-fertilization is possible for control theory. In particular, we will investigate to which extent techniques from epistemic reasoning and game theory can be applied to supervisory control theory, which could lead to the emergence of a unified theoretical framework.

A more detailed research program is attached as appendix 7.1 to the proposal.

2.2 Objectives (for the three years)

R1: Our main objective is to consider robustness for true concurrency models in a context where each process has its own measurement of time. We will start this study with timed variants of Petri nets, building on former results on this model [E-AHJLR12,E-AHJR12], and on experience gained for automata with independently evolving clocks [E-AGMK08]. A first step is to formalize independence of processes and clocks in time(d) Petri nets, and the robustness problems for this model. Several problems could be undecidable, so the next step is to find reasonable restrictions ensuring the existence of (semi) decision procedures.

R2: We want to consider **realistic models for Web-Services**. We have already proposed a session model for Web Services [J-DHM11]. It describes finite sets of agents running an arbitrary number of concurrent transactions. Coverability of some (bad) configuration is decidable for this model. We first want to extend our model and decision procedures to systems with arbitrary numbers of agents. A key challenge is to build realistic but well-structured models, to allow straightforward decidability of interesting safety properties. The second objective is to consider more elaborated properties than coverability, such as **conflicts of interest** between agents, etc. Overall, we wish to propose a highly expressive model together with a decidable logic to reason on this model. The techniques used to reach this goal will build on our knowledge of Well-Structured Transition System [O-FS01], and Petri nets variants. The last and most exploratory objective is **monitoring for session systems**: from a model M of a system, an implementation *I* of this model, and a property to monitor, we want to instrument I with observers (synthesized from M) that raise an alarm when they are sure that the property is violated. Monitoring was studied for pi-calculus [O-HYC08], but it is not yet clear whether the proposed solutions apply to our setting.

R3 : In the next three years, we will develop algorithms to compute precisely probabilities of logical properties, in particular in the presence of imperfect information and/or time. We will build on our work in [J-BG11]. We will also Improve the precision of approximated inference algorithms for distributed (parametrized or not) systems, and deduce formal bounds that guarantee the probability to be in an interval of bounded size. For that, we will develop the techniques we introduced in [J-PAGT11]. In particular, we will provide a decomposition algorithm such that the global approximated probability will be more accurate (through a better accuracy on each component) than by considering the system as a monolith. This has been a major objective in analysis of distributed system, but in general, it cannot be reached exactly. However, approximation gives more freedom for clustering. Last, we will develop approximated verification for logics different from PCTL, leveraging on [J-AAGT12].

R4: The goal of this research direction is a **unified theoretical framework for supervisory control theory**. We will investigate to which extent techniques from epistemic reasoning and game theory can be applied to address control problems for distributed systems. The first milestone will be to reformulate supervisory control in logical and game-theoretical terms. In that respect, epistemic logic should help to handle partial observation. The second milestone will consist in bringing together epistemic logic and imperfect information games to handle individual (i.e. subsystems) knowledge. It is a challenging task because, taking apart control theory issues, the logical foundation of games with imperfect information is an emerging field with only few results [O-GDE11,E-MPB11]. The third milestone will consist in incrementing the previous framework by considering communication mechanisms between the subsystems. In game theory, communication between players is very primitive, whereas in epistemic logic, there are powerful rigorous ways to model effects of atomic communication events on the individual knowledge. It is a challenging task to transfer this apparatus to games and will probably lead to new results in game theory but more importantly, in distributed control. The fourth milestone will consist in studying properties of the developed unified framework, both computational and in terms of expressiveness. For this, we may link the new framework with existing logical formalisms and/or game-based settings.

Our objectives are described with more detail in appendix 7.2.

2.3 Work-program (for the first year)

Robustness issues for Time Petri nets with distributed clocks: (S. Akshay, L. Hélouët, N.K.Kumar, and students). We want to address the timed robustness problem for Petri nets in a context where time measurement may differ for distinct transitions of the considered net. This is an adequate model to represent distribution of tasks on a network of machines. We would like to check whether a net preserves the set of reachable markings, its untimed language, or some logical properties. We plan a visit to Chennai in 2013 to work on these issues.

Well-structure of Session systems over infinite sets of agents: (P. Darondeau, L. Hélouët, M.Mukund, and students). We plan to continue the work of session-based models started in [J-DHM11], to extend this model to an infinite number of agents, and to allow verification techniques for simple logical properties such as conflicts of interest or the Chinese wall property(avoidance of conflict of interest between situations occurring at different times). M. Mukund is expected to come in Rennes early 2013 to work on this topic.

Clustering of large stochastic systems: (S.Akshay, N. Bertrand, B. Genest, K.Lodaya, P.S. Thiagaran and students). The issue we will tackle in the first year is to find a clustering of the large distributed system which makes sense. It means that this clustering needs to ensure that only a *few number* of instances in one cluster directly depends of instance in another cluster (this is a usual decomposition algorithm in qualitative verification), but also that quantitatively, these dependencies we will neglect are weak (this is the new concept). This should be achievable by computing several times the probability, using these dependencies or not, and comparing the result to see whether they influence drastically the result or not. Experiments will be performed in order to assess the accuracy gained in that way.

Logic and control: (G. Aucher, S. Pinchinat, R. Ramanujam, F. Schwarzentruber and students). We plan to reformulate the standard problems of supervisory control theory in terms of standard decision problems of suitable logical framework(s) and in terms of winning conditions of suitable game(s). Then, we plan to investigate how these two approaches can be combined and mapped one to the other, and possibly be integrated. S. Pinchinat and J. Ramanujam will focus in particular on the game-theoretical approach, whereas G. Aucher and F. Schwarzentruber will focus on the logical approach. The overall objective of the year is to merge these two lines of research.

3. Budget

Planned expenses:

Visits: We plan bilateral visits on each research topic. The total amount devoted to visits should hence be 2000 * 4*2 = 16 000 euros.

Internships: We also plan to offer one or two internships for young Indian researchers, and similarly to send two of our young researchers for an internship in India. Each internship should cost around 2000 euros (travel + accommodation)

Dissemination: We expect some publications: topics R2 and R3 are mature enough to foresee publications in 2013. Achieving two publications with our partners seems a reasonable objective. The total cost to present common work is estimated to 1500 + 2000 euros (one presentation in Europe, one in a non-european country).

Expense	Cost	Nb	Total
Visits Rennes-> Partner	2000	4	8000
Visits Partner -> Rennes	2000	4	8000
Internship Rennes	2000	2	4000
Internship India	2000	2	4000
Conference (Europe)	1500	1	1500
Conference (Other countries)	2000	1	2000
Total			27 500
Funding asked from the EA program			20 000

Co-Funding:

LIA Informel: Researchers from Rennes (N. Bertrand, L.Hélouët) are invited in the CNRS LIA, which is an international collaboration between laboratories in Chennai and French laboratories. This LIA can fund several visits each year, from both sides. Researchers in Chennai (CMI & IMSC) benefit from the same funding by the LIA, and will use it to fund visits to France.

ANR IMPRO: The ANR Impro is a French collaboration between LSV at ENS Cachan, IRCCYN in Nantes, and IRISA/INRIA Rennes. The main objectives of this ANR is to consider robustness of models, i.e. how architectural constraints (imprecision of clocks, distribution, scheduling,…) may change properties of a model. The ANR IMPRO can fund participations to conferences and missions. IMPRO will last until 2014.

INRIA/ Rennes 1 University Chair: G. Aucher owns a chair, which gives him funds to attend conferences and for visiting partners. These funds can be used to organize visits to India.

Ministry of Educational Faculty Research : P.S. Thiagarajan received a grant of 62000 SD (39000 euros) from the Ministry of Educational Faculty Research in Singapore to fund studies on “Approximate Analysis of Networks of Dynamical Systems.

Potential sources of funding:

SQUALE : B.Genest, N. Bertrand, E. Fabre, L.Hélouët, and several partners in NUS and In national university of Vietnam in Hanoi applied for a research project funded by the STIC Asie program. The outcome of this application is expected in October. If accepted, it will bring 20 000 euros in 2013 and 2014.

The French teams in Rennes plan to apply for grants from several mobility programs to fund students' travels and internships. Among others, we can already mention: MESR (grants for international mobility and internships), Ulysse (grant from the Conseil regional de Bretagne for internships), Programme Cap Monde (Conseil Général d'Ille-et-Vilaine, Internships), Fondation Rennes 1(Internships). CMI and IMSC will apply for similar opportunities within their institutes.

4. Added value

One of the major outcomes of the DST and CASDS associated teams (2006-2011) was the establishment of long-term collaborations with high-level institutes in India. These past collaborations had a considerable impact on the visibility of research activities in the Distribcom team, and furthermore allowed for the recruitment of excellent post doctoral researchers.

Maintaining this high-level and fruitful collaboration with India and Singapore is a priority of teams in Rennes. Considering the success of past collaborations, we expect similar results in terms of publications and recruitments. Partner researchers in NUS, CMI and IMSC are renowned professors, and collaborating with them has a positive impact on the visibility of researchers in Rennes. Furthermore, such collaboration offers possibilities for young researcher to visit renowned places. Similarly, this collaboration allows us to advertise post-doc positions and internships for highly skilled students in India and Singapore.

In addition to the international collaboration, Distribcom, S4 and Vertecs at INRIA Rennes are restructuring to **build a new team** on the themes of quantitative verification, distributed systems, Web Services. Working together on these topics is an opportunity to progress on existing common research, but also to find new challenges to explore for the new team.

5. Other remarks

The collaboration between Rennes, NUS, CMI and IMSC is very lively. It is now an international network exchanging ideas, students, organizing visits and conferences. We see these collaborations as long term ones. In 2011, the consortium asked for an extension of the DST associated team, which was not accepted. Nevertheless, we found funds to keep these collaborations alive (with the help of the LIA Informel, and of the International relations at INRIA). Several joint papers were published in 2012 (see list of joint publications).

Beyond the common publications, the outcome of the former collaborations led to several exchanges of researchers, PhD students, Post docs. In the future, we want to continue with similar dynamics, and apply for joint projects, for instance by contributing to larger actions via the CEFIPRA.

Involvement of researchers

The table below summarizes the implication of participants to each topic. For each topic, we list participants which are the more likely to contribute, but we of course expect involvement of other researchers.

Topic \ Place	Rennes	Singapore	Chennai
R1	d1, d4	n2	c2, c4,c5
R2	d1, d4, s1		c1, c6
R3	d4,v1, v2	n1,n2,n3,n4	i1,i4
R4	d2,d3,s2,s3		i2,i3

6. References

6.1 Joint publications of the partners

Journals:

[J-AGHY12a] S. Akshay, B. Genest, L. Hélouet, S. Yang. Regular Set of Representatives for Time-Constrained MSC Graphs. IPL, Volume 112, Issues 14–15, 2012.

[J-PABGT12] S. Palaniappan, S. Akshay, L. Bing, B. Genest, P.S. Thiagarajan. A Hybrid Factored Frontier Algorithm for Dynamic Bayesian Networks with a Biopathways Application. IEEE/ACM Trans. Comput. Biology Bioinform (TCBB). 9(5), pages 1352-1365, 2012.

[J-DGTY] P. Darondeau, B. Genest, P.S. Thiagarajan, S. Yang. Quasi-Static Scheduling of Communicating Tasks. Information and Computation 208(10), pages 1154–1168, Elsevier, 2010.

[J-GGHTY] Thomas Gazagnaire, Blaise Genest, Loïc Hélouet, P.S. Thiagarajan, Shaofa Yang. Causal Message Sequence Charts. TCS 410(41), pages 4094-4110, Elsevier, 2009.

[J-YHG08] Shaofa Yang, Loïc Hélouët, Thomas Gazagnaire. Logic-based diagnosis for distributed systems, Perspectives in concurrency, P.S. Thiagarajan's Festchrift, 2008.

Conferences:

[J-AAGT12] M. Agrawal, S. Akshay, Blaise Genest, P. S. Thiagarajan: Approximate Verification of the Symbolic Dynamics of Markov Chains. IEEE/ACM LICS 2012, pages 55-64, 2012.

[J-AGHY12b] S. Akshay, B. Genest, L. Hélouët, S. Yang: *Symbolically Bounding the Drift in Time-Constrained MSC Graphs.* Theoretical Aspects of Computing (ICTAC) 2012, LNCS no 7521, pages 1-15, 2012.

[J BG11] Nathalie Bertrand, Blaise Genest. Minimal Disclosure in Partially Observable Markov Decision Processes. FSTTCS 2011, P 411-422, LIPIcs, 2011.

[J BGG09] N. Bertrand, B. Genest, H. Gimbert. Qualitative Determinacy and Decidability of Stochastic Games with Signals. LICS 2009, pages 319-328, IEEE 978-0-7695-3746-7, 2009.

[J-PAGT12] S. Palaniappan, S. Akshay, Blaise Genest, P.S. Thiagarajan. A Hybrid Factored Frontier Algorithm for Dynamic Bayesian Networks. IEEE/ACM CMSB 2011, pages 35-44, 2011.

[J-DGTY08] P. Darondeau, B. Genest, P.S. Thiagarajan, S. Yang : Quasi-Static Scheduling of Communicating Tasks. CONCUR 2008, LNCS no 5201, pages 310-324, 2008.

[J-DHM11] P. Darondeau, L. Hélouët, M. Mukund. *Assembling Sessions*, In Automated Technology for Verification and Analysis (ATVA), LNCS no 6996, Pages 259-274, Taiwan, 2011.

[J-GGHTY07] T. Gazagnaire, B. Genest, L. Hélouet, P.S. Thiagarajan, S. Yang. Causal Message Sequence Charts. CONCUR 2007, LNCS no 4703, pages 166-180, 2007.

6.2 Main publications of the participants relevant to the project

INRIA Rennes -Distribcom

[E-AHJLR12] S.AKshay, L. Hélouët, C. Jard, D. Lime, O.H.Roux, *Robustness of Time Petri Nets under Architectural Constraints*, in FORMATS 2012, , LNCS no 7595, pages 11-26, 2012.

[E-AHJR12] S. Akshay, L. Hélouët, C. Jard and P.A. Reynier, *Robustness of Time Petri Nets under Guard Enlargement*, in RP 2012, LNCS no 7550, pages 92-107, 2012.

[E-AMS12] G. Aucher, B. Maubert, F. Schwarzentruber: Generalized DEL-Sequents. JELIA 2012, , 13th European Conference on Logics in Artificial Intelligence, LNCS no 7519, pages 54-66, 2012.

[E-A11] G. Aucher: DEL-sequents for progression. Journal of Applied Non-Classical Logics 21(3-4), pages 289-321, 2011.

[E-LS11] Lorini, François Schwarzentruber: A logic for reasoning about counterfactual emotions. Artif. Intell. 175(3-4), pages 814-847, 2011.

INRIA Rennes – s4

[E-MP12] B. Maubert and S. Pinchinat. Uniform strategies. LOFT10, 10th Conference on Logic and the Foundations of Game and Decision Theory, 2012.

[E-RP03] S. Riedweg and S. Pinchinat. Quantified Mu-Calculus for Control Synthesis. MFCS2003, 28th International Symposium on Mathematical Foundations of Computer Science, LNCS no 2747, pages 642-651,2003.

[E-MPB11] B. Maubert, S. Pinchinat and L. Bozzelli. Opacity Issues in Games with Imperfect Information. Gandalf2011, 2nd International Symposium on Games, Automata, Logics and Formal Verification, EPTCS no 54, pages 87-101, 2011.

INRIA Rennes - Vertecs

[E-BGB12] C. Baier, M. Größer, N. Bertrand: Probabilistic ω-automata. Journal of the ACM 59(1): 1, 2012.

[E-BFS12] N. Bertrand, J. Fearnley, S. Schewe: Bounded Satisfiability for PCTL. CSL 2012, LIPIcs 16, pages 92-106, 2012.

IMSC

[E-PR11] S. Paul, R. Ramanujam: Neighbourhood structure in large games, 13th Conference on Theoretical Aspects of Rationality and Knowledge (TARK 2011), pages 121-130, 2011.

[E-GRS10] S. Ghosh, R. Ramanujam, S. E. Simon: Playing Extensive Form Games in Parallel. CLIMA 2010, LNCS no 6245, pages 153-170, 2010.

[E-PL11] M. Praveen, K. Lodaya: Parameterized Complexity Results for 1-safe Petri Nets. CONCUR 2011, LNCS no 6901, pages 358-372, 2011.

CMI

[E-AGMK08] S. Akshay, B. Bollig, P. Gastin, M. Mukund, K. N. Kumar: Distributed Timed Automata with Independently Evolving Clocks. CONCUR 2008, pages 82-97, 2008.

[E-AGMK10] S. Akshay, P. Gastin, M. Mukund, K. N. Kumar: Model checking time-constrained scenario-based specifications. FSTTCS 2010, pages 204-215, 2010.

[E-MS97] M. Mukund, M. A. Sohoni: Keeping Track of the Latest Gossip in a Distributed System. Distributed Computing 10(3), pages 137-148, 1997.

NUS

[E-LHPCCWT12] B. Liu, A. Hagiescu, S. K. Palaniappan, B. Chattopadhyay, Z. Cui, W-F. Wong, P. S. Thiagarajan: Approximate probabilistic analysis of biopathway dynamics. Bioinformatics 28(11): pages 1508-1516, 2012.

[E-LZTHBLSHDT11] B. Liu, J. Zhang, P.Y. Tan, D. Hsu, A. M. Blom, B. Leong, S. Sethi, B. Ho, J. L. Ding, P. S. Thiagarajan: A Computational and Experimental Study of the Regulatory Mechanisms of the Complement System. PLoS Computational Biology 7(1), 2011.

[E-GMW10] B. Genest, A. Muscholl, Z. Wu. Verifying Recursive Active Documents with Positive Data Tree Rewriting. FSTTCS 2010, pages 469-480, LIPIcs, 2010.

6.3 Other references

[O-AHK98] R. Alur, T. Henzinger, and O. Kupferman. Alternating-time temporal logic. In Compositionality: The Significant Difference, pages 23-60, 1998.

[O-BK98] X.Boyen and D.Koller: Tractable inference for complex stochastic processes. in Uncertainty in Artificial Intelligence (UAI-98), pages 33-42, 1998.

[O-BMS11] P. Bouyer, N. Markey, O. Sankur: Robust Model-Checking of Timed Automata via Pumping in Channel Machines. FORMATS 2011, LNCS no 6919, pages 97-112, 2011.

[O-CL08] C.G. Cassandras and S. Lafortune. Introduction to discrete event systems. Springer, 2008.

[O-FS01] A. Finkel, Ph. Schnoebelen: Well-structured transition systems everywhere! Theoretical Computer Science no 256(1-2), pages 63-92, 2001.

[O-GDE11] D. P. Guelev, C. Dima, C. Enea: An alternating-time temporal logic with knowledge, perfect recall and past: axiomatisation and model-checking. Journal of Applied Non-Classical Logics 21(1), pages 93-131, 2011.

[O-HJ94] H. Hansson and B. Jonsson: A logic for reasoning about time and reliability. In Formal Aspects of Computing, 6(5), pages 512-535, 1994.

[O-HM90] J. Halpern and Y. Moses. Knowledge and common knowledge in a distributed environment. Journal of the ACM, 37(3), pages 549-587, 1990.

[O-HS08] A. Herzig, F. Schwarzentruber: Properties of logics of individual and group agency. Advances in Modal Logic 2008, pages 133-149, 2008.

[O-HW02] W. Van Der Hoek and M. Wooldridge. Tractable multiagent planning for epistemic goals. In Proceedings of the first international joint conference on Autonomous agents and multiagent systems: part 3, pages 1167-1174. ACM, 2002.

[O-HYC08] K. Honda, N. Yoshida, and M. Carbone. Multiparty Asynchronous Session Types. In POPL, pages 273–284, 2008.

[O-KVAK10] V. Korthikanti, M. Vishwanathan, G. Agha and Y. Kwon: Reasoning about MDPs as Transformers of probability distributions. QEST'10, pages 199-208, 2010.

[O-MHB11] B. Masson, L. Hélouët, A. Benveniste: Compatibility of Data-Centric Web Services. WS-FM 2011, LNCS no 7176, pages 32-47, 2011.

[O-P00] A. Puri : Dynamical properties of timed systems. Discrete Event Dynamic Systems 10(1-2), pages 87–113, 2000.

[O-RW89] P.J.G. Ramadge and W.M. Wonham. The control of discrete event systems. Proceedings of the IEEE, 77(1), pages 81-98, 1989.

[O-MW01] K.Murphy and Y.Weiss:``The factored frontier algorithm for approximate inference in DBN", in Uncertainty in Artificial Intelligence (UAI 01), pages 378-385, 2001.

7. Detailed scientific program

7.1 Context

This project lays in the general context of formal modeling, and analysis of behaviors of distributed systems. We want to address verification and supervision of distributed systems through formal modeling and automated reasoning on models. By distributed systems, we mean software architectures made of several independent communicating entities. In the 90's the kind of system addressed was mainly telecommunication protocols. Nowadays, distributed systems are frequently web-based systems such as Web Services, but several aspects of distributed systems can be found in biological applications. Within this context, there are several key challenges.

The first one is to ensure that models considered are **realistic**. While modeling can be an abstraction of the represented objects allowing automated reasoning, one should ensure that the distance between a model and reality does not affect important properties, and hence the usefulness of automated verification. It is then important to focus on actual distributed systems (such as Web Services), but also to consider realism of models, i.e. address **robustness** issues to see if a model complies with architectural constraints of real world systems.

A second challenge is to find **quantitative** rather than Boolean answers to formal properties: knowing the probability of some set of paths and the most probable path is usually much more informative than knowing the set of possible path satisfying some property.

The third challenge addressed by this project is the unification of **control** theories for distributed systems: many approaches have been proposed since the original work of Ramage & Wonham [O-RW89]. A challenge in this project is to consider control from several points of view, and to unify existing approaches.

Realistic models

Adressing distributed systems in a formal way means designing adequate models, which can represent ***realistic systems***. A first issue that we want to consider is the modeling of Web-services. Considering that a lot of effort is now spent on building service based architectures, we feel that this category of systems is a first class choice for our studies. Furthermore, Web services have several particularities (use of workflows, absence of ordering in requests,…) that calls for the use of variants of Petri nets, where we have good expertise, and of Well-structured Transition systems [O-FS01]. Hence, a particular challenge is to provide efficient models for Web-services, that is models that are a good tradeoff between expressiveness and decidability. Several solutions were proposed by our teams [O-MHB11, J-DHM11], and should be extended. A lot of efforts is also devoted to formal modeling of services in the pi-calculus community [O-HYC08]. Ensuring realism of models also means ensuring that formal analysis of a model makes sense for the real system that is modeled. For this reason, we want to consider **robustness** of models with respect to architectural constraints imposed by an implementation. In particular, we consider **timed robustness**. In many models (for instance timed automata), the representation of time is idealized: tasks are launched at precise dates, clocks never drift,… Such assumptions cannot be implemented: even very accurate hardware has some imprecision, more especially in a distributed context where components of the system do not share a common global clock. Robust verification addresses the question of whether some property satisfied by a model still holds under the (realistic) assumption that the system is implemented on architectures with imprecise clocks. This problem has been addressed for models such as timed automata [O-Puri00, O-BMS11], but are only in their infancy for true concurrency models, where differences between a perfect semantics and a semantics under imprecision may arise from concurrency (a novelty w.r.t. timed automata). We plan to study **robustness issues for true concurrency models**, that is integrate into verification processes the fact that actions are located on distinct processes with distinct and imprecise clocks. We have started considering robustness issues for variants of Petri Nets [E-AHJLR12, E-AHJR12], and we want to build on this experience.

Quantitative verification

We focus on qualitative (« almost surely, a call to a service is successful") or quantitative ("the average failure rate is lower than 0.01").One possibility to obtain the probability is to compute its exact value. Such questions have answers for markovian models (Markov Chains, Markov Decision Process subsuming finite state systems) and some quantitative logic (PCTL mainly) [O-HJ94]. However, when dealing with very distributed systems (that can be parametric) and/or other logics [O-KVAK10], computing the exact probability may be computationaly infeasible. Dynamic Bayesian Networks for instance allows representing compactly very distributed systems. Approximated inference of the probabilities is a pragmatic solution in that cases [O-BK98,O-MW01].

A Unified Control Theory of Distributed Systems with Communication :

The theory of supervisory control deals with problems related to the existence and the synthesis of supervisors, whose role is to control the behavior of a discrete event system so as to produce a specified behavior [O-RW89,O-CL08]. These problems are addressed under various assumptions like partial control or partial observation of the events, or decentralization of the supervisor. More recently, this theory has started to consider distributed control with communication. In that case, local supervisors can interact, send and receive information from other supervisors and they need to make local decisions without resorting to a central authority gathering all the information from the local supervisors. However, with the current methods that are used, the formalism tends to be quite complex and it is difficult to derive ***automatically*** algorithms that solve the problems of the theory of supervisory control.

Problems of theory of supervisory control connected with problems in other fields:
In parallel, it turns out that several other research fields like distributed artificial intelligence, game theory and recent developments in logic deal with the same kind of situation: a group of agents (alias local supervisors) interact, send and receive information from other agents (alias supervisors) and they need to make local decisions without resorting to a central agent (authority). Independently from the community of supervisory control theory, numerous researchers from these other research fields already gather regularly to address problems dealing with this kind of situations but from a different perspective in conferences such as TARK,LORI,AAMAS,LOFT…. These related research fields traditionally use different methods. For example, several logics like ATL [O-AHK98], ATEL [O-HW02], etc. are the result of the interaction between logicians and game-theorists. They provide formal systems to express perfect information and imperfect information game properties. Also, they provide algorithmic methods to reason ***automatically*** about those properties.

Fusion of the game-logic worlds and theory of supervisory control:
These related fields often cover a larger spectrum of interactive situations and phenomena than the ones usually considered in supervisory control theory. While the study of epistemic reasoning in distributed computing has led to a nuanced understanding of how communication mechanisms enrich and limit co-ordination, incorporating goal orientedness is challenging; on the other hand, game theoretic methods are rich in goal-orientation but communication tends to be primitive. Combining these methods is likely to enrich both paradigms.

2.2 Detailed Objectives for the three years

R1 : Our main objective in this research direction is to provide a clear picture of robustness issues for true concurrency models in a context where each process has its own measurement of time. We intend to start this study with timed variants of Petri nets. This is a challenging task, as we already know that many robustness problems are undecidable for this model [E-AHJLR12,E-AHJR12], and that automata with independently evolving clocks [E-AGMK08] also raise undecidable issues. CMI has a lot of experience on timed models [E-AGMK10, E-AGMK08], and IRISA on time Petri nets [E-AHJLR12, E-AHJR12] and time-constrained scenarios [J-AGHY12a,J-AGHY12b]. A first step is to formalize properly a model of time(d) Petri net with independent processes and the robustness problems we want to address on this model. Then, as we expect any reasonably expressive model to be undecidable, we plan to find restrictions ensuring decision or semi-decision procedures for robustness issues. The questions usually addressed are whether the set of possible configurations of a system is preserved assuming a "realistic" rather than the idealized one, and similar questions on

preservation of languages, of a set of ω-regular properties, etc.

R2: Topic R2 is already a well established research topic. We have already proposed a formal model for data-centric workflows [E-BHM11], and a session+role model for Web Services [J-DHM11] that describes how a finite number of agents can run an arbitrary number of concurrent transactions. For this session model, coverability of some (bad) configuration is decidable, but the model allows only for a finite set of agents. We have two major objectives: the first is to extend our models and decision procedures to be able to describe systems with arbitrary numbers of agents. The tools that will be used to this extent are well-structured transition systems (WSTS) [O-FS01], or variants of Petri nets. Showing that a model is well-structured allows for straightforward decidability of several interesting properties such as coverability of some configuration, which is often enough to prove safety of a system. The second objective is to be able to verify more elaborated properties than only coverability. Our service models are well adapted to the description of transactional systems, for which we would like to detect **conflicts of interest** between agents, or a more elaborated property called the "**Chinese Wall**". This property forbids an agent to take part in a transaction if this is conflicting with some of its former collaborations. Overall, we wish to propose a highly expressive model together with a decidable logic to reason on this model. The techniques used to reach this goal will build on our knowledge of WSTS, on extensions of Petri nets (WSTS are often seen as particular kind of net). The last objective, which is more exploratory, is to propose monitoring techniques for session systems. The problem can be stated as follows: given a model M of a web-based transaction system, an implementation *I* of this model, and a property to monitor, can we instrument I with a set of observers (synthesized from M) that raise an alarm when they are sure that the property is violated. This framework should not raise wrong alarms. There has been a lot of research on this topic using pi-calculus [O-HYC08], but it is not yet clear whether the proposed solutions can be transferred to our setting.

R3 : In the next three years, we will:

- Develop algorithms to compute precisely probabilities of logical properties, in particular in the presence of imperfect information and/or time. We will build on the work already done in [J-BG11] to achieve this goal.
- Improve the precision of approximated inference algorithms for distributed (parametrized or not) systems, and deduce formal bounds that guarantee the probability to be in an interval of bounded size. For that, we will develop the techniques we introduced in [J-PAGT11]. In particular, we will provide a decomposition algorithm such that the global approximated probability will be more accurate (through a better accuracy on each component) than by considering the system as a monolith. This has been a major challenge in analysis of distributed system which cannot be in general reached exactly. However, approximation gives more freedom for clustering.
- Concerning logics different from PCTL, we will leverage on [J-AAGT12] for the development of approximated verification.

R4: Our goal in this project is to develop a unified theoretical framework for supervisory control theory, in line with Halpern and Moses, who hoped in their seminal paper [O-HM90] that "a theory of knowledge, communication, and action will prove rich enough to provide general foundations for a unified theoretical treatment of distributed systems". To that aim, we will investigate to which extent techniques from epistemic reasoning and game theory can be applied to address problems of supervisory control theory for distributed systems. We will have a number of milestones: The first milestone will be to reformulate in logical and game-theoretical terms the core problems of supervisory control theory. In that respect, epistemic logic should help to handle partial observation; note that imperfect information games do provide natural models. The second milestone will consist in bringing together epistemic logic and imperfect information games to handle individual (i.e. subsystems) knowledge. This task may take place in parallel with the first task. It is a challenging task because, taking apart control theory issues, the logical foundation of games with imperfect information is an emerging research field with only few results [O-GDE11, E-MPB11]. The third milestone will consist in incrementing the previous framework by considering communication mechanisms between the subsystems. In game theory, communication between players is very primitive, whereas in (dynamic) epistemic logic, there are powerful rigorous ways to model effects of atomic communication events on the individual knowledge. It is a challenging task to transfer this apparatus to games and will probably lead to genuine new results in game theory but more importantly, in distributed control: indeed, formal approaches for communication in distributed control are rather recent and acknowledged as a difficult and complex question by the community of supervisory control. The fourth milestone will consist in studying properties of the developed unified framework, both computational and in terms of expressiveness. For this, we may link the new framework with existing logical formalisms and/or game-based settings.