IN PARTNERSHIP WITH:
CNRS

Ecole Polytechnique

Activity Report 2011

Project-Team COMETE

Concurrency, Mobility and Transactions

IN COLLABORATION WITH: Laboratoire d'informatique de l'école polytechnique (LIX)

RESEARCH CENTER
Saclay - Île-de-France

THEME
Programs, Verification and Proofs

Table of contents

Project-Team COMETE

Keywords: Concurrency, Information Theory, Quantitative Information Flow, Privacy, Model-Checking

1. Members

Research Scientists

Catuscia Palamidessi [Team Leader, Senior Researcher, HdR]
Frank Valencia [Junior Researcher]
Konstantinos Chatzikokolakis [Junior Researcher]

External Collaborator

Romain Beauxis [Univ. of Tulane, USA. He has been a PhD student in Comète and defended his thesis in 2009]

PhD Students

Andrés Aristizábal [Grant DGA/CNRS. Since 1/10/2009]
Mário Sergio Ferreira Alvim Junior [Grant DGA/CNRS. 1/10/2008 – 30/9/2011]
Nicolás Bordenabe [Grant INRIA/DGA. Since 1/10/2011]
Ivan Gazeau [Grant ANR (CCP). Co-supervised by Dale Miller, INRIA. Since 1/10/2009]
Sophia Knight [Grant INRIA-CORDIS. Since 15/9/2010]
Luis Fernando Pino Duque [Grant INRIA/DGA. Since 1/10/2011]
Lili Xu [Co-supervised by Huimin Li, Chinese Academy of Science, Bejijing. Since 15/10/2011]

Post-Doctoral Fellows

Miguel Andrés [Grant QUALCOMM. Since 27/11/2010]
Jérémy Dubreil [Grant INRIA. 1/12/2009–31/3/2011]
Marco Giunti [Grant ERCIM. Since 1/3/2011]
Ehab ElSalamouni [Grant INRIA. Since 1/10/2011]
Sardaouna Hamadou [Grant ANR (Panda). Since 1/11/2011]

Visiting Scientist

Geoffrey Smith [Professor at Florida International University, USA. He visited comète from 26/8/2011 until 23//12/2011]

Administrative Assistant

Christelle Lievin [SAR]

2. Overall Objectives

2.1. Introduction

Our times are characterized by the massive presence of highly distributed and mobile systems consisting of diverse and specialized devices, forming heterogeneous networks, and providing different services and applications. The resulting computational systems are usually referred to as *Ubiquitous Computing*, (see, e.g., the UK Grand Challenge initiative under the name *Sciences for Global Ubiquitous Computing* [35]). *Security* is one of the fundamental concerns that arises in this setting. The problem of *privacy*, in particular, is exacerbated by orders of magnitude: The frequent interaction between users and electronic devices, and the continuous connection between these devices and the internet, offer to malicious agents the opportunity to gather and store huge amount of information, often without the individual being even aware of it. Mobility is also an additional source of vulnerability, since tracing may reveal significant information. To avoid these hazards, honest agents should use special protocols, called *security protocols*.

The systems above are usually very complex and based on impressive engineering technologies, but they do not always exhibit a satisfactory level of robustness and reliability. The same holds for security protocols: they usually look simple, but the properties that they are supposed to ensure are extremely subtle, and it is also difficult to capture the capabilities of the attacker. As a consequence, even protocols that seem at first "obviously correct" are later (often years later) found to be prone to attacks.

In order to overcome these drawbacks, we need to develop formalisms, reasoning techniques, and tools, to specify systems and protocols, their intended properties, and to guarantee that these intended properties are indeed satisfied. The challenges that we envisage are (a) to find suitably expressive formalisms which capture essential new features such as mobility, probabilistic behavior, presence of uncertain information, and potentially hostile environment, (b) to build suitably representative models in which to interpret these formalisms, and (c) to design efficient tools to perform the verification in presence of these new features.

2.2. Highlights

+ Catuscia Palamidessi has been keynote speaker at the 2011 edition of the conference ICALP (Internetional Colloquium on Automata, Languages and Programming, http://icalp11.inf.ethz.ch/).

3. Scientific Foundations

3.1. Probability and information theory

Participants: Mário Alvim, Miguel Andrés, Nicolás Bordenabe, Konstantinos Chatzikokolakis, Catuscia Palamidessi.

Much of the research of Cométe focuses on security and privacy. In particular, we are interested in the problem of the leakage of secret information through public observables.

Ideally we would like systems to be completely secure, but in practice this goal is often impossible to achieve. Therefore, we need to reason about the amount of information leaked, and the utility that it can have for the adversary, i.e. the probability that the adversary be able to exploit such information.

The recent tendency is to use information theoretic approach to model the problem and define the leakage in a quantitative way. The idea is to consider that system as an information-theoretic *channel*. The input represents the secret, the output represents the observable, and the correlation between the input and output (*mutual information*) represents the information leakage.

Information theory depends on the notion of entropy. Most of the proposals in the literature use *Shannon entropy*, which is the most established measure of uncertainty. From the security point of view, this measure corresponds to a particular model of attack and a particular way of estimating the security threat (vulnerability of the secret). We consider also other notions, in particular the Rényi min-entropy, which seem to be more appropriate for security in common scenarios like the one-try attacks.

3.2. The probabilistic asynchronous π-calculus

Participants: Konstantinos Chatzikokolakis, Marco Giunti, Catuscia Palamidessi, Frank Valencia, Lili Xu.

We will focus our efforts on a probabilistic variant of the asynchronous π-calculus, which is a formalism designed for mobile and distributed computation. A characteristic of our calculus is the presence of both probabilistic and nondeterministic aspects. This combination is essential to represent probabilistic algorithms and protocols, and express their properties in presence of unpredictable (nondeterministic) users and adversaries.

3.3. Expressiveness issues

Participants: Andrés Aristizábal, Catuscia Palamidessi, Luis Fernando Pino Duque, Frank Valencia.

We intend to study models and languages for concurrent, probabilistic and mobile systems, with a particular attention to expressiveness issues. We aim at developing criteria to assess the expressive power of a model or formalism in a distributed setting, to compare existing models and formalisms, and to define new ones according to an intended level of expressiveness, taking also into account the issue of (efficient) implementability.

3.4. Concurrent constraint programming

Participants: Andrés Aristizábal, Sophia Knight, Luis Fernando Pino Duque, Frank Valencia.

Concurrent constraint programming (ccp) is a well-established process calculus [39] for modeling systems where agents interact by adding and asking information in a global store. This information is represented as first-order logic formulae, called constraints, on the shared variables of the system (e.g., $X > 42$). The most distinctive and appealing feature of ccp is perhaps that it unifies in a single formalism the operational view of processes based upon process calculi with a declarative one based upon first-order logic. It also has an elegant denotational semantics that interprets processes as closure operators (over the set of constraints ordered by entailment). In other words, any ccp process can be seen as an idempotent, increasing, and monotonic function from stores to stores. Consequently, ccp processes can be viewed at the same time as computing agents, formulae in the underlying logic, and closure operators. This allows ccp to benefit from the large body of techniques of process calculi, logic and domain theory.

Our research in ccp develops along the following two lines:

1. The study of a bisimulation semantics for ccp. The advantage of bisimulation, over other kinds of semantics, is that it can be efficiently verified.
2. Enriching ccp with epistemic constructs, which will allow to reason about the knowledge of agents.

3.5. Model checking

Participants: Miguel Andrés, Catuscia Palamidessi.

We plan to develop model-checking techniques and tools for verifying properties of systems and protocols specified in the above formalisms.

Model checking addresses the problem of establishing whether the model (for instance, a finite-state machine) of a certain specification satisfies a certain logical formula.

We intend to concentrate our efforts on aspects that are fundamental for the verification of security protocols, and that are not properly considered in existing tools. Namely, we will focus on:

(a) the combination of probability and mobility, which is not provided by any of the current model checkers,

(b) the interplay between nondeterminism and probability, which in security present subtleties that cannot be handled with the traditional notion of scheduler,

(c) the development of a logic for expressing security (in particular privacy) properties.

Concerning the last point (the logic), we should capture both probabilistic and epistemological aspects, the latter being necessary for treating the knowledge of the adversary.

Logics of this kind have been already developed, but the investigation of the relation with the models coming from process calculi, and their utilization in model checking, is still in its infancy.

4. Application Domains

4.1. Security and privacy

Participants: Mário Sergio Ferreira Alvim Junior, Miguel Andrés, Nicolás Bordenabe, Konstantinos Chartzikokolakis, Jérémy Dubreil, Catuscia Palamidessi.

The aim of our research is the specification and verification of protocols used in mobile distributed systems, in particular security protocols. We are especially interested in protocols for *information hiding*.

Information hiding is a generic term which we use here to refer to the problem of preventing the disclosure of information which is supposed to be secret or confidential. The most prominent research areas which are concerned with this problem are those of *secure information flow* and of *privacy*.

Secure information flow refers to the problem of avoiding the so-called *propagation* of secret data due to their processing. It was initially considered as related to software, and the research focussed on type systems and other kind of static analysis to prevent dangerous operations, Nowadays the setting is more general, and a large part of the research effort is directed towards the investigation of probabilistic scenarios and treaths.

Privacy denotes the issue of preventing certain information to become publicly known. It may refer to the protection of *private data* (credit card number, personal info etc.), of the agent's identity (*anonymity*), of the link between information and user (*unlinkability*), of its activities (*unobservability*), and of its *mobility* (*untraceability)*.

The common denominator of this class of problems is that an adversary can try to infer the private information (*secrets*) from the information that he can access (*observables*). The solution is then to obfuscate the link between secrets and observables as much as possible, and often the use randomization, i.e. the introduction of *noise*, can help to achieve this purpose. The system can then be seen as a *noisy channel*, in the information-theoretic sense, between the secrets and the observables.

We intend to explore the rich set of concepts and techniques in the fields of information theory and hypothesis testing to establish the foundations of quantitive information flow and of privacy, and to develop heuristics and methods to improve mechanisms for the protection of secret information. Our approach will be based on the specification of protocols in the probabilistic asynchronous π-calculus, and the application of model-checking to compute the matrices associated to the corresponding channels.

5. Software

5.1. A model checker for the probabilistic asynchronous π-calculus

Participants: Miguel Andrés [correspondant], Catuscia Palamidessi.

In collaborations with Dave Parker and Marta Kwiatkowska, we are developing a model checker for the probabilistic asynchronous π-calculus. Case studies with Fair Exchange and MUTE, an anonymous peer-to-peer file sharing system, are in progress.

Technically we use MMC as a compiler to encode the probabilistic π-calculus into certain PRISM representation, which will then be verified against PCTL using PRISM. The transitional semantics defined in MMC can be reused to derive the symbolic transition graphs of a probabilistic process. The code for derivation will work as an add-on to MMC under XSB and invoke a graph traversal to enumerate all reachable nodes and transitions of the probabilistic process.

In the meanwhile we are also attempting a direct and more flexible approach to the development of a model checker for the probabilistic π-calculus, using OCaml. This should allow to extend the language more easily, to include cryptographic primitives and other features useful for the specification of security protocols. As the result of our preliminary steps in this direction we have developed a rudimentary model checker, available at the following URL: http://vamp.gforge.inria.fr/.

5.2. PRISM model generator

Participants: Konstantinos Chatzikokolakis [correspondant], Catuscia Palamidessi.

This software generates PRISM models for the Dining Cryptographers and Crowds protocols. It can also use PRISM to calculate the capacity of the corresponding channels. More information can be found in [33] and in the file README file width instructions at the URL http://www.lix.polytechnique.fr/comete/software/README-anonmodels.html.

The software can be download at http://www.lix.polytechnique.fr/comete/software/anonmodels.tar.gz. These scripts require Perl to run and have been tested in Linux. The GUI of the corners tool also requires the Perl/TK library. Finally some parts of the model generator tool require PRISM and gnuplot to be installed.

5.3. Calculating the set of corner points of a channel

Participants: Konstantinos Chatzikokolakis [correspondant], Catuscia Palamidessi.

The corner points can be used to compute the maximum probability of error and to improve the Hellman-Raviv and Santhi-Vardy bounds. More information can be found in [34] and in the file README file width instructions at the URL http://www.lix.polytechnique.fr/comete/software/README-corners.html.

The software can be download at http://www.lix.polytechnique.fr/comete/software/corners.tar.gz. These scripts require Perl to run and have been tested in Linux. The GUI of the corners tool also requires the Perl/TK library. Finally some parts of the model generator tool require PRISM and gnuplot to be installed.

5.4. MMCsp, a compiler for the π-calculus

Participants: Peng Wu [correspondant], Catuscia Palamidessi.

MMCsp is a compiler from a simple probabilistic π-calculus to PRISM (http://www.prismmodelchecker.org/manual/Main/Introduction). models. It is built on XSB (http://xsb.sourceforge.net/), a tabled logic programming system, and generates the symbolic semantic representation of a probabilistic pi-calculus term in text. A separate Java program then translates this semantic representation into a probabilistic model for PRISM.

The tool was developed by Peng Wu during his postdoc period in Comète in the context of the collaboration between the teams Comète and PRISM under the INRIA/ARC Project ProNoBib (http://www.lsv.ens-cachan.fr/~goubault/ProNobis/index.html). It is based on the papers [41] and [38].

The source code is free and can be download from http://www.cs.ucl.ac.uk/staff/p.wu/mmc_sp_manual.html.

6. New Results

6.1. Foundations of information hiding

Information hiding refers to the problem of protecting private information while performing certain tasks or interactions, and trying to avoid that an adversary can infer such information.

This is one of the main areas of research in Comète, and two PhD thesis based on this topic have been defended this year in Comète [12], [11] have been defended this year. We are exploring several topics, described below. An overview of our results is contained in [24].

6.1.1. The problem of information hiding in presence of concurrency

The analysis of probabilistic concurrent systems usually relies on the notion of scheduler in order to solve the nondeterminism. Unfortunately the classical notion of scheduler, which is a mathematical functions that chooses the next step depending on the history of the computation, can leak any secret information contained in the history. This creates false positives, and it is known as the problem of the *allmighty scheduler*. One way to solve this problem, already explored in literature, is to fix the strategy of the scheduler beforehand [31]. However this solution is considered too rigid and unrealistic. In [14] we have propose a milder restriction on the schedulers, and we have defined the notion of strong (probabilistic) information hiding under various notions of observables. Furthermore, we have proposed a method, based on the notion of automorphism, to verify that a system satisfies the property of strong information hiding, namely strong anonymity or no-interference, depending on the context.

6.1.2. Modeling the knowledge of the adversary

In [15] we have developed a game semantics for process algebra with two interacting agents. The purpose of our semantics is to make manifest the role of knowledge and information flow in the interactions between agents and to control the information available to interacting agents. We have defined games and strategies on process algebras, so that two agents interacting according to their strategies determine the execution of the process, replacing the traditional scheduler. We have shown that different restrictions on strategies represent different amounts of information being available to a scheduler. We have also shown that a certain class of strategies corresponds to the syntactic schedulers of Chatzikokolakis and Palamidessi [32], which were developed to overcome problems with traditional schedulers modeling interaction. The restrictions on these strategies have an explicit epistemic flavor.

6.1.3. Opacity

Opacity is a security property formalizing the absence of secret information leakage and we have addressed in [30] the problem of synthesizing opaque systems. A secret predicate S over the runs of a system G is opaque to an external user having partial observability over G, if s/he can never infer from the observation of a run of G that the run belongs to S. We have chosen to control the observability of events by adding a device, called a mask, between the system G and the users. We have first investigated the case of static partial observability where the set of events the user can observe is fixed a priori by a static mask. In this context, we have shown that checking whether a system is opaque is PSPACE-complete, which implies that computing an optimal static mask ensuring opacity is also a PSPACE-complete problem. Then, we have introduced *dynamic* partial observability where the set of events the user can observe changes over time and is chosen by a dynamic mask. We have shown how to check that a system is opaque with respect to a dynamic mask and we have also addressed the corresponding synthesis problem: given a system G and secret states S, compute the set of dynamic masks under which S is opaque. Our main result is that the set of such masks can be finitely represented and can be computed in EXPTIME and this is a lower bound. Finally we have also addressed the problem of computing an optimal mask.

6.1.4. Interactive systems

In [13] we have considered systems where secrets and observables can alternate during the computation. We have shown that the information-theoretic approach which interprets such systems as (simple) noisy channels is not valid anymore. However, the principle can be recovered if we consider more complicated types of channels, that in Information Theory are known as channels with memory and feedback. We have shown that there is a complete correspondence between interactive systems and such kind of channels. Furthermore, we have shown that the capacity of the channels associated to such systems is a continuous function of the Kantorovich metric.

6.1.5. Differential privacy

Differential privacy is a notion that has emerged in the community of statistical databases, as a response to the problem of protecting the privacy of the database's participants when performing statistical queries. The idea is that a randomized query satisfies differential privacy if the likelihood of obtaining a certain answer for a database x is not too different from the likelihood of obtaining the same answer on adjacent databases, i.e. databases which differ from x for only one individual.

In [17], [16], we have analyzed critically the notion of differential privacy in light of the conceptual framework provided by the Rényi min information theory. We have shown that there is a close relation between differential privacy and leakage, due to the graph symmetries induced by the adjacency relation. Furthermore, we have considered the utility of the randomized answer, which measures its expected degree of accuracy. We have focused on certain kinds of utility functions called "binary", which have a close correspondence with the Rényi min mutual information. Again, it turns out that there can be a tight correspondence between differential privacy and utility, depending on the symmetries induced by the adjacency relation and by the query. Depending on these symmetries we can also build an optimal-utility randomization mechanism while preserving the required level of differential privacy. Our main contribution is a study of the kind of structures that can be induced by the adjacency relation and the query, and how to use them to derive bounds on the leakage and achieve the optimal utility.

6.2. Concurrent constraint programming

6.2.1. Bisimilarity

Bisimilarity is one of the main representative equivalences for concurrent behaviour. It captures our intuitive notion of process equivalence; two processes are equivalent if they can match each other's moves. Furthermore, it provides an elegant co-inductive proof technique based on the notion of bisimulation. Nevertheless, there have been few attempts to define a notion of bisimilarity for *concurrent constraint programming* (ccp). The ones we were aware of are those in [40] and [36] but they are not completely satisfactory: The first one may tell apart processes with identical observable behaviour, while the second quantifies over all possible inputs from the environment, and hence it is not clear whether it can lead to a feasible proof technique.

Bisimilarity relies on *labelled transitions*: each evolution step of a system is tagged by some information aimed at capturing the possible interactions of a process with the environment. In [18] we have provided a labelled transition system for ccp and we have proposed a notion of ccp bisimilarity. Intuitively, in this transition system the labels represent the minimal information that processes require from the environment to execute. Furthermore we have shown that, unlike previous approaches, our notion of bisimilarity coincides with the standard notion of equivalence for (deterministic) ccp. This way we have provided ccp with an alternative co-inductive proof technique, coherent with previous equivalences, for process behaviour.

When the state space of a system is finite, the ordinary notion of bisimilarity can be computed via the well-known partition refinement algorithm, but unfortunately, this algorithm does not work for ccp bisimilarity. In [19] we have proposed a variation of the partition refinement algorithm for verifying ccp bisimilarity. To the best of our knowledge this is the first work providing for the automatic verification of program equivalence for ccp.

6.2.2. Modeling cellular signaling systems

The molecular mechanisms of cell communication with the environment involve many concurrent processes governing dynamically the cell function. This concurrent behavior makes traditional methods, such as differential equations, unsatisfactory as a modeling strategy since they do not scale well when a more detailed view of the system is required.

In [19] we have described a modeling strategy for cellular signaling systems based on a temporal and probabilistic extension of ccp. Starting from an abstract model, we have built refinements adding further details coming from experimentation or abstract assumptions. The advantages of our approach are: due to the notion of partial information as constraints in CCP, the model can be straightforwardly extended when more information is available; qualitative and quantitative information can be represented by means of probabilistic constructs of the language; finally, the model is a runnable specification and can be executed, thus allowing for the simulation of the system. We have outlined the use of this methodology to model the interaction of G-protein-coupled receptors with their respective G-proteins that activates signaling pathways inside the cell. Finally, we have presented simulation results obtained from an implementation of the framework

6.3. Session types

In [22] we have presented a type checking algorithm for establishing a session-based discipline in the pi calculus of Milner, Parrow and Walker [37]. Our session types are qualified as linear or unrestricted. Linearly typed communication channels are guaranteed to occur in exactly one thread, possibly multiple times; afterwards they evolve as unrestricted channels. Session protocols are described by a type constructor that denotes the two ends of one and the same communication channel. We have proved the soundness of the algorithm by showing that processes consuming all linear resources are accepted by a typing system preserving typings during the computation and that type checking is consistent w.r.t. structural congruence.

7. Partnerships and Cooperations

7.1. National Initiatives

7.1.1. ANR project PANDA: "Analyse du Parallélisme et de la Distribution"

This project is financed by the ANR, for the years 2009-2011. The partners involved are:

- EPIs Comète and Parsifal at INRIA Saclay. Responsible: Catuscia Palamidessi
- CEA Saclay. Responsible: Emmanuel Haucourt
- Pôle Parisien. Responsible: Damiano Mazza
- Pôle Méditerranéen. Responsible: Emmanuel Godard
- Airbus. Responsible: Jean Souyris.

7.1.2. ANR project CPP: Confidence, Proofs and Probabilities

This project is financed by the ANR, for the years 2009-2011. The partners involved are:

- LSV. Responsible: Jean Goubault-Larrecq
- EPIs Comète and Parsifal at INRIA Saclay. Responsible: Catuscia Palamidessi
- CEA LIST. Responsible: Olivier Bouissou
- Supelec SSE. Responsible: Gilles Fleury
- Supelec L2S. Responsible: Michel Kieffer

7.2. International Initiatives

7.2.1. DRI Equipe Associée PRINTEMPS

PRINTEMPS (PRobability and INformation Theory for Modeling Privacy and Secrecy) focuses on the applications of Information Theory to security. We are particularly interested in studying the interactions between Concurrency and Information Theory.

This project has started in January 2006 and includes the following sites:

- INRIA Futurs. Responsible: C. Palamidessi
- McGill University, Canada. Responsible: P. Panangaden

Home page: http://www.lix.polytechnique.fr/comete/Projects/Printemps/.

7.2.1.1. International Partners

- Moreno Falaschi, Dipartimento di Scienze Matematiche e Informatiche, Università di Siena, Italy.
- Camillo Rueda and Carlos Olarte, Pontificia Universidad Javeriana, Colombia.
- Geoffrey Smith, School of Computing and Information Sciences, Florida International University, USA
- Vladimiro Sassone, School of Electronics and Computer Science University of Southampton, United Kingdom.

7.3. Exterior research visitors

7.3.1. Visits of International Scientists

- Geoffrey Smith, Professor at the Florida Florida International University, USA. He visited for four months, from 26/8/2011 until 23/12/2011.

- Moreno Falaschi, professor at the Università di Siena, Italy. He visited for one month, from 1/11/2011 till 30/11/2011.
- Vladimiro Sassone, professor at the University of Southampton, United Kingdom, Italy. He visited for one month, from 1/12/2011 till 31/12/2011.

7.3.2. Internship

- Marco Stronati, master student at the Università di Pisa, Italy. He is visiting for six months, from 1/10/2011 till 31/3/2012. He is doing his master thesis under the co-supervision of Giorgio Levi (Univ. di Pisa) and Catuscia Palamidessi.
- Lili Xu, PhD student at the Academy of Science of Beijing, China. She is visiting for nine months, from 15/10/2011 until 15/7/2012. She is doing her PhD thesis under the co-supervision of Huimin Li (Ch. Academy of Science, Beijing) and Catuscia Palamidessi.

8. Dissemination

8.1. Animation of the scientific community

Note: In this section we include only the activities of the permanent internal members of Comète.

8.1.1. Editorial activity

Catuscia Palamidessi is:

- Member of the Editorial Board of the journal on Mathematical Structures in Computer Science, published by the Cambridge University Press.
- Member of the Editorial Board of the journal on Theory and Practice of Logic Programming, published by the Cambridge University Press.
- Member of the Editorial Board of the Electronic Notes of Theoretical Computer Science, Elsevier Science.
- Co-editor (with Frank Pfenning) of the special issue of Logical Methods in Computer Science dedicated to selected papers of FoSSaCS 2013.
- Co-editor (with Samson Abramsky and Michael Mislove) of the special issue of Theoretical Computer Science dedicated to selected papers of MFPS XXV.
- Co-editor (with Geoffrey Smith) of the special issue of Mathematical Structures in Computer Science dedicated to Quantitative Information Flow.

Frank D. Valencia is:

- Area editor (for the area of Concurrency) of the ALP Newsletter.
- Co-editor of the special issue of Mathematical Structures in Computer Science dedicated to the 17th International Workshop on Expressiveness in Concurrency.

Konstantinos Chatzikokolakis and Catuscia Palamidessi are:

- Co-editors (with Sebastian Mödersheim) of the special issue of the Journal of Computer Security dedicated to selected papers of TOSCA 2011 and SecCo 2011.

8.1.2. Steering Committees

Catuscia Palamidessi is member of:

- The Council of EATCS, the European Association for Theoretical Computer Science. Since 2005.
- The Steering Committee of ETAPS, the European Joint Conferences on Theory and Practice of Software. Since 2006.
- The IFIP Technical Committee 1 – Foundations of Computer Science. Since 2007.
- The IFIP Working Group 2.2 – Formal Description of Programming Concepts. Since 2001.
- The IFIP Working Group 1.7 – Theoretical Foundations of Security Analysis and Design. Since 2010.

8.1.3. Invited Talks

Catuscia Palamidessi has given invited talks at the following conferences and workshops:

- (Keynote speaker) ICALP 2011. The 38th International Colloquium on Automata, Languages and Programming, Zürich, Switzerland, July 2011.
- SecCo 2011. International Workshop on Security Issues in Concurrency. (Part of Aachen Concurrency and Dependability Week.) Aachen, Germany, September 2011.

8.1.4. Organization of workshops and conferences

- Catuscia Palamidessi has served as PC co-chair of QEST 2011. 8th International Conference on Quantitative Evaluation of SysTems. Aachen, Germany, August 2011.
- Catuscia Palamidessi has served as PC co-chairs of the first edition of TOSCA (Theory Of SeCurity and Applications). Associated with the ETAPS conferences, http://www.etaps.org/.
- Catuscia Palamidessi is serving as PC co-chair of TGC 2012. 7th International Symposium on Trustworthy Global Computing. Newcastle, UK, 7-8 September 2012.
- Frank Valencia has co-chaired EXPRESS 2011, the workshop on Expressiveness in Concurrency Theory (http://www.lix.polytechnique.fr/comete/EXPRESS11/. Affiliated to CONCUR 2011, the 22nd International Conference in Concurrency Theory (http://concur2010.inria.fr/). Aachen, Germany, 2011.
- Konstantino Chatzikokolakis has co-chaired SecCo 2011, the 9th International Workshop on Security Issues in Concurrency (http://www.lix.polytechnique.fr/~kostas/SecCo2011/. Affiliated to CONCUR 2011, the 22nd International Conference in Concurrency Theory (http://concur2010.inria.fr/). Aachen, Germany, 2011.

8.1.5. Participation in program committees

Catuscia Palamidessi has been/is a member of the program committees of the following conferences:

- FOSSACS 2013. 16th Int.l Conf. on Foundations of Software Science and Computation Structures. (Part of ETAPS 2013.) Rome, Italy, March 2013.
- QEST 2012. International Conference on Quantitative Evaluation of SysTems. London, UK, September 2012.
- PPDP 2012. International ACM SIGPLAN Conference on Principles and Practice of Declarative Programming. Leuven, Belgium, September 2012.
- CONCUR 2012. 21st International Conference on Concurrency Theory. Newcastle, UK, September 2012.
- CSF 2012. The 25th IEEE Computer Security Foundations Symposium. Cambridge MA, USA, June 2012.

- POST 2012. First Conference on Principles of Security and Trust. Tallin, Estonia, March 2012.
- CALCO 2011. Fourth International Conference on Algebra and Coalgebra in Computer Science. Winchester, UK, August 2011.
- CSF 2011. The 24th IEEE Computer Security Foundations Symposium. Domaine de l'Abbaye des Vaux de Cernay, France, June 2011.
- 27th Int.l Conference on Mathematical Foundations of Programming Semantics. CMU, Pittsburgh, USA, May 2011.

Konstantinos Chatzikokolakis has been/is a member of the program committees of the following conferences and workshops:

- ESOP 2012: 21th European Symposium on Programming.
- ISPEC 2012: 8th International Conference on Information Security Practice and Experience.
- QAPL 2012: 10th Workshop on Quantitative Aspects of Programming Languages .
- TGC 2011: 6th International Symposium on Trustworthy Global Computing.
- QAPL 2011: 9th Workshop on Quantitative Aspects of Programming Languages.
- FAST 2011: The 8th International Workshop on Formal Aspects of Security & Trust.

8.1.6. Participation in other committees

Catuscia Palamidessi has served in the following committees:

- The EAPLS PhD Award (http://eapls.org/pages/phd_award/). 2010-11.

8.1.7. Organization of seminars

- Frank D. Valencia and Andrés Aristizábal are the organizer of the Comète-Parsifal Seminar. This seminar takes place weekly at LIX, and it is meant as a forum where the members of Comète and Parsifal present their current works and exchange ideas. See http://www.lix.polytechnique.fr/comete/seminar/.

8.2. Visitors

- Geoffrey Smith, Professor, Florida International University, USA. He has visited Comète for three months from the end of August until the end of December, 2011.
- Moreno Falaschi, Professor, University of Siena, Italy. He has visited Comète for one month in November 2011.
- Vladimiro Sassone, Professor, University of Southampton, UK. He has visited Comète for one month in December 2011.

8.3. Service

Catuscia Palamidessi has served as:

- Member of the Comité d'Orientation Scientifique et Technique, Groupe de travail Relation Internationales (COST-GTRI). Since November 2007.
- Directrice adjointe du LIX, le Laboratoire d'Informatique de l'Ecole Polytechnique. Since April 2010.
- President of the selection committee for the EATCS Best Paper Award at the ETAPS conferences. Since 2006.
- Member of the Comité de These for Mathematics and Computer Science at the École Polytechnique. Since October 2007.
- Reviewer for the projects proposal for the program PRIN, sponsored by the Italian MIUR ("Ministero dell'Istruzione, dell'Università e della Ricerca"). Since 2004.
- Member of the Comité Academique de l'Ecole Polytechnique. Since November 2010.

Frank Valencia has served as:

- Member of the Evaluation Committee of the LIX/Qualcomm postdoc grants for the year 2011.

8.4. Teaching

Master: Konstantinos Chatzikokolakis has been teaching the course "Concurrence" at the "Master Parisien de Recherche en Informatique" (MPRI) in Paris. Level M2. Total 12 hours.

Doctorat: Catuscia Palamidessi has been teaching a course on Quantitative Information Flow at the 11th International School on Foundations of Security Analysis and Design , Bertinoro, Italy. Total 6 hours.

PhD : Mario Sergio Ferreira Alvim Junior. Ecole Polytechnique. Grant CNRS/DGA. Title of his PhD thesis: *Quantitative Approaches for Information Flow: An Analysis of Interactive Systems and Statistical Databases*. Defended on 12 October 2011. Supervised by Catuscia Palamidessi.

PhD in progress (2011-) Lili Xu. Ecole Polytechnique and Chinese academy of Science, Beijing, China. Co-supervised by Catuscia Palamidessi and Huimin Li.

PhD in progress (2011-) Nicolás Bordenabe. Ecole Polytechnique. Grant INRIA/DGA. Supervised by Catuscia Palamidessi.

PhD in progress (2011-) Luis Fernando Pino Duque. Ecole Polytechnique. Grant INRIA/DGA. Co-supervised by Catuscia Palamidessi and Frank D. Valencia.

PhD in progress (2010-) Sophia Knight. Ecole Polytechnique. Grant INRIA/CORDIS. Co-supervised by Catuscia Palamidessi and Frank D. Valencia.

PhD in progress (2009-) Andrés Aristizábal. Ecole Polytechnique. Grant CNRS/DGA. Co-supervised by Catuscia Palamidessi and Frank D. Valencia.

PhD in progress (2009-) Ivan Gazeau. Ecole Polytechnique. Grant ANR. Co-supervised by Catuscia Palamidessi and Dale Miller.

PhD in progress : Nom du doctorant, titre (provisoire) du mémoire, date du début de la thèse, encadrant(s)

8.4.1. PhD defenses

Catuscia Palamidessi has been "rapporteur" for the thesis of the following PhD students:

- Jacopo Mauro (University of Bologna, Italy). PhD thesis reviewer. Title of the thesis: *Constraints meet Concurrency*. Advised by Maurizio Gabbrielli. Defended in April 2012.
- Morgan Barbier (Ecole Polytechnique, France). Member of the committee board at the PhD defense. Title of the thesis: *Décodage en liste et application à la sécurité de l'information*. Advised by Daniel Augot. Defended in December 2011.
- Giulio Caravagna (University of Pisa, Italy). PhD thesis reviewer. Title of the thesis: *Formal Modeling and Simulation of Biological Systems with Delays.* Defended in December 2011.
- Robert Abo (Conservatoire National des Arts et Métiers, France). PhD thesis reviewer and member of the committee board at the PhD defense. Title of the thesis: *Approches formelles pour l?analyse de la performabilité des systèmes communicants mobiles : Application aux réseaux de capteurs sans fil*. Advised by Kamel Barkaoui. Defended in December 2011.
- Mathieu Sassolas (University of Paris VI, France). PhD thesis reviewer and member of the committee board at the PhD defense. Title of the thesis: *Méthodes qualitatives et quantitatives pour la détection d'information cachée*. Advised by Béatrice Bérard. Defended in November 2011.

9. Bibliography

Major publications by the team in recent years

[1] K. Chatzikokolakis, C. Palamidessi. *Making Random Choices Invisible to the Scheduler*, in "Information and Computation", 2010, vol. 208, n^o 6, p. 694-715 [*DOI :* 10.1016/j.ic.2009.06.006], http://hal.inria.fr/inria-00424860/en.

[2] K. Chatzikokolakis, C. Palamidessi, P. Panangaden. *Anonymity Protocols as Noisy Channels*, in "Information and Computation", 2008, vol. 206, n^o 2–4, p. 378–401 [*DOI :* 10.1016/j.ic.2007.07.003], http://hal.inria.fr/inria-00349225/en/.

[3] K. Chatzikokolakis, C. Palamidessi, P. Panangaden. *On the Bayes risk in information-hiding protocols*, in "Journal of Computer Security", 2008, vol. 16, n^o 5, p. 531–571 [*DOI :* 10.3233/JCS-2008-0333], http://hal.inria.fr/inria-00349224/en/.

[4] Y. Deng, C. Palamidessi. *Axiomatizations for probabilistic finite-state behaviors*, in "Theoretical Computer Science", 2007, vol. 373, n^o 1-2, p. 92–114, http://hal.inria.fr/inria-00200928/en/.

[5] P. Giambiagi, G. Schneider, F. D. Valencia. *On the Expressiveness of Infinite Behavior and Name Scoping in Process Calculi.*, in "Proceedings of FoSSaCS", Lecture Notes in Computer Science, Springer, 2004, vol. 2987, p. 226-240.

[6] S. Hamadou, C. Palamidessi, V. Sassone. *Reconciling Belief and Vulnerability in Information Flow*, in "31st IEEE Symposium on Security and Privacy", Berleley/Oakland, California, USA, IEEE Computer Society, 2010, p. 79-92 [*DOI :* 10.1109/SP.2010.13], http://hal.inria.fr/inria-00548007/en.

[7] C. Palamidessi, O. M. Herescu. *A randomized encoding of the π-calculus with mixed choice*, in "Theoretical Computer Science", 2005, vol. 335, n^o 2-3, p. 73-404, http://hal.inria.fr/inria-00201105/en/.

[8] C. Palamidessi. *Comparing the Expressive Power of the Synchronous and the Asynchronous pi-calculus*, in "Mathematical Structures in Computer Science", 2003, vol. 13, n^o 5, p. 685–719, http://hal.inria.fr/inria-00201104/en/.

[9] C. Palamidessi, V. A. Saraswat, F. D. Valencia, B. Victor. *On the Expressiveness of Linearity vs Persistence in the Asynchronous pi-calculus*, in "Proceedings of the Twenty First Annual IEEE Symposium on Logic in Computer Science (LICS)", IEEE Computer Society, 2006, p. 59–68, http://hal.inria.fr/inria-00201096/en/.

[10] F. D. Valencia. *Decidability of infinite-state timed CCP processes and first-order LTL*, in "Theoretical Computer Science", 2005, vol. 330, n^o 3, p. 577–607.

Publications of the year

Doctoral Dissertations and Habilitation Theses

[11] M. Alvim. *Des approches formelles pour le cachement d'information: Une analyse des systèmes interactifs, contrôle de divulgation statistique, et le raffinement des spécifications*, Ecole Polytechnique X, October 2011, http://hal.inria.fr/tel-00639948/en.

[12] M. ANDRES. *Quantitative Analysis of Information Leakage in Probabilistic and Nondeterministic Systems*, Radboud University, Nijmegen, July 2011, http://hal.inria.fr/tel-00655506/en.

Articles in International Peer-Reviewed Journal

[13] M. ALVIM, M. ANDRES, C. PALAMIDESSI. *Information Flow in Interactive Systems*, in "Journal of Computer Security", 2011, To appear, http://hal.inria.fr/inria-00637356/en.

[14] M. ANDRES, C. PALAMIDESSI, A. SOKOLOVA, P. VAN ROSSUM. *Information Hiding in Probabilistic Concurrent Systems*, in "Journal of Theoretical Computer Science", 2011, vol. 412, n^{o} 28, p. 3072-3089, http://hal.inria.fr/hal-00573447/en.

[15] K. CHATZIKOKOLAKIS, S. KNIGHT, C. PALAMIDESSI, P. PANANGADEN. *Epistemic Strategies and Games on Concurrent Processes*, in "Transactions on Computational Logic", 2011, http://hal.inria.fr/inria-00637160/en.

Invited Conferences

[16] M. ALVIM, M. ANDRES, K. CHATZIKOKOLAKIS, C. PALAMIDESSI. *On the relation between Differential Privacy and Quantitative Information Flow*, in "38th International Colloquium on Automata, Languages and Programming - ICALP 2011", Zurich, Switzerland, L. ACETO, M. HENZINGER, J. SGALL (editors), Lecture Notes in Computer Science, Springer, 2011, vol. 6756, p. 60-76 [*DOI :* 10.1007/978-3-642-22012-8_4], http://hal.inria.fr/inria-00627937/en.

International Conferences with Proceedings

[17] M. ALVIM, M. ANDRES, K. CHATZIKOKOLAKIS, P. DEGANO, C. PALAMIDESSI. *Differential Privacy: on the trade-off between Utility and Information Leakage*, in "The 8th International Workshop on Formal Aspects of Security & Trust (FAST)", Leuven, Belgium, G. BARTHE, A. DATTA, S. ETALLE (editors), Lecture Notes in Computer Science, Springer, March 2011, http://hal.inria.fr/inria-00580122/en.

[18] A. ARISTIZABAL, F. BONCHI, C. PALAMIDESSI, L. PINO, F. D. VALENCIA. *Deriving Labels and Bisimilarity for Concurrent Constraint Programming*, in "Proceedings of the 14th International Conference on Foundations of Software Science an Computation Structures (FOSSACS 2011). ", Saarbrücken, Germany, M. HOFMANN (editor), Lecture Notes in Computer Science, Springer, 2011, vol. 6604, p. 138-152 [*DOI :* 10.1007/ISBN 978-3-642-19804-5], http://hal.inria.fr/hal-00546722/en.

[19] A. ARISTIZABAL, F. BONCHI, L. PINO, F. D. VALENCIA. *Partition Refinement for Bisimilarity in CCP*, in "Proceedings of the 27th ACM Symposium On Applied Computing", Riva del Garda (Trento), Italy, ACM, March 26-30 2011, To appear, http://hal.inria.fr/hal-00641408/en.

[20] D. BAELDE, R. BEAUXIS, S. MIMRAM. *Liquidsoap: a High-Level Programming Language for Multimedia Streaming*, in "SOFSEM 2011: Theory and Practice of Computer Science", Nový Smokovec, Slovakia, I. CERNÁ, T. GYIMÓTHY, J. HROMKOVIC, K. JEFFEREY, R. KRÁLOVIC, M. VUKOLIC, S. WOLF (editors), Lecture Notes in Computer Science, Springer Berlin / Heidelberg, 2011, vol. 6543, p. 99-110 [*DOI :* 10.1007/978-3-642-18381-2_8], http://hal.inria.fr/inria-00585728/en.

[21] R. BEAUXIS, S. MIMRAM. *A Non-Standard Semantics for Kahn Networks in Continuous Time*, in "Computer Science Logic (CSL'11) - 25th International Workshop/20th Annual Conference of the EACSL", Bergen, Norway, M. BEZEM (editor), Leibniz International Proceedings in Informatics (LIPIcs), Schloss

Dagstuhl–Leibniz-Zentrum fuer Informatik, 2011, vol. 12, p. 35–50 [*DOI :* 10.4230/LIPIcs.CSL.2011.35], http://hal.inria.fr/inria-00616968/en.

[22] M. GIUNTI. *A type checking algorithm for qualified session types*, in "7th International Workshop on Automated Specification and Verification of Web Systems", Reykjavik, Iceland, L. KOVÁCS, R. PUGLIESE, F. TIEZZI (editors), Electronic Proceedings in Theoretical Computer Science, August 2011 [*DOI :* 10.4204/EPTCS.61.7], http://hal.inria.fr/hal-00644061/en.

[23] D. HERMITH, C. OLARTE, C. RUEDA, F. D. VALENCIA. *Modeling Cellular Signaling Systems: An Abstraction-Refinement Approach*, in "PACBB", Salamanca, Spain, M. P. ROCHA, J. M. C. RODRIGUEZ, F. FDEZ-RIVEROLA, A. VALENCIA (editors), Advances in Intelligent and Soft Computing, Springer, 2011, vol. 93, p. 321-328 [*DOI :* 10.1007/ISBN 978-3-642-19913-4], http://hal.inria.fr/hal-00641433/en.

Scientific Books (or Scientific Book chapters)

[24] M. ALVIM, M. ANDRES, K. CHATZIKOKOLAKIS, C. PALAMIDESSI. *Quantitative Information Flow and Applications to Differential Privacy*, in "Foundations of Security Analysis and Design VI – FOSAD Tutorial Lectures", A. ALDINI, R. GORRIERI (editors), Lecture Notes in Computer Science, Springer, 2011, vol. 6858, p. 211–230 [*DOI :* 10.1007/978-3-642-23082-0_8], http://hal.inria.fr/hal-00655522/en.

[25] G. LONGO, C. PALAMIDESSI, P. THIERRY. *Some Bridging Results and Challenges in Classical, Quantum and Computational Randomness*, in "Randomness Through Computation", H. ZENIL (editor), World Scientific, 2011, ISBN: 978-981-4327-74-9, http://hal.inria.fr/hal-00445553/en.

Books or Proceedings Editing

[26] K. CHATZIKOKOLAKIS, V. CORTIER (editors). *Proceedings of the 8th International Workshop on Security Issues in Concurrency*, Electronic Proceedings in Theoretical Computer Science, Electronic Proceedings in Theoretical Computer Science, 2011, vol. 51, 51 [*DOI :* 10.4204/EPTCS.51], http://hal.inria.fr/hal-00641020/en.

[27] B. LUTTIK, F. D. VALENCIA (editors). *Proceedings of the Eighth International Conference on Quantitative Evaluation of SysTems*, Electronic Proceedings in Theoretical Computer Science, Electronic Proceedings in Theoretical Computer Science, 2011, vol. 64, 131 [*DOI :* 10.4204/EPTCS.64].

[28] S. MÖDERSHEIM, C. PALAMIDESSI (editors). *Post-proceedings of TOSCA – Theory of Security and Applications*, Lecture Notes in Computer Science, Springer, 2011, vol. 6993, 235, http://hal.inria.fr/hal-00655523/en.

[29] C. PALAMIDESSI, A. RISKA (editors). *Proceedings of the Eighth International Conference on Quantitative Evaluation of SysTems*, IEEE, 2011, 276, http://hal.inria.fr/hal-00655524/en.

Research Reports

[30] F. CASSEZ, J. DUBREIL, H. MARCHAND. *Synthesis of Opaque Systems with Static and Dynamic Masks.*, INRIA, 2011, Submitted to Formal Methods in System Design (FORM).

References in notes

[31] R. CANETTI, L. CHEUNG, N. LYNCH, O. PEREIRA. *On the Role of Scheduling in Simulation-Based Security*, 2007, Cryptology ePrint Archive, Report 2007/102.

[32] K. CHATZIKOKOLAKIS, C. PALAMIDESSI. *Making Random Choices Invisible to the Scheduler*, in "Information and Computation", 2010, vol. 208, n^{o} 6, p. 694-715 [*DOI :* 10.1016/J.IC.2009.06.006], http://hal.inria.fr/inria-00424860/en.

[33] K. CHATZIKOKOLAKIS, C. PALAMIDESSI, P. PANANGADEN. *Anonymity Protocols as Noisy Channels*, in "Inf. and Comp.", 2008, vol. 206, n^{o} 2–4, p. 378–401 [*DOI :* 10.1016/J.IC.2007.07.003], http://hal.inria.fr/inria-00349225/en/.

[34] K. CHATZIKOKOLAKIS, C. PALAMIDESSI, P. PANANGADEN. *On the Bayes risk in information-hiding protocols*, in "Journal of Computer Security", 2008, vol. 16, n^{o} 5, p. 531–571 [*DOI :* 10.3233/JCS-2008-0333], http://hal.inria.fr/inria-00349224/en/.

[35] T. HOARE, R. MILNER. *Grand Challenges for Computing Research*, in "Computer Journal", 2005, vol. 48, n^{o} 1, p. 49-52.

[36] N. P. MENDLER, P. PANANGADEN, P. J. SCOTT, R. A. G. SEELY. *A Logical View of Concurrent Constraint Programming*, in "Nord. J. Comput.", 1995, vol. 2, n^{o} 2, p. 181-220.

[37] R. MILNER, J. PARROW, D. WALKER. *A Calculus of Mobile Processes, I and II*, in "Information and Computation", 1992, vol. 100, n^{o} 1, p. 1–40 & 41–77, A preliminary version appeared as Technical Reports ECF-LFCS-89-85 and -86, University of Edinburgh, 1989..

[38] G. NORMAN, C. PALAMIDESSI, D. PARKER, P. WU. *Model checking probabilistic and stochastic extensions of the π-calculus*, in "IEEE Transactions of Software Engineering", 2009, vol. 35, n^{o} 2, p. 209–223, http://hal.archives-ouvertes.fr/inria-00424856/en/.

[39] V. A. SARASWAT, M. RINARD, P. PANANGADEN. *Semantic foundations of concurrent constraint programming*, in "Conference Record of the Eighteenth Annual ACM Symposium on Principles of Programming Languages", ACM Press, 1991, p. 333–352.

[40] V. A. SARASWAT, M. C. RINARD. *Concurrent Constraint Programming*, in "POPL", ACM Press, 1990, p. 232-245.

[41] P. WU, C. PALAMIDESSI, H. LIN. *Symbolic Bisimulation for Probabilistic Systems*, in "Proceedings of 4th International Conference on the Quantitative Evaluation of SysTems (QEST)", IEEE Computer Society, 2007, p. 179-188, http://www.lix.polytechnique.fr/~catuscia/papers/Wu/qest2.pdf.