INRIA

INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET EN AUTOMATIQUE

Team Comète

Concurrence, Mobilité et Transactions

Futurs

THEME COM

Table of contents

1. Team

Joint team with LIX (Laboratoire d'Informatique de l'École Polytechnique) and CNRS.

Project-team Leader
Catuscia Palamidessi [Research Director (DR) INRIA, HdR]

Administrative assistant
Lydie Fontaine [Secretary (SAR) INRIA]

Staff members CNRS
Frank Valencia [Research Associate (CR) CNRS]

Ph. D. students
Kostas Chatzikokolakis [Allocataire École Polytechnique - Ministère. Till 26/10/2007]
Romain Beauxis [Allocataire Region Ile de France]
Sylvain Pradalier [Allocataire ENS Cachan. Co-supervised by Cosimo Laneve, University of Bologna, Italy]
Carlos Olarte [Allocataire INRIA/CORDIS]
Jesus Aranda [Co-supervised by Juan Francisco Diaz, Universidad del Valle, Colombia]
Christelle Braun [Allocataire École Polytechnique - Ministère. From 1/10/2007]

Visitors
Robin Milner [Professor, University of Cambridge, UK. One year visit.]
Moreno Falaschi [Professor, University of Siena, Italy. Two months visit.]
Antonino Salibra [Professor, University of Venice, Italy. Two months visit.]
Cosimo Laneve [Professor, University of Bologna, Italy. One month visit.]
Diletta Romana Cacciagrano [Assistant Professor, University of Camerino, Italy. One month visit.]
Sergio Maffeis [PostDoc, Imperial College, UK. One month visit]
Andrea Turrini [PhD student, University of Verona, Italy. Three months visit.]
Troels C. Damgaard [PhD student, IT University of Copenhagen, Denmark. Four months visit.]

Post-doctoral fellows
Peng Wu [Post Doctorant CNRS. Till 30/9/2006]
Angelo Troina [Post Doctorant INRIA. Till 31/10/2006]
Simon Kramer [Post Doctorant INRIA. From 1/11/2006]

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* [43]). *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*.

These systems 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 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, computer scientists 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 of the year

- Robin Milner visits the Comète team for one year. His visit is supported by a Blaise Pascal chair. Milner is an outstanding scientist who has given many fundamental contributions to the fields of Functional Languages and of Concurrency. He has received many prestigious recognitions, among which the Turing award.
- Catuscia Palamidessi gives various invited talks and tutorials at international conferences and workshops on the work done with Prakash Panangaden and Kostas Chatzikokolakis in the context of the INRIA/DREI project Printemps.

3. Scientific Foundations

3.1. Probabilistic aspects

Keywords: *Probability*.

Participants: Romain Beauxis, Christelle Braun, Kostas Chatzikokolakis, Catuscia Palamidessi, Sylvain Pradalier, Angelo Troina, Peng Wu.

The need to deal with probabilities can arise for various reasons:

First, algorithms for distributed systems and security protocols often use randomization.

Second, the modeling of the physical world frequently requires coping with uncertain and approximate information (for example, the number of the requests that are received by a web server during various times of the day), which one can refine by statistical measurements, and which can then be naturally represented using a probabilistic formalism.

Third, reality can sometimes be too complicated to be represented and analyzed in detail; probabilistic models offer then a convenient abstraction mechanism.

3.2. Expressiveness issues

Keywords: *Expressiveness*.

Participants: Jesus Aranda, Romain Beauxis, Catuscia Palamidessi, Carlos Olarte, 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.

3.3. The probabilistic asynchronous π-calculus

Keywords: *Process calculi*.

Participants: Jesus Aranda, Romain Beauxis, Catuscia Palamidessi, Carlos Olarte, Angelo Troina, Frank Valencia.

We will focus our efforts on a probabilistic variant of the asynchronous π-calculus, that 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.

4. Application Domains

4.1. Security

Keywords: *anonymity*, *privacy*, *security*, *unobservability*.

Participants: Romain Beauxis, Christelle Braun, Kostas Chatzikokolakis, Simon Kramer, Catuscia Palamidessi, Angelo Troina.

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 *privacy*, because they exhibit features that require the kind of concepts and approach in which we feel most competent. It is likely, however, that the instruments and tools developed having privacy in mind can later be useful and adaptable also to other domains of security, like *Secure Information flow*. Privacy is a generic term which denotes the issue of preventing certain information to become known to an agent, except in case that agent is explicitly allowed to be informed. 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 purpose of privacy protocols is then to obfuscate the link between secrets and observables as much as possible, and they often use randomization to achieve this purpose, i.e. to introduce *noise*. The protocol can therefore 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 privacy, and to develop heuristics and methods to improve protocols for privacy. 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.

4.2. Model checking

Keywords: *automatic verification*.

Participants: Romain Beauxis, Kostas Chatzikokolakis, Catuscia Palamidessi, Peng Wu.

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. These are (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. 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.

5. Software

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

Participants: Catuscia Palamidessi, Kostas Chatzikokolakis, Peng Wu.

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/.

6. New Results

6.1. Expressive power of models and formalisms for concurrency

Participants: Jesus Aranda, Romain Beauxis, Catuscia Palamidessi, Carlos Olarte, Frank Valencia.

6.1.1. Infinite Behavior and Name Scoping in Process Calculi

Process calculi differ in the constructs for the specification of infinite behavior and in the scoping rules for channel names.

In [20] we have surveyed various notions of scope and infinite behavior proposed in literature, and we have pointed out their impact in the expressive power of concurrent formalisms.

In [19] we have proved that the CCS variant with replication mentioned above can faithfully (deterministically) encode regular languages but not context-free ones. We have also proved that the languages generated by the processes of this variant are context-sensitive.

6.1.2. Synchronous vs Asynchronous Communication

One of the early results about the asynchronous π-calculus which significantly contributed to its popularity is the capability of encoding the output prefix of the (choiceless) π-calculus in a natural and elegant way. Encodings of this kind were proposed by Honda and Tokoro [44], and by Boudol [36]. In [13], we have investigated whether the above encodings preserve De Nicola and Hennessy's testing semantics. It turns out that, under some general conditions, no encoding of output prefix is able to preserve the must testing. This negative result is due to (a) the non atomicity of the sequences of steps which are necessary in the asynchronous π-calculus to mimic synchronous communication, and (b) testing semantics's sensitivity to divergence. The preservation of testing semantics is however ensured if we assume some form of fairness.

6.1.3. Fairness

In [22] we have defined fair computations in the π-calculus. We have followed Costa and Stirling's approach for CCS-like languages [39], [40] but exploited a more natural labeling method of process actions to filter out unfair process executions. The new labeling allowed us to prove all the significant properties of the original one, such as unicity, persistence and disappearance of labels. It also turned out that the labeled π-calculus is a conservative extension of the standard one. We contrasted the existing fair testing notions [37], [46] with those that naturally arise by imposing weak and strong fairness. This comparison provides the expressiveness of the various fair testing-based semantics and emphasizes the discriminating power of the one already proposed in the literature.

6.1.4. Distributed Agreement

In [18] we have systematized a collection of results on the expressiveness of process calculi obtained by the means of impossibility results in the field of distributed computing. In particular, we have focused on the *symmetric leader election problem* which allows to classify languages based on their capability of achieving a distributed agreement.

6.1.5. Linearity vs Persistence

In [21] we have compared the expressive power of linear and persistent communication in the context of weak bisimilarity. We have considered four fragments of the π-calculus, corresponding to combinations of linearity/persistence also present in other frameworks such as Concurrent Constraint Programming and several calculi for security. The study is presented by providing (or proving the non-existence of) encodings among the fragments, a processes-as-formulae interpretation and a reduction from Minsky machines.

6.2. Semantics of probabilistic systems

Participants: Romain Beauxis, Catuscia Palamidessi, Angelo Troina.

One of the goals of Comète is to investigate the foundations of probabilistic calculi, and in particular the probabilistic asynchronous π-calculus.

6.2.1. Bisimulation semantics

In [16] we have studied a process calculus which combines both nondeterministic and probabilistic behavior in the style of Segala and Lynch's probabilistic automata. We have considered various strong and weak behavioral equivalences, and we have provided complete axiomatizations for finite-state processes, restricted to guarded definitions in case of the weak equivalences. We conjecture that in the general case of unguarded recursion the "natural" weak equivalences are undecidable.

This has been the first work, to our knowledge, to provide a complete axiomatization for weak equivalences in the presence of recursion and both nondeterministic and probabilistic choice.

6.2.2. Parametric Probabilities

In [17] we have developed a model of Parametric Probabilistic Transition Systems, where probabilities associated with transitions may be parameters. We have showed how to find instances of the parameters that satisfy a given property and instances that either maximize or minimize the probability of reaching a certain state. As an application, we have modeled a probabilistic non–repudiation protocol with a Parametric Probabilistic Transition System. The theory we have developed allows us to find instances that maximize the probability that the protocol ends in a fair state (i.e. no participant has an advantage over the others).

6.3. Specification and verification of security protocols

Participants: Kostas Chatzikokolakis, Catuscia Palamidessi, Carlos Olarte, Frank Valencia.

6.3.1. Probabilistic protocols

Probabilistic security protocols involve *probabilistic choices* and are used for many purposes including signing contracts, sending certified email and protecting the anonymity of communication agents. Some probabilistic protocols rely on specific random primitives such as the *Oblivious Transfer* [51]. There are various examples in this category, notably the contract signing protocol in [41] and the privacy-preserving auction protocol in [45].

A large effort has been dedicated to the formal verification of security protocols, and several approaches based on process-calculi techniques have been proposed. However, in the particular case of probabilistic protocols, only few attempts of this kind have been made.

In [14] we have developed a framework for analyzing probabilistic security protocols using a probabilistic extension of the π-calculus inspired by the work in [42], [48]. In order to express security properties in this calculus, we have extended the notion of testing equivalence [47] to the probabilistic setting. We have applied these techniques to verify the Partial Secret Exchange, a protocol which uses a randomized primitive, the Oblivious Transfer, to achieve fairness of information exchange between two parties.

6.3.2. Universal Timed Concurrent Constraint Programming

In [32] we have introduced the *Universal Timed Concurrent Constraint Programming* (`utcc`) process calculus; a generalisation of Timed Concurrent Constraint Programming. The `utcc` calculus allows for the specification of mobile behaviours in the sense of Milner's π-calculus: Generation and communication of private channels or links. We first endow `utcc` with an *operational* semantics and then with a *symbolic* semantics to deal with problematic operational aspects involving infinitely many substitutions and divergent internal computations. The novelty of the symbolic semantics is to use *temporal constraints* to represent finitely infinitely-many substitutions. We also show that `utcc` has a strong connection with Pnueli's Temporal Logic. This connection can be used to prove reachability properties of `utcc` processes. As a compelling example, we have used `utcc` to exhibit the secrecy flaw of the Needham-Schroeder security protocol.

In [27] we introduced a framework for the declarative debugging of tcc programs. We expect to adapt this framework to our work in [32] and use them to debug security protocols specified in `utcc`.

6.4. Foundations of anonymity and privacy

Participants: Kostas Chatzikokolakis, Catuscia Palamidessi, Peng Wu.

6.4.1. Probability and nondeterminism

The systems for ensuring anonymity often use random mechanisms which can be described probabilistically, while the agents' interest in performing the anonymous action may be totally unpredictable, irregular, and hence expressible only nondeterministically. In the past, formal definitions of the concept of anonymity have been investigated either in a totally nondeterministic framework, or in a purely probabilistic one.

In [35], [50], [49] we have proposed a notion of strong anonymity which combines both probability and nondeterminism, and which is suitable for describing the most general situation in which both the systems and the user can have both probabilistic and nondeterministic behavior. We have also investigated the properties of the definition for the particular cases of purely nondeterministic users and purely probabilistic users. One interesting feature of our approach is that in the purely probabilistic case, strong anonymity turns out to be independent from the probability distribution of the users.

In [26] we have also investigated a notion of weak anonymity. This is a more realistic notion in the sense that it is more likely to be satisfied by the anonymity protocols used in practice.

Our notions of anonymity are defined in terms of observables for processes in the probabilistic π-calculus. As one of the goals of the project is to develop a model checker and other verification tools for this calculus, that will provide also a way to check automatically that the protocols satisfy the intended anonymity properties.

6.4.2. The problem of the scheduler

It has been observed recently that in security the combination of nondeterminism and probability can be harmful, in the sense that the resolution of the nondeterminism can reveal the outcome of the probabilistic choices even though they are supposed to be secret [38]. This is known as the problem of the *information-leaking scheduler*. In [23] we have developed a linguistic (process-calculus) approach to this problem, and we have shown how to apply it to control the behavior of the scheduler in various anonymity examples.

6.4.3. Information-Theory

In [15] we have proposed a framework in which anonymity protocols are interpreted as particular kinds of channels, and the degree of anonymity provided by the protocol as the converse of the channel's capacity. We have then illustrated how various notions of anonymity can be expressed in this framework, and showed the relation with some definitions of probabilistic anonymity in literature. Finally, we have discussed how to compute the channel matrix on the basis of the transition system associated to the protocol, and how to perform the computation automatically using a model-checker like PRISM.

In [24] we have investigated how the adversary can test the system to try to infer the user's identity, and we have studied how the probability of error depends on the characteristics of the channel. In particular we have considered the Bayes approach, and we have been able to characterize the associated probability of error (Bayes risk) in terms of the solution of certain systems of equations derived from the channel. This has allowed us to compute tight bounds for the Bayes risk, thus improving long-standing results in literature.

The PhD thesis of Kostas Chatzikokolakis, which has been defended on October 26, 2007, is largely based on the results described in this section.

6.5. The probabilistic applied π-calculus

Participants: Catuscia Palamidessi, Angelo Troina.

In order to obtain a language suitable for the specification and verification of a large class of security protocols, we aim at enriching the probabilistic π-calculus with value passing, encryption and decryption, other primitive functions, and data types, along the lines of the *applied π-calculus* [34].

Some preliminary work in this direction is represented by [28]. We have investigated an extension of the Applied π-calculus obtained by introducing nondeterministic and probabilistic choice operators. The semantics of the resulting model, in which probability and nondeterminism are combined, is given by Segala's Probabilistic Automata driven by schedulers which resolve the nondeterministic choice among the probability distributions over target states. We have provided notions of static and observational equivalence for the enriched calculus. In order to model the possible interaction of a process with its surrounding environment, we have given a labeled semantics together with a notion of weak bisimulation which is shown to coincide with the observational equivalence. Finally, we have proved that results in the probabilistic framework are preserved in a purely nondeterministic setting.

6.6. Model checking

Participants: Kostas Chatzikokolakis, Catuscia Palamidessi, Peng Wu.

Model checking is the main tool that we aim at developing for the verification of security protocols.

In [33] we have introduced a weak symbolic bisimulation for the probabilistic π-calculus to overcome the infinite branching problem in checking ground bisimulations between probabilistic systems. The definition of weak symbolic bisimulation does not rely on the random capability of adversaries and suggests a solution to the open problem on the axiomatization for weak bisimulation in the case of unguarded recursion. Furthermore, we have presented an efficient characterization of symbolic bisimulations for the calculus, which allows the "on-the-fly" instantiation of bound names and dynamic construction of equivalence relations for quantitative evaluation. This has directly resulted in a local decision algorithm that can explore just a minimal portion of the state spaces of the probabilistic processes in question.

In [30], in collaboration with the PRISM team at Oxford, we have established the basis for an implementation of model checking for the probabilistic π-calculus. Building upon the (non-probabilistic) π-calculus model checker MMC [52], we have developed an automated procedure for constructing a Markov decision process representing a probabilistic π-calculus process. This representation can then be verified using existing probabilistic model checkers such as PRISM. Secondly, we have demonstrated how for a large class of systems an efficient, compositional approach can be applied, which uses our extension of MMC on each parallel component of the system and then translates the results into a higher-level model description for the PRISM tool.

6.7. Modeling biological systems

Participants: Jesus Aranda, Sylvain Pradalier, Frank Valencia.

6.7.1. Timed Concurrent Constraint Programming

Quantitative and partial information may help to better describe the behavior of many real-life systems. In the particular case of biological ones, the former is fundamental for description and experimentation purposes, and the latter allows to represent those facts that are not precisely known. Moreover, the dynamic nature of these systems makes the use of time in system descriptions a mandatory requirement. In [29] we have proposed ntcc, a timed concurrent constraint process calculus, as a convenient language to model biological systems. ntcc allows to describe both non-deterministic and asynchronous behavior, useful features for describing many scenarios such as unpredictable biological events. A crucial advantage of using ntcc is that interesting properties of biological models can be verified by appealing to its associated proof system. The advantages of following this approach are demonstrated by modelling the Sodium-Potassium pump, a cellular mechanism present in many live organisms.

6.7.2. The nanoκ *calculus*

In [25] we have developed a process calculus – the nanoκ calculus – for modeling, analyzing and predicting the properties of molecular devices. The nanoκ calculus is equipped with a simple stochastic model, that we use to model and simulate the behaviour of a molecular shuttle, a basic nano device currently used for building more complex systems.

7. Other Grants and Activities

7.1. Actions nationales

7.1.1. Project INRIA/ARC PRONOBIS

The project PRONOBIS started in January 2006 and will end in December 2007. The consorsium is composed as follows:

- ENS Cachan. Responsible: J. Gobault-Larrecq
- INRIA Futurs. Responsible: C. Palamidessi
- Oxford University, UK. Responsible: M. Kwiatkowska
- University of Verona, Italy. Responsible: R. Segala

The goal of the ProNobis project is to explore mixing probability and non-determinism in the semantics of transition systems, and also of programming languages. We plan to keep one eye on applications to typical computer related problems, in particular to problems stemming from security. Several interesting verification problems related to security involve proving that two processes are contextually equivalent. This usally uses notions such as bisimulation, which need to be better understood in a setting where probabilities, external non-determinism (choosing which action to fire in Markov decision processes), and internal non-determinism (where no visible action distinguishes between the various alternatives).

Home Page: http://www.lsv.ens-cachan.fr/~goubault/ProNobis/pronobis1index.html.

Some publications representative of this collaboration are [30], [28].

7.1.2. LIX project on Distributed, Mobile and Secure Complex Systems

This project is finanged by the DGA, for the years 2007 and 2008. The teams involved are:

- Hipercom. Responsible: Philippe Jacquet
- Comète. Responsible: C. Palamidessi
- Algorithmes et Optimisation. Responsible: Philippe Baptiste
- MAX. Responsible: Michel Fliess. 2007-2008.

7.2. Actions internationales

7.2.1. DREI Equipe Associée PRINTEMPS

The project has started in December 2005 and includes the following sites:

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

PRINTEMPS focuses on the applications of Information Theory to security. We are particularly interested in studying the interactions between Concurrency and Information Theory.

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

Some publications representative of this collaboration are [15], [24].

7.2.2. REACT: Robust theories for Emerging Applications in Concurrency Theory

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

- Pontificia Universidad Javeriana, Colombia. Responsible: C. Rueda
- INRIA Futurs. Responsible: F. Valencia
- IRCAM, France.

Home page: http://cic.puj.edu.co/wiki/doku.php?id=grupos:avispa:react.

A publication representative of this collaboration is [29].

7.2.3. PAI project MONACO: MOdels for New Applications of COncurrency

. The project has started in January 2007 and will end in December 2008. It involves the following sites:

- Imperial College, UK. Responsible I. Phillips
- INRIA Futurs. Responsible: C. Palamidessi
- Technische Universität Berlin, Germany. Responsible: U. Nestmann.

A publications representative of this collaboration is [18].

8. Dissemination

8.1. Contribution to scientific events and activities

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.
- Catuscia Palamidessi is member of the Editorial Board of the journal on Theory and Practice of Logic Programming, published by the Cambridge University Press.
- Catuscia Palamidessi is member of the Editorial Board of the Electronic Notes of Theoretical Computer Science, Elsevier Science.
- Frank D. Valencia is area editor (for the area of Concurrency) of the ALP Newsletter.

8.1.2. Steering Committees

Catuscia Palamidessi is member of:

- The IFIP Technical Committee 1 – Foundations of Computer Science. Since 2007
- The Council of EATCS, the European Association for Theoretical Computer Science. Since 2005
- The IFIP Working Group 2.2 – Formal Description of Programming Concepts. Since 2001

8.1.3. Invited Talks

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

- Workshop on the Interplay of Programming Languages and Cryptography. Sophia Antipolis, France, 7 November 2007.
- Dagstuhl seminar on Formal Protocol Verification Applied. Dagstuhl, Germany, 14-19 October 2007.
- PLID'07. Programming Language Interference and Dependence. Kongens Lyngby, Denmark, 21 August 2007.
- PAuL'07. 2nd International Workshop on Probabilistic Automata and Logics. Wroclaw, Poland, 9 July 2007.
- PERAD 2007. Pervasive Adaptive Joint FET - EATCS Workshop. Brussels, Belgium, 26 January 2007.

8.1.4. Organization of workshops and conferences

- Catuscia Palamidessi has been the co-organizer of the workshop SecCo 2007, the 5th International Workshop on Security Issues in Concurrency. Lisboa, Portugal, September 2007. See http://www.dsi.uniroma1.it/~gorla/SecCo07/.

8.1.5. Participation in program committees

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

- QEST'08. International Conference on Quantitative Evaluation of SysTems. Saint Malo, France, September 2008.
- CONCUR'08. 19th International Conference on Concurrency Theory. Toronto, Canada, August 2008.
- CiE 2008: Logic and Theory of Algorithms. Athens, Greece. June 2008.
- FICS'08. Foundations of Informatics, Computing and Software. Shanghai, China, June 2008.
- LICS 2008. 23rd Symposium on Logic in Computer Science. Pittsburgh, USA. June 2008.
- MFPS XXIV. Twenty-fourth Conference on the Mathematical Foundations of Programming Semantics. University of Pennsylvania, Philadelphia, USA, May 2008.

- ESOP 2008. 17th European Symposium on Programming. (Part of ETAPS 2008.) Budapest, Hungary, March - April 2008.
- VMCAI 2008. 9th International Conference on Verification, Model Checking, and Abstract Interpretation. San Francisco, USA. January 2008.
- QEST'07. International Conference on Quantitative Evaluation of Systems. Edinburgh, UK, September 2007.
- CONCUR 2007. 18th International Conference on Concurrency Theory. Lisbon, Portugal, September 2007.
- FCT 2007. 16th International Symposium on Fundamentals of Computation Theory. Budapest, Hungary, August 2007.
- ESOP 2007. 16th European Symposium on Programming. (Part of ETAPS 2007.) Braga, Portugal, 24 March - 1 April, 2007.

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

- TFIT 2008. The Fourth Taiwanese-French Conference on Information Technology. Taipei, Taiwan, March 2008.
- FInCo 2007. Workshop on the Foundations of Interactive Computation. (Satellite event of ETAPS 2007). Braga, Portugal, March - April, 2007.

8.1.6. Organization of seminars

- Frank D. Valencia and Carlos Olarte 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. Service

- Catuscia Palamidessi has served as a member of the committee for the evaluation of the candidates to the INRIA Futurs positions of CR2 in the 2007 competition.
- Catuscia Palamidessi is a member of the INRIA GTRI (Group de Travail Relations Internationales) from November 2007 till October 2009.
- Catuscia Palamidessi is a member of the Comité de These for Mathematics and Computer Science at the École Polytechnique. From October 2007.

8.3. Teaching

8.3.1. Postgraduate courses:

- Catuscia Palamidessi and Frank Valencia are teaching (together with Francesco Zappa Nardelli and Roberto Amadio) the course "Concurrence" at the "Master Parisien de Recherche en Informatique" (MPRI) in Paris. Winter semester 2007-08.

8.3.2. Undergraduate courses:

- Frank D. Valencia has been a lecturer on "Concurrency Theory" at Universidad Javeriana de Cali. July 2007.

8.4. Advising

8.4.1. PhD students

Catuscia Palamidessi has supervised the following PhD students during 2007:

- Kostas Chatzikokolakis. Allocataire École Polytechnique - Ministère.
- Romain Beauxis. Allocataire Region Ile de France.
- Christelle Braun. Allocataire École Polytechnique - Ministère.
- Sylvain Pradalier. Allocataire ENS Cachan. Co-supervised by Cosimo Laneve, University of Bologna, Italy.

Catuscia Palamidessi and Frank Valencia have co-supervised the following PhD students

- Carlos Olarte. Allocataire INRIA/CORDIS.
- Jesus Aranda. Co-supervised by Juan Francisco Diaz, Universidad del Valle, Colombia.

8.4.2. PhD defenses

Catuscia Palamidessi has been “rapporteur” at the following PhD thesis defenses during 2007:

- Florent Garnier (Loria). PhD thesis on *Terminaison en temps moyen fini de systèmes de règles probabilistes*. Defended on 17 September, 2007. Advised by Claude Kirchner.
- Rémy Haemmerlé (INRIA Rocquencourt). PhD thesis on *Fermetures et Modules dans les langages concurrents avec contraintes fondés sur la logique linèaire*. Defended on December, 2007. Advised by François Fages.

9. Bibliography

Major publications by the team in recent years

[1] M. BHARGAVA, C. PALAMIDESSI. *Probabilistic Anonymity*, in "Proceedings of CONCUR", M. ABADI, L. DE ALFARO (editors), Lecture Notes in Computer Science, vol. 3653, Springer, 2005, p. 171–185, http://www.lix.polytechnique.fr/~catuscia/papers/Anonymity/concur.pdf.

[2] K. CHATZIKOKOLAKIS, C. PALAMIDESSI. *Probable Innocence Revisited*, in "Theoretical Computer Science", vol. 367, n^{o} 1-2, 2006, p. 123–138, http://www.lix.polytechnique.fr/~catuscia/papers/Anonymity/reportPI.pdf.

[3] K. CHATZIKOKOLAKIS, C. PALAMIDESSI, P. PANANGADEN. *Probability of Error in Information-Hiding Protocols*, in "Proceedings of the 20th IEEE Computer Security Foundations Symposium (CSF20)", IEEE Computer Society, 2007, p. 341-354, http://www.lix.polytechnique.fr/~catuscia/papers/ProbabilityError/full.pdf.

[4] Y. DENG, C. PALAMIDESSI. *Axiomatizations for probabilistic finite-state behaviors*, in "Theoretical Computer Science", vol. 373, n^{o} 1-2, 2007, p. 92–114, http://www.lix.polytechnique.fr/~catuscia/papers/Prob_Axiom/tcs.pdf.

[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, vol. 2987, Springer, 2004, p. 226-240, http://www.brics.dk/~fvalenci/papers/fossacs04.pdf.

[6] M. NIELSEN, C. PALAMIDESSI, F. VALENCIA. *Temporal Concurrent Constraint Programming: Denotation, Logic and Applications*, in "Nordic Journal of Computing", vol. 9, 2002, p. 145–188, http://www.lix.polytechnique.fr/~catuscia/papers/Ntcc/njc02.ps.

[7] C. PALAMIDESSI, O. M. HERESCU. *A randomized encoding of the π-calculus with mixed choice*, in "Theoretical Computer Science", vol. 335, n^o 2-3, 2005, p. 73-404, http://www.lix.polytechnique.fr/~catuscia/papers/prob_enc/report.pdf.

[8] C. PALAMIDESSI. *Comparing the Expressive Power of the Synchronous and the Asynchronous pi-calculus*, in "Mathematical Structures in Computer Science", vol. 13, n^o 5, 2003, p. 685–719, http://www.lix.polytechnique.fr/~catuscia/papers/pi_calc/mscs.pdf.

[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://www.lix.polytechnique.fr/~catuscia/papers/Frank/LICS_06/main.pdf.

[10] F. D. VALENCIA. *Decidability of infinite-state timed CCP processes and first-order LTL*, in "Theoretical Computer Science", vol. 330, n^o 3, 2005, p. 577–607, http://www.brics.dk/~fvalenci/papers/tcs.pdf.

Year Publications

Books and Monographs

[11] G. F. ITALIANO, C. PALAMIDESSI (editors). *Special issue of Theoretical Computer Science dedicated to a selection of the best papers presented at ICALP'05. 380(1-2)*, vol. 380, n^o 1-2, Elsevier, 2007, p. 1-218.

Doctoral dissertations and Habilitation theses

[12] K. CHATZIKOKOLAKIS. *Probabilistic and Information-Theoretic Approaches to Anonymity*, Ph. D. Thesis, LIX, École Polytechnique, October 2007, http://www.lix.polytechnique.fr/~kostas/thesis.pdf.

Articles in refereed journals and book chapters

[13] D. CACCIAGRANO, F. CORRADINI, C. PALAMIDESSI. *Separation of synchronous and asynchronous communication via testing*, in "Theoretical Computer Science", To appear, 2007, http://www.lix.polytechnique.fr/~catuscia/papers/Diletta/Must/tcs.pdf.

[14] K. CHATZIKOKOLAKIS, C. PALAMIDESSI. *A Framework for Analyzing Probabilistic Protocols and its Application to the Partial Secrets Exchange*, in "Theoretical Computer Science", To appear, 2007, http://www.lix.polytechnique.fr/~catuscia/papers/PartialSecrets/TCSreport.pdf.

[15] K. CHATZIKOKOLAKIS, C. PALAMIDESSI, P. PANANGADEN. *Anonymity Protocols as Noisy Channels*, in "Information and Computation", To appear, 2007, http://www.lix.polytechnique.fr/~catuscia/papers/Anonymity/Channels/full.pdf.

[16] Y. DENG, C. PALAMIDESSI. *Axiomatizations for probabilistic finite-state behaviors*, in "Theoretical Computer Science", vol. 373, n^o 1-2, 2007, p. 92–114, http://www.lix.polytechnique.fr/~catuscia/papers/Prob_Axiom/tcs.pdf.

[17] R. LANOTTE, A. MAGGIOLO-SCHETTINI, A. TROINA. *Parametric probabilistic transition systems for system design and analysis*, in "Formal Aspects of Computing", vol. 19, n^o 1, 2007, p. 93–109.

[18] I. PHILLIPS, M. G. VIGLIOTTI, C. PALAMIDESSI. *Expressiveness via Leader Election Problems*, in "Theoretical Computer Science", To appear, 2007.

Publications in Conferences and Workshops

[19] J. ARANDA, C. D. GIUSTO, M. NIELSEN, F. VALENCIA. *CCS with Replication in the Chomsky Hierarchy: The Expressive Power of Divergence*, in "Proc. of The Fifth ASIAN Symposium on Programming Languages (APLAS'07)", LNCS, To appear, Springer, 2007, http://www.lix.polytechnique.fr/~fvalenci/papers/aplas07.pdf.

[20] J. ARANDA, C. D. GIUSTO, C. PALAMIDESSI, F. VALENCIA. *Expressiveness of Recursion, Replication and Scope Mechanisms in Process Calculi*, in "Postproceedings of the 5th International Symposium on Formal Methods for Components and Objects (FMCO'06)", F. DE BOER, M. BONSANGUE (editors), LNCS, To appear, Springer, 2007.

[21] D. CACCIAGRANO, F. CORRADINI, J. ARANDA, F. VALENCIA. *Persistence and Testing Semantics in the Asynchronous Pi Calculus*, in "Proc. of 14th International Workshop on Expressiveness of Concurrency, (EXPRESS'07)", R. AMADIO, T. HILDENBRANDT (editors), ENTCS, To appear, Elsevier, 2007, http://www.lix.polytechnique.fr/~fvalenci/papers/express07.pdf.

[22] D. CACCIAGRANO, F. CORRADINI, C. PALAMIDESSI. *Fair* Π, in "Proceedings of the 13th International Workshop on Expressiveness in Concurrency (EXPRESS'06)", Electronic Notes in Theoretical Computer Science, vol. 175 (3), Elsevier Science B.V., 2007, p. 3–26, http://www.lix.polytechnique.fr/~catuscia/papers/Diletta/FairPi/express06.pdf.

[23] K. CHATZIKOKOLAKIS, C. PALAMIDESSI. *Making Random Choices Invisible to the Scheduler*, in "Proceedings of CONCUR'07", L. CAIRES, V. T. VASCONCELOS (editors), Lecture Notes in Computer Science, vol. 4703, Springer, 2007, p. 42–58, http://www.lix.polytechnique.fr/~catuscia/papers/Scheduler/report.pdf.

[24] K. CHATZIKOKOLAKIS, C. PALAMIDESSI, P. PANANGADEN. *Probability of Error in Information-Hiding Protocols*, in "Proceedings of the 20th IEEE Computer Security Foundations Symposium (CSF20)", IEEE Computer Society, 2007, p. 341-354, http://www.lix.polytechnique.fr/~catuscia/papers/ProbabilityError/full.pdf.

[25] A. CREDI, M. GARAVELLI, C. LANEVE, S. PRADALIER, S. SILVI, G. ZAVATTARO. *Modelization and Simulation of Nano Devices in* $\mathtt{nano}\kappa$ *Calculus*, in "Proceedings of the International Conference on Computational Methods in Systems Biology, (CMSB)", M. CALDER, S. GILMORE (editors), Lecture Notes in Computer Science, vol. 4695, Springer, 2007, p. 168–183.

[26] Y. DENG, C. PALAMIDESSI, J. PANG. *Weak Probabilistic Anonymity*, in "Proceedings of the 3rd International Workshop on Security Issues in Concurrency (SecCo)", Electronic Notes in Theoretical Computer Science, vol. 180 (1), Elsevier Science B.V., 2007, p. 55–76, http://www.lix.polytechnique.fr/~catuscia/papers/Anonymity/report_wa.pdf.

[27] M. FALASCHI, C. OLARTE, C. PALAMIDESSI, F. D. VALENCIA. *Declarative Diagnosis of Temporal Concurrent Constraint Programs*, in "Proceedings of The 23rd International Conference in Logic Programming (ICLP'07)", V. DAHL, I. NIEMELÄ (editors), Lecture Notes in Computer Science, vol. 4670, Springer, 2007, p. 271–285, http://www.lix.polytechnique.fr/~catuscia/papers/Carlos/iclp07.pdf.

[28] J. GOUBAULT-LARRECQ, C. PALAMIDESSI, A. TROINA. *A Probabilistic Applied Pi-Calculus*, in "Proceedings of the 5th Asian Symposium on Programming Languages and Systems (APLAS'07)", LNCS, To appear, Springer, 2007, http://www.lix.polytechnique.fr/~catuscia/papers/Angelo/aplas.pdf.

[29] J. GUTIERREZ, J. PEREZ, C. RUEDA, FRANK D. VALENCIA. *Timed Concurrent Constraint Programming for Analyzing Biological Systems.*, in "Proceedings of Workshop on Membrane Computing and Biologically Inspired Process Calculi.", Electronic Notes in Theoretical Computer Science, vol. 171 (2), Elsevier Science B.V., 2007, p. 117–137.

[30] G. NORMAN, C. PALAMIDESSI, D. PARKER, P. WU. *Model checking the probabilistic pi-calculus*, in "4th International Conference on the Quantitative Evaluation of SysTems (QEST)", IEEE Computer Society Press, 2007, p. 169-178, http://www.lix.polytechnique.fr/~catuscia/papers/Wu/qest1.pdf.

[31] C. OLARTE, C. PALAMIDESSI, F. D. VALENCIA. *Universal Timed Concurrent Constraint Programming*, in "Proceedings of the 23rd International Conference in Logic Programming (ICLP'07)", V. DAHL, I. NIEMELÄ (editors), Lecture Notes in Computer Science, vol. 4670, Springer, 2007, p. 464–465, http://www.lix.polytechnique.fr/~catuscia/papers/Carlos/iclp07DC.pdf.

[32] C. OLARTE, F. D. VALENCIA. *Universal Concurrent Constraint Programing: Symbolic Semantics and Applications to Security*, in "Proceedings of the 23rd ACM Symposium on Applied Computing (SAC)", To appear, ACM, 2007, http://www.lix.polytechnique.fr/~colarte/colarte/Publications_files/sac08.pdf.

[33] P. WU, C. PALAMIDESSI, H. LIN. *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.

References in notes

[34] M. ABADI, C. FOURNET. *Mobile Values, New Names, and Secure Communication*, in "28th Annual Symposium on Principles of Programming Languages (POPL)", ACM, January 2001, p. 104–115.

[35] M. BHARGAVA, C. PALAMIDESSI. *Probabilistic Anonymity*, in "Proceedings of CONCUR", M. ABADI, L. DE ALFARO (editors), Lecture Notes in Computer Science, vol. 3653, Springer, 2005, p. 171–185, http://www.lix.polytechnique.fr/~catuscia/papers/Anonymity/report.pdf.

[36] G. BOUDOL. *Asynchrony and the π-calculus (Note)*, Rapport de Recherche, n^{o} 1702, INRIA, Sophia-Antipolis, 1992, http://hal.inria.fr/inria-00076939.

[37] E. BRINKSMA, A. RENSINK, W. VOGLER. *Fair Testing*, in "Proceedings of the 6th International Conference on Concurrency Theory (CONCUR)", I. LEE, S. A. SMOLKA (editors), Lecture Notes in Computer Science, vol. 962, Springer-Verlag, 1995, p. 313–327.

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

[39] G. COSTA, C. STIRLING. *A Fair Calculus of Communicating Systems*, in "Acta Informatica", vol. 21, 1984, p. 417–441.

[40] G. COSTA, C. STIRLING. *Weak and Strong Fairness in CCS*, in "Information and Computation", vol. 73, n^o 3, June 1987, p. 207–244.

[41] S. EVEN, O. GOLDREICH, A. LEMPEL. *A randomized protocol for signing contracts*, in "Commun. ACM", vol. 28, n^o 6, 1985, p. 637–647.

[42] O. M. HERESCU, C. PALAMIDESSI. *Probabilistic Asynchronous π-Calculus*, in "Proceedings of FOSSACS 2000 (Part of ETAPS 2000)", J. TIURYN (editor), Lecture Notes in Computer Science, vol. 1784, Springer, 2000, p. 146–160, http://www.lix.polytechnique.fr/~catuscia/papers/Prob_asy_pi/fossacs.ps.

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

[44] K. HONDA, M. TOKORO. *An Object Calculus for Asynchronous Communication*, in "Proceedings of the European Conference on Object-Oriented Programming (ECOOP)", P. AMERICA (editor), Lecture Notes in Computer Science, vol. 512, Springer, 1991, p. 133–147.

[45] M. NAOR, B. PINKAS, R. SUMNER. *Privacy preserving auctions and mechanism design*, in "Proceedings of the 1st ACM Conference on Electronic Commerce", ACM Press, 1999, p. 129–139.

[46] V. NATARAJAN, R. CLEAVELAND. *Divergence and Fair Testing*, in "Proceedings of the 22nd International Colloquium on Automata, Languages and Programming (ICALP)", Z. FÜLÖP, F. GÉCSEG (editors), Lecture Notes in Computer Science, vol. 944, Springer, 1995, p. 648–659.

[47] R. D. NICOLA, M. C. B. HENNESSY. *Testing equivalences for processes*, in "Theoretical Computer Science", vol. 34, n^o 1-2, 1984, p. 83–133.

[48] C. PALAMIDESSI, O. M. HERESCU. *A randomized encoding of the π-calculus with mixed choice*, in "Theoretical Computer Science", vol. 335, n^o 2-3, 2005, p. 373-404, http://www.lix.polytechnique.fr/~catuscia/papers/prob_enc/report.pdf.

[49] C. PALAMIDESSI. *Anonymity in probabilistic and nondeterministic system*, in "Proceedings of the Workshop on “Algebraic Process Calculi: The First Twenty Five Years and Beyond”, Bertinoro, Italy", Electronic Notes in Theoretical Computer Science, To appear, Elsevier Science B.V., 2005, http://www.lix.polytechnique.fr/~catuscia/papers/Anonymity/Bertinoro/paper.pdf.

[50] C. PALAMIDESSI. *Probabilistic and nondeterministic aspects of Anonymity*, in "Proceedings of the 21st Conference on the Mathematical Foundations of Programming Semantics, Birmingham, UK", Electronic Notes in Theoretical Computer Science, to appear, Elsevier Science B.V., 2005, http://www.lix.polytechnique.fr/~catuscia/papers/Anonymity/MFPS/paper.pdf.

[51] M. O. RABIN. *How to exchange secrets by oblivious transfer*, in "Technical Memo TR-81, Aiken Computation Laboratory, Harvard University", 1981.

[52] P. YANG, C. R. RAMAKRISHNAN, S. A. SMOLKA. *A logical encoding of the pi-calculus: model checking mobile processes using tabled resolution*, in "International Journal on Software Tools for Technology Transfer", vol. 6, n^o 1, 2004, p. 38–66.