IN PARTNERSHIP WITH:
CNRS

Ecole Polytechnique

Activity Report 2012

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, Constraints, Information Theory, Quantitative Information Flow, Privacy

Creation of the Project-Team: January 01, 2008 .

1. Members

Research Scientists

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

External Collaborator

Jérémy Dubreil [Monoidics Ltd, UK. He has been a postdoc in Comète from December 1, 2009, until February 28, 2011]

PhD Students

Andrés Aristizábal [Grant DGA/CNRS. From October 1, 2009, until September 30, 2012]
Nicolás Bordenabe [Grant Inria/DGA. Since October 1, 2011]
Luis Fernando Pino Duque [Grant Inria/DGA. Since October 1, 2011]
Ivan Gazeau [Grant ANR (CCP). Co-supervised by Dale Miller, Inria. Since October 1, 2009]
Sophia Knight [Grant Inria-CORDIS. Since September 15, 2010]
Marco Stronati [Grant EDX "Monge". Since October 1, 2012]
Lili Xu [Grant ANR (PANDA). Co-supervised by Huimin Li, Chinese Academy of Science, Bejijing. Since October 15, 2011]

Post-Doctoral Fellows

Miguel Andrés [Grant QUALCOMM. From November 27, 2010, until September 30, 2012]
Ehab ElSalamouny [Grant Inria. From October 1, 2011, until November 11, 2012]
Marco Giunti [Grant ERCIM. From March 1, 2011, until February 29, 2012]
Sardaouna Hamadou [Grant ANR (Panda). From November 1, 2011]
Matteo Mio [Grant ERCIM. From March 1, 2012]

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* [41]). *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 of the Year

Mário Alvim, an ex PhD student of Comète who defended his thesis in October 2011, has been nominated for the "Prix de thèse ParisTech 2012".

3. Scientific Foundations

3.1. Probability and information theory

Participants: Miguel Andrés, Nicolás Bordenabe, Konstantinos Chatzikokolakis, Ehab ElSalamouny, Sardaouna Hamadou, Catuscia Palamidessi, Marco Stronati.

Much of the research of Comète 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 [43] 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: Miguel Andrés, Nicolás Bordenabe, Konstantinos Chatzikokolakis, 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 [39] 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 [40] 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. models. It is built on XSB, 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 2005-2007, in the context of the collaboration between the teams Comète and PRISM under the Inria/ARC Project ProNoBis. It is based on the papers [44] and [42].

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; we are exploring several topics, described below.

6.1.1. Measuring information leakage

A fundamental concern in computer security is to control information flow, whether to protect confidential information from being leaked, or to protect trusted information from being tainted. In view of the pragmatic difficulty of preventing undesirable flows completely, there is now much interest in theories that allow information flow to be quantified, so that "small" leaks can be tolerated. In [19] we introduced g-leakage, a rich generalization of the min-entropy model of quantitative information flow. In g-leakage, the benefit that an adversary derives from a certain guess about a secret is specified using a gain function g. Gain functions allow a wide variety of operational scenarios to be modeled, including those where the adversary benefits from guessing a value close to the secret, guessing a part of the secret, guessing a property of the secret, or guessing the secret within some number of tries. We proved important properties of g-leakage, including bounds between min-capacity, g-capacity, and Shannon capacity. We also showed a deep connection between a strong leakage ordering on two channels, C1 and C2, and the possibility of factoring C1 into C2 C3 , for some C3 . Based on this connection, we proposed a generalization of the Lattice of Information from deterministic to probabilistic channels.

6.1.2. Interactive systems

In [12] 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.

6.1.3. Unlinkability

Unlinkability is a privacy property of crucial importance for several systems (such as RFID or voting systems). Informally, unlinkability states that, given two events/items in a system, an attacker is not able to infer whether they are related to each other. However, in the literature we find several definitions for this notion, which are apparently unrelated and shows a potentially problematic lack of agreement. In [22] we shed new light on unlinkability by comparing different ways of defining it and showing that in many practical situations the various definitions coincide. It does so by (a) expressing in a unifying framework four definitions of unlinkability from the literature (b) demonstrating how these definitions are different yet related to each other and to their dual notion of "inseparability" and (c) by identifying conditions under which all these definitions become equivalent. We argued that the conditions are reasonable to expect in identification systems, and we prove that they hold for a generic class of protocols.

6.1.4. A compositional method to compute the sensitivity of differentially private queries

Differential privacy is a modern approach in privacy-preserving data analysis to control the amount of information that can be inferred about an individual by querying a database. The most common techniques are based on the introduction of probabilistic noise, often defined as a Laplacian parametric on the sensitivity of the query. In order to maximize the utility of the query, it is crucial to estimate the sensitivity as precisely as possible.

In [28] we considered relational algebra, the classical language for expressing queries in relational databases, and we proposed a method for computing a bound on the sensitivity of queries in an intuitive and compositional way. We used constraint-based techniques to accumulate the information on the possible values for attributes provided by the various components of the query, thus making it possible to compute tight bounds on the sensitivity.

6.1.5. A differentially private mechanism of optimal utility for a region of priors

Differential privacy (already introduced in the previous section) is usually achieved by using mechanisms that add random noise to the query answer. Thus, privacy is obtained at the cost of reducing the accuracy, and therefore the utility, of the answer. Since the utility depends on the user's side information, commonly modeled as a prior distribution, a natural goal is to design mechanisms that are optimal for every prior. However, it has been shown in the literature that such mechanisms do not exist for any query other than counting queries.

Given the above negative result, in [38] we considered the problem of identifying a restricted class of priors for which an optimal mechanism does exist. Given an arbitrary query and a privacy parameter, we geometrically characterized a special region of priors as a convex polytope in the priors space. We then derived upper bounds for utility as well as for min-entropy leakage for the priors in this region. Finally we defined what we call the tight-constraints mechanism and we discussed the conditions for its existence. This mechanism has the property of reaching the bounds for all the priors of the region, and thus it is optimal on the whole region.

6.1.6. Differential privacy with general metrics

Differential privacy, already described above, is a formal privacy guarantee that ensures that sensitive information relative to individuals cannot be easily inferred by disclosing answers to aggregate queries. If two databases are adjacent, i.e. differ only for an individual, then querying them should not allow to tell them apart by more than a certain factor. The transitive application of this property induces a bound also on the

distinguishability of two generic databases, which is determined by their distance on the Hamming graph of the adjacency relation.

In [37] we lifted the restriction relative to the Hamming graphs and we explored the implications of differential privacy when the indistinguishability requirement depends on an arbitrary notion of distance. We showed that we can express, in this way, (protection against) kinds of privacy threats that cannot be naturally represented with the standard notion. We gave an intuitive characterization of these threats in terms of Bayesian adversaries, which generalizes the characterization of (standard) differential privacy from the literature. Next, we revisited the well-known result on the non-existence of universally optimal mechanisms for any query other than counting queries. We showed that in our setting, for certain kinds of distances, there are many more queries for which universally optimal mechanisms exist: Notably sum, average, and percentile queries. Finally, we showed some applications in various domains: statistical databases where the units of protection are groups (rather than individuals), geolocation, and smart metering.

6.1.7. Privacy for location-based systems

The growing popularity of location-based systems, allowing unknown/untrusted servers to easily collect and process huge amounts of users' information regarding their location, has recently started raising serious concerns about the privacy of this kind of sensitive information. In [36] we studied geo-indistinguishability, a formal notion of privacy for location-based systems that protects the exact location of a user, while still allowing approximate information - typically needed to obtain a certain desired service - to be released.

Our privacy definition formalizes the intuitive notion of protecting the user's location within a radius r with a level of privacy that depends on r. We presented three equivalent characterizations of this notion, one of which corresponds to a generalized version [37] of the well-known concept of differential privacy. Furthermore, we presented a perturbation technique for achieving geo-indistinguishability by adding controlled random noise to the user's location, drawn from a planar Laplace distribution. We demonstrated the applicability of our technique through two case studies: First, we showed how to enhance applications for location-based services with privacy guarantees by implementing our technique on the client side of the application. Second, we showed how to apply our technique to sanitize location-based sensible information collected by the US Census Bureau.

6.1.8. Compositional analysis of information hiding

Systems concerned with information hiding often use randomization to obfuscate the link between the observables and the information to be protected. The degree of protection provided by a system can be expressed in terms of the probability of error associated to the inference of the secret information. In [15] we considered a probabilistic process calculus to specify such systems, and we studied how the operators affect the probability of error. In particular, we characterized constructs that have the property of not decreasing the degree of protection, and that can therefore be considered safe in the modular construction of these systems. As a case study, we applied these techniques to the Dining Cryptographers, and we derived a generalization of Chaum's strong anonymity result.

In [29], a similar framework was proposed for reasoning about the degree of differential privacy provided by such systems. In particular, we investigated the preservation of the degree of privacy under composition via the various operators. We illustrated our idea by proving an anonymity-preservation property for a variant of the Crowds protocol for which the standard analyses from the literature are inapplicable. Finally, we made some preliminary steps towards automatically computing the degree of privacy of a system in a compositional way.

6.1.9. Anonymous and route-secure communication systems

Incentives to Cooperation. Anonymity systems have a broad range of users, ranging from ordinary citizens who want to avoid being profiled for targeted advertisements, to companies trying to hide information from their competitors, to entities requiring untraceable communication over the Internet. With these many potential users, it would seem that anonymity services based on a consumer/provider users will naturally be well-resourced and able to operate efficiently. However, cooperation cannot be taken for granted. Current deployed

systems show that some users will indeed act selfishly, and only use the system to send their messages whilst ignoring the requests to forward others' messages. Obviously, with not enough cooperative users, the systems will hardly operate at all, and will certainly not be able to afford adequate anonymity guarantees. It is therefore vital that these systems are able to deploy incentives to encourage users' cooperation and so make the anonymity provision effective. Some interesting approaches to achieve that have been proposed, such as make running relays easier and provide better forwarding performance.

To evaluate whether these approaches are effective, we need a framework which empowers us to analyze them, as well as provide guidelines and some mechanism design principles for incentive schemes. This much we have provided in [30], exploiting notions and techniques from Game Theory. We proposed a game theoretic framework and used it to analyze users' behaviours and also predict what strategies users will choose under different circumstances and according to their exact balance of preferences among factors such as anonymity, performance (message delivery time) and cost. Significantly, we also used the model to assess the effectiveness of the gold-star incentive mechanism, which was introduced in Tor network to encourage users to act as cooperative relays, and thus enhance the service performance for well-behaved forwarders.

Trust in anonymity networks. Trust metrics are used in anonymity networks to support and enhance reliability in the absence of verifiable identities, and a variety of security attacks currently focus on degrading a user's trustworthiness in the eyes of the other users. In [16] we have presented an enhancement of the Crowds anonymity protocol via a notion of trust which allows crowd members to route their traffic according to their perceived degree of trustworthiness of each other member of the crowd. Such trust relations express a measure of an individual's belief that another user may become compromised by an attacker, either by a direct attempt to corrupt or by a denial-of-service attack. Our protocol variation has the potential of improving the overall trustworthiness of data exchanges in anonymity networks, which cannot normally be taken for granted in a context where users are actively trying to conceal their identities. Using such formalization, in the paper we have then analyzed quantitatively the privacy properties of the protocol under standard and adaptive attacks.

6.2. Foundations of Concurrency

Distributed systems have changed substantially in the recent past with the advent of phenomena like social networks and cloud computing. In the previous incarnation of distributed computing the emphasis was on consistency, fault tolerance, resource management and related topics; these were all characterized by *interaction between processes*. Research proceeded along two lines: the algorithmic side which dominated the Principles Of Distributed Computing conferences and the more process algebraic approach epitomized by CONCUR where the emphasis was on developing compositional reasoning principles. What marks the new era of distributed systems is an emphasis on managing access to information to a much greater degree than before.

The *Concurrent constraint programming (ccp)* paradigm focuses on information access and therefore it is suited for this new era of concurrent systems. Ccp singles out the fundamental aspects of asynchronous systems whose agents (or processes) evolve by accessing information in a global medium. In the works [20], [21], [31], [26] described below we developed algorithms and extended the foundations of ccp.

6.2.1. Spatial and Epistemic Modalities for Constraint-based Calculi

Epistemic concepts were crucial in distributed computing as was realized in the mid 1980s with Halpern and Moses' groundbreaking paper on common knowledge. This led to a flurry of activity in the next few years with many distributed protocols being understood from an epistemic point of view. The impact of epistemic ideas in the concurrency theory community was slower in coming. We believe that epistemic ideas need to be exploited more by concurrency theorists and we did so in the following works.

In [26] we introduced spatial and epistemic process calculi for reasoning about spatial information and knowledge distributed among the agents of a system. We also introduced domain-theoretical structures to represent spatial and epistemic information. Finally we provided operational and denotational techniques for reasoning about the potentially infinite behaviour of spatial and epistemic processes. We also gave compact

representations of infinite objects that can be used by processes to simulate announcements of common knowledge and global information. We also developed an interpreter of these calculi in [31].

6.2.2. Bisimilarity for Constraint-based Calculi

Bisimilarity is a standard behavioural equivalence in concurrency theory, but a well-behaved notion of bisimilarity for ccp has been proposed only recently. 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 [20] we 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.

In [20] we only studied the strong version of bisimilarity. Weak bisimiliarity is obtained from the strong case by taking into account only the actions that are observable in the system. Typically, the standard partition refinement can also be used for deciding weak bisimilarity simply by using Milner's reduction from weak to strong bisimilarity; a technique referred to as saturation. In [21] we showed that, because of its involved labeled transitions, the above-mentioned saturation technique does not work for ccp. We also gave an alternative reduction from weak ccp bisimilarity to the strong one that allows us to use the ccp partition refinement algorithm for deciding this equivalence.

In the more traditional setting of the pi-calculus we have also proposed an approach to restrict access to information.

6.2.3. Locality in the Pi-Calculus

In [25] we enriched the pi-calculus with an operator for confidentiality (hide), whose main effect is to restrict the access to the object of the communication, thus representing confidentiality in a natural way. The hide operator is meant for local communication, and it differs from new in that it forbids the extrusion of the name and hence has a static scope. Consequently, a communication channel in the scope of a hide can be implemented as a dedicated channel, and it is more secure than one in the scope of a new. To emphasize the difference, we introduced a spy context that represents a side-channel attack and breaks some of the standard security equations for new. To formally reason on the security guarantees provided by the hide construct, we also introduced an observational theory and establish stronger equivalences by relying on a proof technique based on bisimulation semantics.

6.2.4. Foundations of Probabilistic Concurrent Systems

In [17] we have solved an open problem in the literature by proving that two known semantics for the probabilistic mu-calculus, a denotational semantics and a two-player stochastic game semantics, coincide on all models.

In [18] we have improved the result of [17] by introducing a new logic called probabilistic mu-calculus with independent product. We have proved that two semantics coincide in all models: a denotational semantics and a two-player game semantics based on a novel class of concurrent games. Furthermore, we have shown how the new logic is strictly more expressive than the other. This allows the encoding of other important temporal logics for probabilistic concurrent systems such as PCTL.

In [27] we have introduced a proof system designed for supporting human-aided verification of properties (expressed as probabilistic mu-calculus formulas ([17]) of concurrent probabilistic processes described by SOS-style operational semantics.

6.2.5. Interference metrics for Mobile ad-hoc networks (MANETs)

Mobile ad-hoc networks consist of a collection of nodes that communicate with each other through wireless links without a pre-established networking infrastructure. A common feature of most of these networks is free node mobility. Each device will therefore change its links to other devices frequently. These frequent changes in the network topology can cause the nodes to continuously enter and exit each other transmission area. Hence, highly dynamic routing algorithms are needed to ensure the connectivity. Moreover, mobile devices

may have strict requirements on the energy consumption because their expected life-time often depends on the energy stored in a battery or other exhaustible power sources. For these reasons, finding a good trade-off between network connectivity, power saving and interference reduction is one of the most critical challenges in managing mobile ad hoc networks. In [23], we have proposed an effective framework for analysing protocol connectivity and measuring the level of interference and, based on that for developing novel interference-aware communication strategies. Though other models exist in the literature, to our best knowledge, our framework is the most comprehensive and effective for the behavioral analysis and a quantitative assessment of interference for wireless networks in the presence of node mobility.

7. Partnerships and Cooperations

7.1. National Initiatives

7.1.1. ANR projects

7.1.1.1. ANR-09-BLAN-0169-01

Project acronym: PANDA

Project title: Analysis of Parallelism and Distribution

Duration: October 2009 - March 2013

URL: http://lipn.univ-paris13.fr/~mazza/Panda/

Coordinator: Catuscia Palamidessi, Inria Saclay

Other PI's and partner institutions: Dale Miller, EPIs Parsifal at Inria Saclay. Emmanuel Haucourt, CEA Saclay. Damiano Mazza, Pôle Parisien (ENS Cachan, Paris VII and Paris XIII). Emmanuel Godard, Pôle Méditerranéen (ENS Lyon and the University of Marseille). Jean Souyris, Airbus.

Abstract: The aim of PANDA is to bring together different mathematical models of parallel and concurrent computation (geometric models, rewriting theory, higher category theory, stochastic processes), along with theoretical frameworks for static analysis (spatial logics, proof construction), in order to guide the development of software tools that meet industrial needs of program specification and verification (in particular, fault detection of parallel programs involved in avionics).

7.1.1.2. ANR-09-BLAN-0345-02

Project acronym: CCP

Project title: Confidence, Proof and Probabilities

Duration: October 2009 - March 2013

URL: http://www.lix.polytechnique.fr/~bouissou/cpp/

Coordinator: Jean Goubault-Larrecq, ENS Cachan

Other PI's and partner institutions: Catuscia Palamidessi, Inria. Olivier Bouissou, CEA LIST. Gilles Fleury, Supelec SSE. Michel Kieffer, Supelec L2S.

Abstract: In the context of proofs of safety properties for critical software, The CPP project proposes to study the joint use of probabilistic and formal (deterministic) semantics and analysis methods, in a way to improve the applicability and precision of static analysis methods on numerical programs.

7.1.2. Large-scale initiatives

Project acronym: CAPPRIS

Project title: Collaborative Action on the Protection of Privacy Rights in the Information Society

Duration: October 2011 - September 2015

Coordinator: Daniel Le Metayer, Inria Grenoble

Other partner institutions: The project involves four Inria research centers (Saclay, Saphia-Antipolis, Rennes and Grenoble), CNRS-LAAS, Eurecom and the university of Namur. Besides computer scientists, the consortium also includes experts in sociology and in law, thus covering the complementary areas of expertise required to reach the objectives.

Abstract: The goal of this project is to study the challenges related to privacy in the modern information society, trying to consider not only the technical, but also the social and legal ones, and to develop methods to enhance the privacy protection.

7.2. European Initiatives

7.2.1. FP7 Projects

Program: FP7-PEOPLE-2011-IRSES

Project acronym: MEALS

Project title: Mobility between Europe and Argentina applying Logic to Systems

Duration: October 2011 - September 2005

URL: http://www.meals-project.eu/

Coordinator: Holger Hermans, Saarland University, Germany

Other partner institutions: Rheinisch-Westfälische Technische Hochschule Aachen, Germany. Technische Universität Dresden, Germany. Inria, France. Imperial College of Science, Technology and Medicine, UK, University of Leicester, UK. Technische Universiteit Eindhoven, NL. Universidad Nacional de Cordoba, AR. Universidad de Buenos Aires, AR. Instituto Tecnologico de Buenos Aires, AR. Universidad Nacional de Rio Cuarto, AR.

Abstract: In this project we focus on three aspects of formal methods: specification, verification, and synthesis. We consider the study of both qualitative behavior and quantitative behavior (extended with probabilistic information). We aim to study formal methods in all their aspects: foundations (their mathematical and logical basis), algorithmic advances (the conceptual basis for software tool support) and practical considerations (tool construction and case studies).

7.3. International Initiatives

7.3.1. International Partners

Geoffrey Smith. School of Computing and Information Sciences, Florida International University, USA.

Vladimiro Sassone. School of Electronics and Computer Science, University of Southampton, UK.

Camilo Rueda. Department of Computer Science, Pontificia Universidad Javeriana, Colombia.

7.3.2. Participation in International Programs

Program: ANR Blanc International

Project acronym: LOCALI

Project title: Logical Approach to Novel Computational Paradigms

Duration: October 2011 - September 2015

Coordinator: Gilles Dowek, Inria Rocquencourt

Other PI's and partner institutions: Catuscia Palamidessi, Inria Saclay. Thomas Erhard, Paris VII. Ying Jiang , Chinese Academy of Science in Beijin (China).

Abstract: This project aims at exploring the interplays between logic and sequential/distributed computation in formalisms like the lambda calculus and the π calculus. Going back to the fundamentals of the definitions of these calculi, the project plans to design new programming languages and proof systems via a logical approach.

7.4. International Research Visitors

7.4.1. Visits of International Scientists

Carlos Olarte. Associate professor at the Pontificia Universidad Javeriana, Colombia. He visited for one month in July 2012, funded by the Ecole Polytechnique.

Moreno Falaschi. Full professor at the Università di Siena, Italy. He visited for one month in June 2012, funded by the Ecole Polytechnique.

Elaine Pimentel. Associate professor at the Universidade Federal de Minas Gerais, Belo Horizonte, Brazil. She visited for one month in July 2012, funded by the Ecole Polytechnique/Digiteo.

Linda Brodo. Assistant professor at the Università di Sassari, Italy. She visited for one month in June 2012, funded by the Ecole Polytechnique/Digiteo.

Vladimiro Sassone. Full professor at the University of Southampton, UK. He visited for two months in October and November 2012, funded by the Ecole Polytechnique/Digiteo.

Camilo Rueda. Full professor at the Pontificia Universidad Javeriana,Colombia. He visited for two months in October and November 2012, funded by the Ecole Polytechnique.

7.4.2. Internships

Name: Lili Xu
Duration: From October 2011 until October 2012)
Subject: Compositionality of privacy on a probabilistic process calculus
Institution: Chinese Academy of Sciences of Beijin (China)
Support: ANR project PANDA, Inria, and Chinese Academy of Sciences

Name: Marco Stronati
Duration: From October 2011 until March 2013
Subject: Compositional analysis of queries' sensitivity
Institution: University of Pisa, Italy
Support: Ecole Polytechnique and University of Pisa

Name: Fernán Martinelli
Duration: From September 2012 until March 2013
Subject: Computation of bounds on the information flow
Institution: University of Rio Cuarto, Argentina
Support: FP7 project MEALS

Name: Michela Paolini
Duration: From September 2012 until December 2012
Subject: Compositionality of privacy on a probabilistic process calculus.
Institution: IMT Institute for Advanced Studies, Lucca, Italy
Support: Grant from IMT

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 Mathematical Structures in Computer Science, 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 Geoffrey Smith) of the special issue of Mathematical Structures in Computer Science dedicated to Quantitative Information Flow.

Co-editor (with Samson Abramsky and Michael Mislove) of the special issue of Theoretical Computer Science dedicated to selected papers of MFPS XXV. [33]

Co-editor (with Sebastian Mödersheim) of the proceedings of TOSCA 2011, Theory of Security and Applications. [34]

Co-editor (with Mark Ryan) of the proceedings of TGC 2012, Trustworthy Global Computing. [35]

Frank D. Valencia is:

Co-editor of the special issue of Mathematical Structures in Computer Science dedicated to the 18th International Workshop on Expressiveness in Concurrency.

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.

Frank D. Valencia member of:

The steering committee of the International Workshop in Concurrency EXPRESS. Since 2010.

8.1.3. Invited Talks

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

Grande Region Security and Reliability Day. Nancy, France. March 2012.

COW 2012. The 19th CREST Open Workshop on *Interference and Dependence* on 30st April - 1st May 2012.

VECoS 2012. 6th International Workshop on Verification and Evaluation of Computer and Communication Systems. CNAM, Paris, France. August 27-28, 2012.

8.1.4. Organization of workshops and conferences

Catuscia Palamidessi has served as PC co-chair (together with Mark Ryan) of TGC 2012, the 7th International Symposium on Trustworthy Global Computing. Newcastle, UK, 7-8 September 2012.

Catuscia Palamidessi has co-organized (together with Boris Köpf and Pasquale Malacaria) the Dagstuhl seminar on Quantitative Security Analysis. Dagstuhl, Germany, 25-30 November 2012.

8.1.5. Participation in program committees

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

TGC 2013. The 8th International Symposium on Trustworthy Global Computing. Buenos Aires, Argentina, 30-31 August 2013.

ICALP 2013 Track B. The 40th International Colloquium on Automata, Languages and Programming. Riga, Latvia, 8-12 July 2013.

CSF 2013. The 26th IEEE Computer Security Foundations Symposium. Tulane University, New Orleans, Louisiana, USA, 26-28 June 2013.

LICS 2013. The Twenty-Eighth Annual ACM/IEEE Symposium on Logic in Computer Science. Tulane University, New Orleans, Louisiana, USA, 25-28 June 2013.

FOSSACS 2013. The 16th Int.l Conf. on Foundations of Software Science and Computation Structures. (Part of ETAPS 2013.) Rome, Italy, March 2013.

SOFSEM 2013. 39th International Conference on Current Trends in Theory and Practice of Computer Science. Špindlerův Mlýn, Czech Republic, January 26–31, 2013.

CARDIS 2012. The Eleventh Smart Card Research and Advanced Application Conference. Graz, Austria, 28-30 November 2012.

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.

Frank D. Valencia has been/is a member of the program committees of the following conferences and workshops:

CONCUR 2013. The 24th International Conference on Concurrency Theory. Buenos Aires, Argentina, 27-30 August 2013.

EXPRESS 2012: Combined 19th International Workshop on Expressiveness in Concurrency and 9th Workshop on Structural Operational Semantics. Newcastle upon Tyne, UK, 3 September 2012.

ICE 2012. The 5th International Workshop on Interaction and Concurrency Experience. Stockholm, Sweden, 16 June 2012.

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

ISPEC 2013: 9th International Conference on Information Security Practice and Experience.

QAPL 2013: 11th Workshop on Quantitative Aspects of Programming Languages.

ESOP 2012: 21th European Symposium on Programming.

TGC 2012: 7th International Symposium on Trustworthy Global Computing.

ISPEC 2012: 8th International Conference on Information Security Practice and Experience.

QAPL 2012: 10th Workshop on Quantitative Aspects of Programming Languages .

8.1.6. Participation in other committees

Catuscia Palamidessi serves in the following committees:

President of the selection committee for the EATCS Best Paper Award at the ETAPS conferences. Since 2006.

Member of the EAPLS PhD Award committee. Since 2010.

8.1.7. Organization of seminars

Frank D. Valencia, Luis Fernando Pino Duque, 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.

8.2. Service

Catuscia Palamidessi serves 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.

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.

Frank Valencia has served as:

Member of the Evaluation Committee of the LIX/Qualcomm postdoc grants for the year 2012.

8.3. 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.

Master: Miguel E. Andrés, Konstantinos Chatzikokolakis, and Catuscia Palamidessi have been teaching an advanced course on Quantitative Information Flow and on Differential Privacy at RIO 2012, the Summer School on Informatics Río Cuarto, Argentine. Total 20 hours. 13-18 February 2012.

Master. Frank D. Valencia has been teaching an advanced course on Process Modeling at Master Program in Computer Science of the Pontificia Universidad Javeriana de Cali, Colombia. Total 30 hours. A.Y. 2011-12.

PhD (2009-12) Andrés Aristizábal. Ecole Polytechnique. Grant CNRS/DGA. Title of the thesis: *Bisimulation Techniques and Algorithms for Concurrent Constraint Programming*. Defended on 17 October 2012. Co-supervised by Catuscia Palamidessi and Frank D. Valencia.

PhD in progress (2012-) Marco Stronati. Ecole Polytechnique. Grant EDX Monge. Co-supervised by Catuscia Palamidessi and Konstantinos Chatzikokolakis.

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

PhD in progress (2011-) Nicolás E. Bordenabe. Ecole Polytechnique. Grant Inria/DGA. Co-supervised by Catuscia Palamidessi and Konstantinos Chatzikokolakis.

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-) Ivan Gazeau. Ecole Polytechnique. Grant ANR. Co-supervised by Catuscia Palamidessi and Dale Miller.

8.3.1. PhD defenses

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

James Jerson Ortiz Vega (Universidad del Valle, Cali, Colombia). Title of the thesis: *Formal Methods for the Specification and Verification of Distributed and Timed Systems*. Advised Juan Francisco Dias Frias. Defended in September 2013.

Thomas Given-Wilson (University of Technology, Sydney, Australia). Title of the thesis: *Concurrent Pattern Unification*. Advised by Barry Jay. Defended in August 2012.

Jacopo Mauro (University of Bologna, Italy). PhD thesis reviewer. Title of the thesis: *Constraints meet Concurrency*. Advised by Maurizio Gabbrielli. Defended in April 2012.

Catuscia Palamidessi has been member of the committee at the HDR defense of Daniele Varacca, University of Paris VII, December 2012.

9. Bibliography

Major publications by the team in recent years

[1] M. ALVIM, K. CHATZIKOKOLAKIS, C. PALAMIDESSI, G. SMITH. *Measuring Information Leakage using Generalized Gain Functions*, in "Computer Security Foundations", Cambridge MA, United States, IEEE, 2012, p. 265-279 [*DOI :* 10.1109/CSF.2012.26], http://hal.inria.fr/hal-00734044.

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

[3] K. CHATZIKOKOLAKIS, C. PALAMIDESSI, P. PANANGADEN. *Anonymity Protocols as Noisy Channels*, in "Information and Computation", 2008, vol. 206, n° 2–4, p. 378–401 [*DOI :* 10.1016/J.IC.2007.07.003], http://hal.inria.fr/inria-00349225/en/.

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

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

[6] 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, http://www.lix.polytechnique.fr/~fvalenci/papers/fossacs04.pdf.

[7] 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.

[8] S. KNIGHT, C. PALAMIDESSI, P. PANANGADEN, F. D. VALENCIA. *Spatial and Epistemic Modalities in Constraint-Based Process Calculi*, in "CONCUR 2012 - Concurrency Theory - 23rd International Conference, CONCUR 2012", Newcastle upon Tyne, United Kingdom, September 2012, vol. 7454, p. 317-332 [*DOI :* 10.1007/978-3-642-32940-1], http://hal.inria.fr/hal-00761116.

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

[10] 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/.

Publications of the year

Doctoral Dissertations and Habilitation Theses

[11] A. ARISTIZÁBAL. *Techniques de Bisimulation et Algorithmes pour la Programmation Concurrente par Contraintes*, Ecole Polytechnique X, October 2012, http://hal.inria.fr/pastel-00756952.

Articles in International Peer-Reviewed Journals

[12] M. ALVIM, M. ANDRÉS, C. PALAMIDESSI. *Quantitative Information Flow in Interactive Systems*, in "Journal of Computer Security", 2012, vol. 20, n° 1, p. 3-50, http://hal.inria.fr/inria-00637356.

[13] F. CASSEZ, J. DUBREIL, H. MARCHAND. *Synthesis of opaque systems with static and dynamic masks*, in "Formal Methods in System Design", 2012, vol. 40, n° 1, p. 88-115 [*DOI :* 10.1007/s10703-012-0141-9], http://hal.inria.fr/hal-00662539.

[14] K. CHATZIKOKOLAKIS, S. KNIGHT, C. PALAMIDESSI, P. PANANGADEN. *Epistemic Strategies and Games on Concurrent Processes*, in "Transactions on Computational Logic", October 2012, vol. 13, n° 4, p. 28:1-28:35 [*DOI :* 10.1145/2362355.2362356], http://hal.inria.fr/inria-00637160.

[15] K. CHATZIKOKOLAKIS, C. PALAMIDESSI, C. BRAUN. *Compositional Methods for Information-Hiding*, in "Mathematical Structures in Computer Science", 2013, to appear, http://hal.inria.fr/hal-00760596.

[16] S. HAMADOU, V. SASSONE, M. YANG. *An analysis of trust in anonymity networks in the presence of adaptive attackers*, in "Mathematical Structures in Computer Science", 2013, page : to appear, http://hal.inria.fr/hal-00760437.

[17] M. MIO. *On the equivalence of game and denotational semantics for the probabilistic μ-calculus*, in "Logical Methods in Computer Science", June 2012, vol. 8, n° 2, http://hal.inria.fr/hal-00763454.

[18] M. MIO. *Probabilistic modal μ-calculus with independent product*, in "Logical Methods in Computer Science", November 2012, vol. 8, n° 4, http://hal.inria.fr/hal-00763451.

International Conferences with Proceedings

[19] M. ALVIM, K. CHATZIKOKOLAKIS, C. PALAMIDESSI, G. SMITH. *Measuring Information Leakage using Generalized Gain Functions*, in "Computer Security Foundations", Cambridge MA, United States, IEEE, 2012, p. 265-279 [*DOI :* 10.1109/CSF.2012.26], http://hal.inria.fr/hal-00734044.

[20] A. ARISTIZÁBAL, F. BONCHI, L. PINO, F. D. VALENCIA. *Partition Refinement for Bisimilarity in CCP*, in "27th ACM Symposium On Applied Computing", Trento, Italy, 2012, 6, http://hal.inria.fr/hal-00641408.

[21] A. ARISTIZÁBAL, F. BONCHI, L. PINO, F. D. VALENCIA. *Reducing Weak to Strong Bisimilarity in CCP*, in "Fifth Interaction and Concurrency Experience", Stockholm, Sweden, December 2012, p. 2-16, To appear in EPTCS [*DOI :* 10.4204/EPTCS.104], http://hal.inria.fr/hal-00761611.

[22] M. BRUSÓ, K. CHATZIKOKOLAKIS, S. ETALLE, J. DEN HARTOG. *Linking Unlinkability*, in "7th International Symposium on Trustworthy Global Computing (TGC)", Newcastle upon Tyne, United Kingdom, 2012, To appear, http://hal.inria.fr/hal-00760150.

[23] M. BUGLIESI, L. GALLINA, A. MARIN, S. ROSSI, S. HAMADOU. *Interference-Sensitive Preorders for MANETs*, in "Ninth International Conference on Quantitative Evaluation of Systems, QEST 2012", London, United Kingdom, IEEE Computer Society, 2012, p. 189-198, http://hal.inria.fr/hal-00760455.

[24] I. GAZEAU, D. MILLER, C. PALAMIDESSI. *A non-local method for robustness analysis of floating point programs*, in "QAPL - Tenth Workshop on Quantitative Aspects of Programming Languages", Tallinn, Estonia, M. MASSINK, H. WIKLICKY (editors), March 2012 [*DOI :* 10.4204/EPTCS.85.5], http://hal.inria.fr/hal-00665995.

[25] M. GIUNTI, C. PALAMIDESSI, F. D. VALENCIA. *Hide and New in the Pi-Calculus*, in "Combined 19th International Workshop on Expressiveness in Concurrency and 9th Workshop on Structured Operational Semantics (EXPRESS/SOS 2012)", New Castle upon Tyne, United Kingdom, August 2012, p. 65-80 [*DOI :* 10.4204/EPTCS.89], http://hal.inria.fr/hal-00761118.

[26] S. KNIGHT, C. PALAMIDESSI, P. PANANGADEN, F. D. VALENCIA. *Spatial and Epistemic Modalities in Constraint-Based Process Calculi*, in "CONCUR 2012 - Concurrency Theory - 23rd International Conference, CONCUR 2012", Newcastle upon Tyne, United Kingdom, September 2012, vol. 7454, p. 317-332 [*DOI :* 10.1007/978-3-642-32940-1], http://hal.inria.fr/hal-00761116.

[27] M. MIO, A. SIMPSON. *A Proof System for Compositional Verification of Probabilistic Concurrent Processes*, in "FoSSaCS", Rome, Italy, F. PFENNING (editor), March 2013, 15, http://hal.inria.fr/hal-00766384.

[28] C. PALAMIDESSI, M. STRONATI. *Differential privacy for relational algebra: improving the sensitivity bounds via constraint systems*, in "QAPL - Tenth Workshop on Quantitative Aspects of Programming Languages", Tallin, Estonia, H. WIKLICKY, M. MASSINK (editors), Open Publishing Association, 2012, vol. 85, p. 92-105 [*DOI :* 10.4204/EPTCS.85.7], http://hal.inria.fr/hal-00760688.

[29] L. XU. *Modular Reasoning about Differential Privacy in a Probabilistic Process Calculus*, in "7th International Symposium on Trustworthy Global Computing (TGC)", Newcastle upon Tyne, United Kingdom, 2013, page : to appear, http://hal.inria.fr/hal-00691284.

[30] M. YANG, V. SASSONE, S. HAMADOU. *A Game-Theoretic Analysis of Cooperation in Anonymity Networks*, in "Principles of Security and Trust - First International Conference, POST 2012", Tallinn, Estonia, P. DEGANO, J. D. GUTTMAN (editors), Springer, 2012, vol. 7215, p. 269-289, http://hal.inria.fr/hal-00760445.

Conferences without Proceedings

[31] A. BARCO, S. KNIGHT, F. D. VALENCIA. *K-Stores: A Spatial and Epistemic Concurrent Constraint Interpreter*, in "21st Workshop on Functional and (Constraint) Logic Programming (WFLP2012)", Nagoya, Japan, May 2012, http://hal.inria.fr/hal-00761679.

Scientific Books (or Scientific Book chapters)

[32] S. KNIGHT, R. MARDARE, P. PANANGADEN. *Combining Epistemic Logic and Hennessy-Milner Logic*, in "Logic and Program Semantics - Essays Dedicated to Dexter Kozen on the Occasion of His 60th Birthday", April 2012, p. 219-243 [*DOI :* 10.1007/978-3-642-29485-3_14], http://hal.inria.fr/hal-00760967.

Books or Proceedings Editing

[33] S. ABRAMSKY, M. W. MISLOVE, C. PALAMIDESSI (editors). *Special issue dedicated to a selection of papers from Mathematical Foundations of Programming Semantics (MFPS XXV). Theoretical Computer Science 430 (1-2)*, Elsevier B.V., 2012, 125, http://hal.inria.fr/hal-00778536.

[34] S. MÖDERSHEIM, C. PALAMIDESSI (editors). *Proceedings of the Joint Workshop on Theory of Security and Applications (TOSCA)*, Lecture Notes in Computer Science, Springer, 2012, vol. 6993, 224 [*DOI :* 10.1007/978-3-642-27375-9], http://hal.inria.fr/hal-00778535.

[35] M. D. RYAN, C. PALAMIDESSI (editors). *Proceedings of the 7th International Symposium on Trustworthy Global Computing (TGC)*, Lecture Notes in Computer Science, Springer, 2013, To appear, http://hal.inria.fr/hal-00778538.

Research Reports

[36] M. ANDRÉS, N. E. BORDENABE, K. CHATZIKOKOLAKIS, C. PALAMIDESSI. *Geo-Indistinguishability: Differential Privacy for Location-Based Systems*, 2012, http://hal.inria.fr/hal-00766821.

[37] K. CHATZIKOKOLAKIS, M. ANDRÉS, N. E. BORDENABE, C. PALAMIDESSI. *Enhancing Differential Privacy: from Hamming to General Metrics*, 2012, http://hal.inria.fr/hal-00767210.

Other Publications

[38] E. ELSALAMOUNY, K. CHATZIKOKOLAKIS, C. PALAMIDESSI. *A differentially private mechanism of optimal utility for a region of priors*, This paper is to appear in the proceedings of POST 2013 (Principles of Security and Trust)., http://hal.inria.fr/hal-00760735.

References in notes

[39] 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/.

[40] 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/.

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

[42] 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/.

[43] 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.

[44] 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.