IN PARTNERSHIP WITH:
CNRS

Ecole Polytechnique

Activity Report 2013

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
Security and Confidentiality

Table of contents

1. Members **1**
2. Overall Objectives **2**
3. Research Program **2**
3.1. Probability and information theory 2
3.2. Expressiveness of Concurrent Formalisms 3
3.3. Concurrent constraint programming 3
3.4. Model checking 3
4. Application Domains **3**
5. Software and Platforms **4**
5.1. Location Guard 4
5.2. PRISM model generator 4
5.3. Calculating the set of corner points of a channel 5
5.4. MMCsp, a compiler for the π-calculus 5
6. New Results **5**
6.1. Foundations of information hiding 5
6.1.1. Differential privacy with general metrics. 5
6.1.2. Privacy for location-based services. 6
6.1.3. Relation between differential privacy and quantitative information flow. 6
6.1.4. A differentially private mechanism of optimal utility for a region of priors 6
6.1.5. Compositional analysis of information hiding 6
6.1.6. Preserving differential privacy under finite-precision semantics 7
6.1.7. Metrics for differential privacy in concurrent systems 7
6.1.8. Unlinkability 7
6.1.9. Trust in anonymity networks 8
6.2. Foundations of Concurrency 8
6.2.1. Models and Emerging Trends of Concurrent Constraint Programming 8
6.2.2. Efficient computation of program equivalence for confluent concurrent constraint programming 8
6.2.3. Abstract Interpretation of Temporal Concurrent Constraint Programs 9
6.2.4. Foundations of Probabilistic Concurrent Systems 9
7. Partnerships and Cooperations **9**
7.1. National Initiatives 9
7.1.1. ANR projects 9
7.1.1.1. ANR-09-BLAN-0169-01 9
7.1.1.2. ANR-09-BLAN-0345-02 9
7.1.2. Large-scale initiatives 10
7.2. European Initiatives 10
7.3. International Initiatives 11
7.3.1. Inria Associate Teams 11
7.3.2. Inria International Partners 11
7.3.3. Participation In other International Programs 11
7.3.3.1. PACE 11
7.3.3.2. LOCALI 11
7.4. International Research Visitors 12
7.4.1. Visits of International Scientists 12
7.4.2. Internships 12
7.4.2.1. Xiao Wang 12
7.4.2.2. Fernán Martinelli 12
7.4.3. Visits to International Teams 12

8. Dissemination **12**
8.1. Scientific Animation 12
8.1.1. Editorial activity 13
8.1.2. Steering Committees 13
8.1.3. Invited Talks 13
8.1.4. Organization of workshops and conferences 13
8.1.5. Participation in program committees 13
8.1.6. Participation in other committees 14
8.1.7. Organization of seminars 15
8.1.8. Service 15
8.2. Teaching - Supervision - Juries 15
8.2.1. Teaching 15
8.2.2. Supervision 15
8.2.3. Other didactical duties 15
9. Bibliography **16**

Project-Team COMETE

Keywords: Concurrency, Constraints, Information Theory, Quantitative Information Flow, Privacy

Creation of the Project-Team: 2008 January 01.

1. Members

Research Scientists

Catuscia Palamidessi [Team leader, Inria, Senior Researcher]
Konstantinos Chatzikokolakis [CNRS, Researcher]
Frank Valencia [CNRS, Researcher]

PhD Students

Nicolas Bordenabe [Inria, grant by Inria-DGA]
Yamil Salim Perchy [Inria, grant by Digicosme, from Nov 2013]
Luis Pino [Inria, grant by Inria-DGA]
Marco Stronati [Ecole Polytechnique, grant Monge]
Lili Xu [Inria, grant by the ANR LOCALI project]
Sophia Knight [Inria, grant by Inria-CORDIS , until Sep 2013]

Post-Doctoral Fellows

Thomas Given-Wilson [Inria, grant by the ANR PACE project, from Sep 2013]
Sardaouna Hamadou [Inria, grant by the ANR LOCALI project, until Nov 2013]
Tobias Heindel [Inria, grant by the CEA, from Jun 2013 until Nov 2013]
Yusuke Kawamoto [Inria, grant by the Digital Society Institute project, from Jul 2013]
Matteo Mio [Grant by ERCIM, until Feb 2013]

Visiting Scientists

Nikita Borisov [Associate Professor, University of Illinois at Urbana-Champaign, from Nov 2013 until Dec 2013]
Moreno Falaschi [Professor, University of Siena, from Sep 2013 until Sep 2013]
Mario Ferreira Alvim Junior [Assistant Professor, Federal University of Minas Gerais, from Nov 2013 until Dec 2013]
Fabio Gadducci [Associate Professor, University of Pisa, from Jun 2013 until Aug 2013]
Dominik Luecke [Postdoc, from Apr 2013 until Apr 2013]
Annabelle Mciver [Associate Professor, Macquarie University, from Dec 2013 until Dec 2013]
Charles Carroll Morgan [Professor, University of New South Wales, from Dec 2013 until Dec 2013]
Carlos Olarte [Associate Professor, Universidad Javeriana Cali, from June 2013 until Jul 2013]
Camilo Rueda [Professor, Universidad Javeriana Cali, from Nov 2013 until Dec 2013]
Vladimiro Sassone [Professor, University of Southampton, from Apr 2013 until May 2013]
Mauricio Toro Bermudez [Postdoc, University of Cyprus, from Jun 2013 until Jun 2013]

Administrative Assistant

Christelle Liévin [Inria]

Others

Fernan Gabriel Martinelli [Master student, from Sep 2012 until Apr 2013]
Xiao Wang [Master student, Ecole Polytechnique, from May 2013 until Aug 2013]

2. Overall Objectives

2.1. Overall Objectives

Our times are characterized by the massive presence of highly *distributed systems* consisting of diverse and specialized devices, forming heterogeneous networks, and providing different services and applications. Revolutionary phenomena such as *social networks* and *cloud computing* are examples of such systems.

In Comète we study emerging concepts of this new era of computing. *Security* and *privacy* are some of the fundamental concerns that arise in this setting. In particular, in the modern digital world the problem of keeping information secret or confidential 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 malicious agents the opportunity to gather and store huge amount of information, often without the individual even being aware of it. Mobility is an additional source of vulnerability, since tracing may reveal significant information. To avoid these kinds of hazards, *security protocols* and various techniques for privacy protection have been designed. However, the properties that they are supposed to ensure are rather subtle, and, furthermore, it is difficult to foresee all possible expedients that a potential attacker may use. As a consequence, even protocols that seem at first "obviously correct" are later (often years later) found to be prone to attacks.

In addition to the security problems, the problems of correctness, robustness and reliability are made more challenging by the complexity of these systems, since they are highly concurrent and distributed. Despite being based on impressive engineering technologies, they are still prone to faulty behavior due to errors in the software design.

To overcome these drawbacks, we need to develop formalisms, reasoning techniques, and verification methods, to specify systems and protocols, their intended properties, and to guarantee that these intended properties of correctness and security are indeed satisfied.

In Comète we study formal computational frameworks for specifying these systems, theories for defining the desired properties of correctness and security and for reasoning about them, and methods and techniques for proving that a given system satisfies the intended properties.

3. Research Program

3.1. Probability and information theory

Participants: Nicolas Bordenabe, Konstantinos Chatzikokolakis, Thomas Given-Wilson, Sardaouna Hamadou, Yusuke Kawamoto, 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 is able to exploit such information.

The recent tendency is to use an information theoretic approach to model the problem and define the leakage in a quantitative way. The idea is to consider the 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 as a 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). Most of the proposals in the literature use Shannon entropy, which is the most established notion of entropy in information theory. We, however, consider also other notions, in particular Rényi min-entropy, which seems to be more appropriate for security in common scenarios like one-try attacks.

3.2. Expressiveness of Concurrent Formalisms

Participants: Catuscia Palamidessi, Luis Pino, Frank Valencia.

We study computational models and languages for distributed, 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, also taking into account the issue of (efficient) implementability.

3.3. Concurrent constraint programming

Participants: Sophia Knight, Luis Pino, Frank Valencia.

Concurrent constraint programming (ccp) is a well established process calculus for modeling systems where agents interact by posting and asking information in a store, much like in users interact in *social networks*. 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 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. **(a)** 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. **(b)** The extension of ccp with constructs to capture emergent systems such as those in social networks and cloud computing.

3.4. Model checking

Participants: Konstantinos Chatzikokolakis, Catuscia Palamidessi.

Model checking addresses the problem of establishing whether a given specification satisfies a certain property. We are interested in developing model-checking techniques for verifying concurrent systems of the kind explained above. In particular, we focus on security and privacy, i.e., on the problem of proving that a given system satisfies the intended security or privacy properties. Since the properties we are interested in have a probabilistic nature, we use probabilistic automata to model the protocols. A challenging problem is represented by the fact that the interplay between nondeterminism and probability, which in security presents subtleties that cannot be handled with the traditional notion of a scheduler,

4. Application Domains

4.1. Security and privacy

Participants: Nicolas Bordenabe, Konstantinos Chatzikokolakis, Catuscia Palamidessi, Marco Stronati.

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 and Platforms

5.1. Location Guard

Participants: Konstantinos Chatzikokolakis [correspondant], Marco Stronati.

The purpose of *Location Guard* is to implement obfuscation techniques for achieving location privacy, in a an easy and intuitive way that makes them available to the general public. Various modern applications, running either on smartphones or on the web, allow third parties to obtain the user's location. A smartphone application can obtain this information from the operating system using a system call, while web application obtain it from the browser using a javascript call.

Although both mobile operating systems and browsers require the user's permission to disclose location information, the user faces an "all-or-nothing" choice: either disclose his exact location and give up his privacy, or stop using the application. This forces many users to disclose their location, although ideally they would like to enjoy some privacy.

The API level of a browser or an operating system would be an ideal place for integrating a location obfuscation technique, in a way that is easy to understand for the average user, and readily available to all applications. When an application asks for the user's location, the browser or operating system can ask the user's permission, but including the option to provide an obfuscated location instead of the real one! Different levels of obfuscation can be also offered, so that the user can chose to provide more accurate location to applications that really need it, and more noisy location to those that don't.

A prototype of Location Guard has been already implemented for Google Chrome. In the future we plan to extend it to other desktop and mobile browsers (Firefox, Internet Explorer, etc), as well as to implement it in modern mobile operating systems, primarily on Android.

https://github.com/chatziko/location-guard

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 [29] 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 [30] 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

Participant: Catuscia Palamidessi [correspondant].

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 [32] and [31].

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. Differential privacy with general metrics.

Differential privacy can be interpreted as a bound on the distinguishability of two generic databases, which is determined by their Hamming distance: the distance in the graph determined by the adjacency relation (two databases are adjacent if they differ for one individual).

In [21] 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.

6.1.2. Privacy for location-based services.

The growing popularity of location-based services, 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 [19] we studied geo-indistinguishability, a formal notion of privacy for location-based services 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 [21] 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.

6.1.3. Relation between differential privacy and quantitative information flow.

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

In [13], we analyzed critically the notion of differential privacy in light of the conceptual framework provided by the Rényi min information theory. We proved that there is a close relation between differential privacy and leakage, due to the graph symmetries induced by the adjacency relation. Furthermore, we considered the utility of the randomized answer, which measures its expected degree of accuracy. We focused on certain kinds of utility functions called "binary", which have a close correspondence with the Rényi min mutual information. Again, it turns out that there can be a tight correspondence between differential privacy and utility, depending on the symmetries induced by the adjacency relation and by the query. Depending on these symmetries we can also build an optimal-utility randomization mechanism while preserving the required level of differential privacy. Our main contribution was a study of the kind of structures that can be induced by the adjacency relation and the query, and how to use them to derive bounds on the leakage and achieve the optimal utility.

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

Differential privacy (already introduced in the previous sections) 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 [22] 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.

6.1.5. 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 [14] 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 [26], 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.

6.1.6. Preserving differential privacy under finite-precision semantics

The approximation introduced by finite-precision representation of continuous data can induce arbitrarily large information leaks even when the computation using exact semantics is secure. Such leakage can thus undermine design efforts aimed at protecting sensitive information. For instance, the standard approach to achieve differential privacy (introduced in previous sections) is the addition of noise to the true (private) value. To date, this approach has been proved correct only in the ideal case in which computations are made using an idealized, infinite-precision semantics. In [23], we analyzed the situation at the implementation level, where the semantics is necessarily finite-precision, i.e. the representation of real numbers and the operations on them are rounded according to some level of precision. We showed that in general there are violations of the differential privacy property, and we studied the conditions under which we can still guarantee a limited (but, arguably, totally acceptable) variant of the property, under only a minor degradation of the privacy level. Finally, we illustrated our results on two cases of noise-generating distributions: the standard Laplacian mechanism commonly used in differential privacy, and a bivariate version of the Laplacian recently introduced in the setting of privacy-aware geolocation.

6.1.7. Metrics for differential privacy in concurrent systems

Many protocols for protecting confidential information have involved randomized mechanisms and a nondeterministic behavior (such as the Dining Cryptographers protocol or the Crowds protocol). In [28], we investigate techniques for proving differential privacy in the context of concurrent systems which contain both probabilistic and nondeterministic behaviors. Our motivation stems from the work of Tschantz et al., who proposed a verification method based on proving the existence of a stratified family of bijections between states, that can track the privacy leakage, ensuring that it does not exceed a given leakage budget. We improve this technique by investigating state properties which are more permissive and still imply differential privacy. We consider three pseudometrics on probabilistic automata: The first one is essentially a reformulation of the notion proposed by Tschantz et al. The second one is a more liberal variant, still based on the existence of a family of bijections, but relaxing the relation between them by integrating the notion of amortization, which results into a more parsimonious use of the privacy budget. The third one aims at relaxing the bijection requirement, and is inspired by the Kantorovich-based bisimulation metric proposed by Desharnais et al. We cannot adopt the latter notion directly because it does not imply differential privacy. Thus we propose a multiplicative variant of it, and prove that it is still an extension of weak bisimulation. We show that for all the pseudometrics the level of differential privacy is continuous on the distance between the starting states, which makes them suitable for verification. Moreover we formally compare these three pseudometrics, proving that the latter two metrics are indeed more permissive than the first one, but incomparable with each other, thus constituting two alternative techniques for the verification of differential privacy.

6.1.8. 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 [20] 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.

6.1.9. 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.

6.2.1. Models and Emerging Trends of Concurrent Constraint Programming

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, represented as constraints over the variables of the system. Agents communicate by posting and querying partial information in the medium. This covers a vast variety of systems as those arising in biological phenomena, reactive systems, net- centric computing and the advent of social networks and cloud computing. In [17] we surveyed the main applications, developments and current trends of ccp.

6.2.2. Efficient computation of program equivalence for confluent concurrent constraint programming

The development of algorithms and automatic verification procedures for ccp have hitherto been far too little considered. To the best of our knowledge there is only one existing verification algorithm for the standard notion of ccp program (observational) equivalence. In [25] we first showed that this verification algorithm has an exponential-time complexity even for programs from a representative sub-language of ccp; the summation-free fragment (ccp+). We then significantly improved on the complexity of this algorithm by providing two alternative polynomial-time decision procedures for ccp+ program equivalence. Each of these two procedures has an advantage over the other. One has a better time complexity. The other can be easily adapted for the full language of ccp to produce significant state space reductions. The relevance of both procedures derives from the importance of ccp+. This fragment, which has been the subject of many theoretical studies, has strong ties to first-order logic and an elegant denotational semantics, and it can be used to model real-world situations. Its most distinctive feature is that of confluence, a property we exploit to obtain our polynomial procedures.

6.2.3. Abstract Interpretation of Temporal Concurrent Constraint Programs

Timed concurrent constraint programming (tcc) is a declarative model for concurrency offering a logic for specifying reactive systems, i.e. systems that continuously interact with the environment. The universal tcc formalism (utcc) is an extension of tcc with the ability to express mobility. Here mobility is understood as communication of private names as typically done for mobile systems and security protocols. In [15] we considered the denotational semantics for tcc, and we extended it to a "collecting" semantics for utcc based on closure operators over sequences of constraints. Relying on this semantics, we formalized a general framework for data flow analyses of tcc and utcc programs by abstract interpretation techniques. The concrete and abstract semantics we proposed are compositional, thus allowing us to reduce the complexity of data flow analyses. We showed that our method is sound and parametric with respect to the abstract domain. Thus, different analyses can be performed by instantiating the framework. We illustrated how it is possible to reuse abstract domains previously defined for logic programming to perform, for instance, a groundness analysis for tcc programs. We showed the applicability of this analysis in the context of reactive systems. Furthermore, we made also use of the abstract semantics to exhibit a secrecy flaw in a security protocol. We also showed how it is possible to make an analysis which may show that tcc programs are suspension free. This can be useful for several purposes, such as for optimizing compilation or for debugging.

6.2.4. Foundations of Probabilistic Concurrent Systems

In [24] we introduced a formal proof system for compositional verification of probabilistic concurrent processes. Properties are expressed using a probabilistic modal μ-calculus, and the proof system is formulated as a sequent calculus in which sequents are given a quantitative interpretation. A key feature is that the probabilistic scenario is handled by introducing the notion of Markov proof, by which each proof in the system is interpreted as a Markov Decision Process, with the proof only considered valid in the case that the value of the MDP is zero.

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
- **URL:** https://cappris.inria.fr/
- **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

7.2.1.1. MEALS

Program: FP7-PEOPLE-2011-IRSES

Project acronym: MEALS

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

Duration: October 2011 - September 2015

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

Coordinator: Holger Hermans, Saarland University, Germany

Coordinator for the Inria sites: Catuscia Palamidessi, Inria Saclay

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. Inria Associate Teams

7.3.1.1. PRINCESS

Title: Protecting privacy while preserving data access

Inria principal investigator: Catuscia Palamidessi

International Partners:

Geoffrey Smith, Florida International University (United States)

Andre Scedrov, University of Pennsylvania (United States)

Duration: 2013 - 2016

URL: http://www.lix.polytechnique.fr/comete/Projects/Princess/

Abstract: PRINCESS is an Inria associated team focusing on the protection of privacy and confidential information. In particular, we study the issues related to the leakage of confidential information through public observables.

We aim at developing a meaningful notion of measure in order to quantify the leakage of information, and to design mechanisms to limit the amount of leakage, without interfering too severely with the utility of the information that is meant to be disclosed.

The main topics currently investigated are quantitative information flow, where we are developing a decision-theoretic approach, and differential privacy, where we are developing an extension which lifts the basic notion of privacy meant for databases to arbitrary domains.

7.3.2. Inria International Partners

7.3.2.1. Informal International Partners

- **Charles Carroll Morgan,** Professor, University of New South Wales
- **Moreno Falaschi,** Professor, University of Siena
- **Mario Ferreira Alvim Junior,** Assistant Professor, Federal University of Minas Gerais
- **Annabelle Mciver,** Associate Professor, Macquarie University
- **Carlos Olarte,** Associate Professor, Universidad Javeriana Cali

7.3.3. Participation In other International Programs

7.3.3.1. PACE

- **Program:** ANR Blanc International
- **Project title:** Beyond plain Processes: Analysis techniques, Coinduction and Expressiveness
- **Duration:** January 2013 - December 2016
- **URL:** http://perso.ens-lyon.fr/daniel.hirschkoff/pace/
- **Coordinator:** Daniel Hirschkoff , Ecole Normale Supérieure de Lyon
- **Other PI's and partner institutions:** Catuscia Palamidessi, Inria Saclay. Davide Sangiorgi, University of Bologna (Italy). Yuxi Fu, Shanghai Jiao Tong University (China).
- **Abstract:** This project objective is to enrich and adapt these methods, techniques, and tools to much broader forms of interactive models, well beyond the realm of "traditional" processes.

7.3.3.2. LOCALI

- **Program:** ANR Blanc International
- **Project title:** Logical Approach to Novel Computational Paradigms
- **Duration:** October 2011 - September 2015
- **URL:** http://lcs.ios.ac.cn/~locali2013/

- **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

- **Nikita Borisov,** Associate Professor, University of Illinois at Urbana-Champaign, from Nov 2013 until Dec 2013
- **Moreno Falaschi,** Professor, University of Siena, from Sep 2013 until Sep 2013
- **Mario Ferreira Alvim Junior,** Assistant Professor, Federal University of Minas Gerais, from Nov 2013 until Dec 2013
- **Fabio Gadducci,** Associate Professor, University of Pisa, from Jun 2013 until Aug 2013
- **Dominik Luecke,** Postdoc, from Apr 2013 until Apr 2013
- **Annabelle Mciver,** Associate Professor, Macquarie University, from Dec 2013 until Dec 2013
- **Charles Carroll Morgan,** Professor, University of New South Wales, from Dec 2013 until Dec 2013
- **Carlos Olarte,** Associate Professor, Universidad Javeriana Cali, from June 2013 until Jul 2013
- **Camilo Rueda,** Professor, Universidad Javeriana Cali, from Nov 2013 until Dec 2013
- **Vladimiro Sassone,** Professor, University of Southampton, from Apr 2013 until May 2013
- **Mauricio Toro Bermudez,** Postdoc, University of Cyprus, from Jun 2013 until Jun 2013

7.4.2. Internships

7.4.2.1. Xiao Wang

- **Duration:** From May 2013 until August 2013
- **Subject:** Differential privacy and applications of privacy protection in location-based services
- **Institution:** LIX, Ecole Polytechnique

7.4.2.2. 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

7.4.3. Visits to International Teams

Catuscia Palamidessi visited the team of Andre Scedrov and Benjamin Pierce at the University of Pennsylvania. July 2013.

8. Dissemination

8.1. Scientific Animation

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

8.1.1. Editorial activity

Catuscia Palamidessi is/has been:

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 Franck van Breughel, Elham Kashefi and Jan Rutten) of a festschrift dedicated to Prakash Panagaden. Special issue of Lecture Notes in Computer Science.

Co-editor (with Geoffrey Smith) of the special issue of Mathematical Structures in Computer Science dedicated to Quantitative Information Flow.

Co-editor (with Mark Ryan) of the proceedings of TGC 2012, Trustworthy Global Computing. [27]

Frank D. Valencia has been:

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 have been:

Co-editors (with Sebastian Mödersheim and Jun Pang) 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:

Workshop on Logic, Language, Information and Computation. TU Darmstadt, Germany. August 2013.

Forum des jeunes mathématiciennes. ENS Lyon. Novembre 2013.

8.1.4. Organization of workshops and conferences

Catuscia Palamidessi is serving as PC co-chair (together with Erika Ábrahám) of FORTE 2014: the 34th IFIP International Conference on Formal Techniques for Distributed Objects, Components and Systems. Berlin, Germany, 3-6 June 2014. Co-located with DisCoTec 2014.

8.1.5. Participation in program committees

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

QEST 2014. The 11th International Conference on Quantitative Evaluation of Systems. Florence, Italy, 8-12 September 2014.

POST 2014. The 3rd Conference on Principles of Security and Trust. Grenoble, 5-13 April 2014.

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.

Frank D. Valencia is/has been 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.

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

ICFEM 2014: The 6th International Conference on Formal Engineering Methods.

PETS 2014: The 14th Privacy Enhancing Technologies Symposium.

HotPETs 2014: 7th Workshop on Hot Topics in Privacy Enhancing Technologies.

QAPL 2014: 12th Workshop on Quantitative Aspects of Programming Languages.

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

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

HotPETs 2013: 6th Workshop on Hot Topics in Privacy Enhancing Technologies.

8.1.6. Participation in other committees

Catuscia Palamidessi has been serving in the following committees:

Member of the Swedish Research Council Committee for Computer Science, 2013. The main duty of this committee is to evaluate and select the grant applications.

Member of the committee for the ACM SIGSAC 2014 Doctoral Dissertation Award for Outstanding PhD Thesis in Computer and Information Security.

Member of the committee for the Ackermann Award 2013: The EACSL outstanding dissertation award for logic in Computer Science.

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 Nicolás Bordenabe 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.

8.1.8. 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.

8.2. Teaching - Supervision - Juries

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

Master. Frank D. Valencia has been teaching an advanced course on Process Modeling at the Master Program in Computer Science of the Pontificia Universidad Javeriana de Cali, Colombia. Total 30 hours. A.Y. 2012-13.

8.2.2. Supervision

PhD (2010-2013) Sophia Knight. Ecole Polytechnique. Grant Inria/CORDIS. Title of the thesis: *The Epistemic Dimension of Concurrency Theory*. Defended on 20 Sep 2013. Co-supervised by Catuscia Palamidessi and Frank D. Valencia.

PhD (2009-2013) Ivan Gazeau. Ecole Polytechnique. Grant ANR. Title of the thesis: *Safe Programming in finite precision: Controling the errors and information leaks*. Defended on 14 Oct 2013. Co-supervised by Catuscia Palamidessi and Dale Miller.

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

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

PhD in progress (2011-) Nicolás 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 (2013-) Salim Percy. Ecole Polytechnique. Grant Digiteo-Digicosme. Co-supervised by Frank D. Valencia and Stefan Haar.

8.2.3. Other didactical duties

Catuscia Palamidessi is:

- Co-responsible of the Master 2 course on Concurrency since 2003, first at the DEA in Theoretical Computer Science (Paris) and then at the MPRI.
- External member of the scientific council for the PhD in Computer Science at the University of Pisa, Italy. Since 2012.
- Member of the advising committee for the PhD of Andrea Margheri, University of Florence, Italy.

9. Bibliography

Major publications by the team in recent years

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

[2] 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, pp. 265-279 [*DOI :* 10.1109/CSF.2012.26], http://hal.inria.fr/hal-00734044

[3] M. ANDRÉS, N. BORDENABE, K. CHATZIKOKOLAKIS, C. PALAMIDESSI. *Geo-Indistinguishability: Differential Privacy for Location-Based Systems*, in "20th ACM Conference on Computer and Communications Security", Berlin, Allemagne, ACM Press, 2013, pp. 901-914, DGA, Inria large scale initiative CAPPRIS [*DOI :* 10.1145/2508859.2516735], http://hal.inria.fr/hal-00766821

[4] K. CHATZIKOKOLAKIS, M. ANDRÉS, N. BORDENABE, C. PALAMIDESSI. *Broadening the Scope of Differential Privacy Using Metrics*, in "The 13th Privacy Enhancing Technologies Symposium", Bloomington, Indiana, États-Unis, E. DE CRISTOFARO, M. WRIGHT (editors), Springer, 2013, vol. 7981, pp. 82-102, DGA, Inria large scale initiative CAPPRIS [*DOI :* 10.1007/978-3-642-39077-7], http://hal.inria.fr/hal-00767210

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

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

[7] 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, pp. 531–571 [*DOI :* 10.3233/JCS-2008-0333], http://hal.inria.fr/inria-00349224/en/

[8] 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, pp. 226-240, http://www.lix.polytechnique.fr/~fvalenci/papers/fossacs04.pdf

[9] 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, pp. 317-332 [*DOI :* 10.1007/978-3-642-32940-1], http://hal.inria.fr/hal-00761116

[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, pp. 59–68, http://hal.inria.fr/inria-00201096/en/

Publications of the year

Doctoral Dissertations and Habilitation Theses

[11] I. GAZEAU. , *Programmation sûre en précision finie : Contrôler les erreurs et les fuites d'informations*, Ecole Polytechnique X, October 2013, http://hal.inria.fr/pastel-00913469

[12] S. KNIGHT. , *Le point de vue epistémique de théorie de la concurrence*, Ecole Polytechnique X, September 2013, http://hal.inria.fr/tel-00940413

Articles in International Peer-Reviewed Journals

[13] M. ALVIM, M. ANDRÉS, K. CHATZIKOKOLAKIS, P. DEGANO, C. PALAMIDESSI. *On the information leakage of differentially-private mechanisms*, in "Journal of Computer Security", 2014, to appear, http://hal.inria.fr/hal-00940425

[14] K. CHATZIKOKOLAKIS, C. PALAMIDESSI, C. BRAUN. *Compositional Methods for Information-Hiding*, in "Mathematical Structures in Computer Science", 2014, to appear, http://hal.inria.fr/hal-00760596

[15] M. FALASCHI, C. OLARTE, C. PALAMIDESSI. *Abstract Interpretation of Temporal Concurrent Constraint Programs*, in "Theory and Practice of Logic Programming", 2014, http://hal.inria.fr/hal-00945462

[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, to appear, http://hal.inria.fr/hal-00760437

[17] C. OLARTE, C. RUEDA, F. D. VALENCIA. *Models and Emerging Trends of Concurrent Constraint Programming*, in "Constraints", October 2013, vol. 18, n^{o} 4, pp. 535-578 [*DOI :* 10.1007/s10601-013-9145-3], http://hal.inria.fr/hal-00869192

Invited Conferences

[18] C. PALAMIDESSI. *Quantitative Approaches to Information Protection*, in "WOLLIC", Darmstadt, Germany, L. LIBKIN (editor), 2013, http://hal.inria.fr/hal-00945678

International Conferences with Proceedings

[19] M. ANDRÉS, N. BORDENABE, K. CHATZIKOKOLAKIS, C. PALAMIDESSI. *Geo-Indistinguishability: Differential Privacy for Location-Based Systems*, in "20th ACM Conference on Computer and Communications Security", Berlin, Germany, ACM Press, 2013, pp. 901-914 [*DOI :* 10.1145/2508859.2516735], http://hal.inria.fr/hal-00766821

[20] 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, M. R. CATUSCIA PALAMIDESSI (editor), Springer, 2013, vol. 8191, pp. 129-144 [*DOI :* 10.1007/978-3-642-41157-1], http://hal.inria.fr/hal-00760150

[21] K. CHATZIKOKOLAKIS, M. ANDRÉS, N. BORDENABE, C. PALAMIDESSI. *Broadening the Scope of Differential Privacy Using Metrics*, in "The 13th Privacy Enhancing Technologies Symposium", Bloomington, Indiana, United States, E. DE CRISTOFARO, M. WRIGHT (editors), Springer, 2013, vol. 7981, pp. 82-102 [*DOI :* 10.1007/978-3-642-39077-7], http://hal.inria.fr/hal-00767210

[22] E. ELSALAMOUNY, K. CHATZIKOKOLAKIS, C. PALAMIDESSI. *A differentially private mechanism of optimal utility for a region of priors*, in "POST-2nd Conference on Principles of Security and Trust", Rome, Italy, Springer Berlin Heidelberg, 2013, vol. 7796, pp. 41-62 [*DOI :* 10.1007/978-3-642-36830-1_3], http://hal.inria.fr/hal-00760735

[23] I. GAZEAU, D. MILLER, C. PALAMIDESSI. *Preserving differential privacy under finite-precision semantics*, in "QAPL - 11th International Workshop on Quantitative Aspects of Programming Languages and Systems", Rome, Italy, L. BORTOLUSSI, H. WIKLICKY (editors), Electronic Proceedings in Theoretical Computer Science, Open Publishing Association, 2013, vol. 117, pp. 1-18 [*DOI :* 10.4204/EPTCS.117.1], http://hal.inria.fr/hal-00780774

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

[25] L. PINO, F. BONCHI, F. D. VALENCIA. *Efficient computation of program equivalence for confluent concurrent constraint programming*, in "15th International Symposium on Principles and Practice of Declarative Programming, PPDP '13, Madrid, Spain, September 16-18, 2013", France, 2013, pp. 263-274, http://hal.inria.fr/hal-00909394

[26] 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, M. R. CATUSCIA PALAMIDESSI (editor), Springer, 2013, vol. 8191, pp. 198-212, http://hal.inria.fr/hal-00691284

Books or Proceedings Editing

[27] C. PALAMIDESSI, M. D. RYAN (editors). , *Proceedings of the 7th International Symposium on Trustworthy Global Computing (TGC)*, Lecture Notes in Computer Science, Springer, 2013, vol. 8191, pp. 1-212 [*DOI :* 10.1007/978-3-642-41157-1], http://hal.inria.fr/hal-00778538

Research Reports

[28] L. XU, K. CHATZIKOKOLAKIS, H. LIN, C. PALAMIDESSI. , *Metrics for Differential Privacy in Concurrent Systems*, October 2013, http://hal.inria.fr/hal-00879140

References in notes

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

[30] 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, pp. 531–571 [*DOI :* 10.3233/JCS-2008-0333], http://hal.inria.fr/inria-00349224/en/

[31] 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, pp. 209–223, http://hal.archives-ouvertes.fr/inria-00424856/en/

[32] 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, pp. 179-188, http://www.lix.polytechnique.fr/~catuscia/papers/Wu/qest2.pdf