IN PARTNERSHIP WITH:
CNRS

Ecole Polytechnique

Activity Report 2014

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

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, until Sep 2014]
Michell Guzman [Inria, grant Inria CORDI-S, from Nov 2014]
Luis Pino [Inria, grant by Inria-DGA, until Sep 2014]
Yamil Salim Perchy [Inria, grant Digiteo-Digicosme]
Marco Stronati [Ecole Polytechnique, grant Monge]
Lili Xu [Inria, grant ANR DEDUCTEAM-LOCALI project, until Jan 2014]

Post-Doctoral Fellows

Thomas Given-Wilson [Inria, grant Caisse des Dépôts et Consignations]
Yusuke Kawamoto [Inria, grant Fondation de Cooperation Scientifique Campus Paris Saclay-DIGITEO]

Visiting Scientists

Moreno Falashi [Professor at the University of Siena, Italy, from Jul 2014 until Aug 2014]
Maurizio Gabbrielli [Professor, University of Bologna, Italy, from July 2014 until Aug 2014]
Carlos Olarte [Junior Professor at the Universidade Federal do Rio Grande do Norte, Brazil, Jul 2014]

Administrative Assistant

Martine Thirion [Inria]

Other

Raphaelle Crubille [Master student, ENS Lyon, Mar 2014 to July 2014]

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, 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: Michell Guzman, Yamil Salim Perchy, 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. New Software and Platforms

5.1. Location Guard

Participants: Konstantinos Chatzikokolakis [correspondant], Marco Stronati.

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

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 is 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.

Location Guard was created as a prototype for Google Chrome at the end for 2013. In 2014, Location Guard matured from a prototype to a high quality software, supporting both desktop and mobile browsers:

- Google Chrome / Chromium
- Mozilla Firefox and Firefox for Android
- Opera

After only a short period online, the extension has more than 8500 daily users, and it was presented in an article by the popular technology news site Ghacks. Our experience so far shows that end users do care about location privacy, and geo-indistinguishability is a practical approach for providing it.

In the future we plan to make Location Guard more widely available on smartphones, supporting more mobile browsers as well as providing direct integration into the operating system, primarily on Android.

5.2. libqif - A Quantitative Information Flow C++ Toolkit Library

Participants: Konstantinos Chatzikokolakis [correspondant], Marco Stronati.

https://github.com/chatziko/libqif

The goal of libqif is to provide an efficient C++ toolkit implementing a variety of techniques and algorithms from the area of quantitative information flow and differential privacy. We plan to implement all techniques produced by Comète in recent years, as well as several ones produced outside the group, giving the ability to privacy researchers to reproduce our results and compare different techniques in a uniform and efficient framework.

Some of these techniques were previously implemented in an ad-hoc fashion, in small, incompatible with each-other, non-maintained and usually inefficient tools, used only for the purposes of a single paper and then abandoned. We aim at reimplementing those – as well as adding several new ones not previously implemented – in a structured, efficient and maintainable manner, providing a tool of great value for future research. Of particular interest is the ability to easily re-run evaluations, experiments and case-studies from all our papers, which will be of great value for comparing new research results in the future.

The library is still in under heavy development but substantial progress has been made in 2014. Some of the techniques already implemented are:

- Standard leakage measures: Shannon, min-entropy, guessing entropy
- Measures from the g-leakage framework [32]
- Channel factorization
- Standard differential privacy mechanisms from the literature
- The planar Laplace mechanism of [33]
- The standard Kantorovich metric as well as the multiplicative variant from [19]
- All operations are supported for both doubles (for precision) and floats (for memory efficiency)
- All operations involving only rational quantities are supported using arbitrary precision rational arithmetic, allowing to obtain exact results
- Native linear programing for rationals

Many more are scheduled to be added in the near future.

5.3. LeakWatch: Estimating Information Leakage from Java Programs

Participant: Yusuke Kawamoto.

http://www.cs.bham.ac.uk/research/projects/infotools/leakwatch/

Comète contributed to the development of LeakWatch, a quantitative information leakage analysis tool for the Java programming language, created by several people at the University of Birmingham.

LeakWatch is based on a flexible "point-to-point" information leakage model, where secret and publicly-observable data may occur at any time during a program's execution. LeakWatch repeatedly executes a Java program containing both secret and publicly-observable data and uses robust statistical techniques to provide estimates, with confidence intervals, for min-entropy leakage (using a new theoretical result from [23]) and mutual information.

6. New Results

6.1. Highlights of the Year

- Prix de thèse de l'Ecole Polytechnique 2014 for the thesis "The Epistemic View of Concurrency Theory" by Sophia Knight (Defended 20 September, 2013).
- Catuscia Palamidessi has been invited keynote speaker at the joint conferences CONCUR 2014 and TGC 2014. Rome, September 2014.

6.2. 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.

6.2.1. Additive and multiplicative notions of leakage, and their capacities

Protecting sensitive information from improper disclosure is a fundamental security goal. It is complicated, and difficult to achieve, often because of unavoidable or even unpredictable operating conditions that can lead to breaches in planned security defences. An attractive approach is to frame the goal as a quantitative problem, and then to design methods that measure system vulnerabilities in terms of the amount of information they leak. A consequence is that the precise operating conditions, and assumptions about prior knowledge, can play a crucial role in assessing the severity of any measured vunerability.

In [20] we developed this theme by concentrating on vulnerability measures that are *robust* in the sense of allowing general leakage bounds to be placed on a program, bounds that apply whatever its operating conditions and whatever the prior knowledge might be. In particular we proposed a theory of channel capacity, generalising the Shannon capacity of information theory, that can apply both to additive and to multiplicative forms of a recently-proposed measure known as g-leakage. Further, we explored the computational aspects of calculating these (new) capacities: one of these scenarios can be solved efficiently by expressing it as a Kantorovich distance, but another turns out to be NP-complete.

We also found capacity bounds for arbitrary correlations with data not directly accessed by the channel, as in the scenario of Dalenius's Desideratum.

6.2.2. Compositionality Results for Quantitative Information Flow

In the min-entropy approach to quantitative information flow, the leakage is defined in terms of a minimization problem, which, in case of large systems, can be computationally rather heavy. The same happens for the recently proposed generalization called g-vulnerability. In [28] we studied the case in which the channel associated to the system can be decomposed into simpler channels, which typically happens when the observables consist of several components. Our main contribution was the derivation of bounds on the g-leakage of the whole system in terms of the g-leakages of its components.

6.2.3. LeakWatch: Estimating Information Leakage from Java Programs

Programs that process secret data may inadvertently reveal information about those secrets in their publicly-observable output. In [23] we presented LeakWatch, a quantitative information leakage analysis tool for the Java programming language; it is based on a flexible "point-to-point" information leakage model, where secret and publicly-observable data may occur at any time during a program's execution. LeakWatch repeatedly executes a Java program containing both secret and publicly-observable data and uses robust statistical techniques to provide estimates, with confidence intervals, for min-entropy leakage (using a new theoretical result presented in this paper) and mutual information. We demonstrated how LeakWatch can be used to estimate the size of information leaks in a range of real-world Java programs.

6.2.4. On the information leakage of differentially-private mechanisms

Differential privacy aims at protecting the privacy of participants in statistical databases. Roughly, a mechanism satisfies differential privacy if the presence or value of a single individual in the database does not significantly change the likelihood of obtaining a certain answer to any statistical query posed by a data analyst. Differentially-private mechanisms are often oblivious: first the query is processed on the database to produce a true answer, and then this answer is adequately randomized before being reported to the data analyst. Ideally, a mechanism should minimize leakage—i.e., obfuscate as much as possible the link between reported answers and individuals' data—while maximizing utility—i.e., report answers as similar as possible to the true ones. These two goals are, however, conflicting, and a trade-off between privacy and utility is imposed.

In [13] we used quantitative information flow principles to analyze leakage and utility in oblivious differentially-private mechanisms. We introduced a technique that exploits graph-symmetries of the adjacency relation on databases to derive bounds on the min-entropy leakage of the mechanism. We evaluated utility using identity gain functions, which are closely related to min-entropy leakage, and we derived bounds for it. Finally, given some graph-symmetries, we provided a mechanism that maximizes utility while preserving the required level of differential privacy.

6.2.5. Metric-based approaches for privacy in concurrent systems

In a series of two papers we investigated metric-based techniques for varifying differential privacy in the context of concurrent systems.

The first work [30] was motivated from the one 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 improved this technique by investigating state properties which are more permissive and still imply differential privacy. We introduced a new pseudometric, still based on the existence of a family of bijections, but relaxing the relation between them by integrating the notion of amortization, and showed that this results to a more parsimonious use of the privacy budget. We also showed that for the new pseudometric the level of differential privacy is continuous on the distance between the starting states, which makes it suitable for verification.

Continuing this line of work, we studied the pseudometric based on the Kantorovich lifting, which is one of the most popular notions of distance between probabilistic processes proposed in the literature. However, its application in verification is limited to linear properties. In [19], we proposed a generalization which allows to deal with a wider class of properties, such as those used in security and privacy. More precisely, we proposed a family of pseudometrics, parametrized on a notion of distance which depends on the property we want to verify. Furthermore, we showed that the members of this family still characterize bisimilarity in terms of their kernel, and provided a bound on the corresponding distance between trace distributions. Finally, we studied the instance corresponding to differential privacy, and we showed that it has a dual form, easier to compute. We also proved that the typical process-algebra constructs are non-expansive, thus paving the way to a modular approach to verification.

6.2.6. Optimal Geo-Indistinguishable Mechanisms for Location Privacy

With location-based services becoming increasingly more popular, serious concerns are being raised about the potential privacy breaches that the disclosure of location information may induce. In [21] we considered two approaches that have been proposed to limit and control the privacy loss: one is the *geo-indistinguishability* notion developed within Comète, which is inspired by differential privacy, and like the latter it is independent from the side knowledge of the adversary and robust with respect to composition of attacks. The other one is the mechanism of Shokri et al., which offers an optimal trade-off between the loss of quality of service and the privacy protection with respect to a given Bayesian adversary.

We showed that it is possible to combine the advantages of the two approaches: given a minimum threshold for the degree of geo-indistinguishability, we construct a mechanism that offers the maximal utility, as the solution of a linear program. Thanks to the fact that geo-indistinguishability is insensitive to the remapping of a Bayesian adversary, the mechanism so constructed is optimal also in the sense of Shokri et al. Furthermore we proposed a method to reduce the number of constraints of the linear program from cubic to quadratic (with respect to the number of locations), maintaining the privacy guarantees without affecting significantly the utility of the generated mechanism. This lowers considerably the time required to solve the linear program, thus enlarging significantly the size of location sets for which the optimal trade-off mechanisms can still be computed.

6.2.7. A Predictive Differentially-Private Mechanism for Mobility Traces

With the increasing popularity of GPS-enabled handheld devices, location based applications and services have access to accurate and real-time location information, raising serious privacy concerns for their millions of users. Trying to address these issues, the notion of *geo-indistinguishability* was recently introduced, adapting the well-known concept of Differential Privacy to the area of location-based systems. A Laplace-based obfuscation mechanism satisfying this privacy notion works well in the case of a *sporadic* use; Under repeated use, however, *independently* applying noise leads to a quick loss of privacy due to the correlation between the location in the trace.

In [22] we showed that correlations in the trace can be in fact exploited in terms of a *prediction function* that tries to guess the new location based on the previously reported locations. The proposed mechanism tests the quality of the predicted location using a private test; in case of success the prediction is reported otherwise the location is sanitized with new noise. If there is considerable correlation in the input trace, the extra cost of the test is small compared to the savings in budget, leading to a more efficient mechanism.

We evaluated the mechanism in the case of a user accessing a location-based service while moving around in a city. Using a simple prediction function and two budget spending strategies, optimizing either the utility or the budget consumption rate, we showed that the predictive mechanism can offer substantial improvements over the independently applied noise.

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

Differential privacy is a notion of privacy that was initially designed for statistical databases, and has been recently extended to a more general class of domains. Both differential privacy and its generalized version can be achieved by adding random noise to the reported data. Thus, privacy is obtained at the cost of reducing the data's accuracy, and therefore their *utility*.

In [31] we considered the problem of identifying *optimal* mechanisms for generalized differential privacy, i.e. mechanisms that maximize the utility for a given level of privacy. The utility usually depends on a prior distribution of the data, and naturally it would be desirable to design mechanisms that are *universally optimal*, i.e., optimal for all priors. However it is already known that such mechanisms do not exist in general. We then characterized maximal *classes of priors* for which a mechanism which is optimal for all the priors of the class *does exist*. We showed that such classes can be defined as convex polytopes in the priors space.

As an application, we considered the problem of privacy that arises when using, for instance, location-based services, and we showed how to define mechanisms that maximize the quality of service while preserving the desired level of geo-indistinguishability.

6.2.9. 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.

6.3. 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.

6.3.1. A Concurrent Pattern Calculus

In [16] we detailed how Concurrent pattern calculus (CPC) drives interaction between processes by comparing data structures, just as sequential pattern calculus drives computation. By generalising from pattern matching to pattern unification, interaction becomes symmetrical, with information flowing in both directions. CPC provides a natural language to express trade where information exchange is pivotal to interaction. The

unification allows some patterns to be more discriminating than others; hence, the behavioural theory must take this aspect into account, so that bisimulation becomes subject to compatibility of patterns. Many popular process calculi can be encoded in CPC; this allows for a gain in expressiveness, formalised through encodings.

6.3.2. An Intensional Concurrent Faithful Encoding of Turing Machines

The benchmark for computation is typically given as Turing computability; the ability for a computation to be performed by a Turing Machine. Many languages exploit (indirect) encodings of Turing Machines to demonstrate their ability to support arbitrary computation. However, these encodings are usually by simulating the entire Turing Machine within the language, or by encoding a language that does an encoding or simulation itself. This second category is typical for process calculi that show an encoding of lambda-calculus (often with restrictions) that in turn simulates a Turing Machine. Such approaches lead to indirect encodings of Turing Machines that are complex, unclear, and only weakly equivalent after computation. In [25] we developed an approach to encoding Turing Machines into intensional process calculi that is faithful, reduction preserving, and structurally equivalent. The encoding is demonstrated in a simple asymmetric concurrent pattern calculus before generalised to simplify infinite terms, and to show encodings into Concurrent Pattern Calculus and Psi Calculi.

6.3.3. Expressiveness via Intensionality and Concurrency

Computation can be considered by taking into account two dimensions: extensional versus intensional, and sequential versus concurrent. Traditionally sequential extensional computation can be captured by the lambda-calculus. However, recent work shows that there are more expressive intensional calculi such as SF-calculus. Traditionally process calculi capture computation by encoding the lambda-calculus, such as in the pi-calculus. Following this increased expressiveness via intensionality, other recent work has shown that concurrent pattern calculus is more expressive than pi-calculus. In [26] we formalised the relative expressiveness of all four of these calculi by placing them on a square whose edges are irreversible encodings. This square is representative of a more general result: that expressiveness increases with both intensionality and concurrency.

6.3.4. On the Expressiveness of Intensional Communication

The expressiveness of communication primitives has been explored in a common framework based on the pi-calculus by considering four features: synchronism (asynchronous vs synchronous), arity (monadic vs polyadic data), communication medium (shared dataspaces vs channel-based), and pattern-matching (binding to a name vs testing name equality). In [27] pattern-matching is generalised to account for terms with internal structure such as in recent calculi like Spi calculi, Concurrent Pattern Calculus and Psi calculi. This exploreD intensionality upon terms, in particular communication primitives that can match upon both names and structures. By means of possibility/impossibility of encodings, we showed that intensionality alone can encode synchronism, arity, communication-medium, and pattern-matching, yet no combination of these without intensionality can encode any intensional language.

6.3.5. Weak CCP Bisimilarity with Strong Procedures

Concurrent constraint programming (CCP) is a well-established model for concurrency that singles out the fundamental aspects of asynchronous systems whose agents (or processes) evolve by posting and querying (partial) information in a global medium. Bisimilarity is a standard behavioral equivalence in concurrency theory. However, only recently a well-behaved notion of bisimilarity for CCP, and a CCP partition refinement algorithm for deciding the strong version of this equivalence have been proposed. Weak bisimilarity is a central behavioral equivalence in process calculi and it 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 [17] we demonstrated that, because of its involved labeled transitions, the above-mentioned saturation technique does not work for CCP. We 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.

6.3.6. Efficient Algorithms for Program Equivalence for Confluent Concurrent Constraint Programming

While the foundations and principles of CCP e.g., semantics, proof systems, axiomatizations, have been thoroughly studied for over the last two decades. In contrast, 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 [18] 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 exploited to obtain our polynomial procedures.

6.3.7. A Behavioral Congruence for Concurrent Constraint Programming with Nondeterministic Choice

Weak bisimilarity is one of the most representative notions of behavioral equivalence for models of concurrency. As we mentioned earlier, a notion of weak bisimilarity, called weak saturated barbed bisimilarity (wsbb), was recently proposed for CCP. This equivalence improves on previous bisimilarity notions for CCP that were too discriminating and it is a congruence for the choice-free fragment of CCP. In [29], however, we showed that wsbb is not a congruence for CCP with nondeterministic choice. We then introduced a new notion of bisimilarity, called weak full bisimilarity (wfb), and showed that it is a congruence for the full language of CCP. We also showed the adequacy of wfb by establishing that it coincides with the congruence induced by closing wsbb under all contexts. The advantage of the new definition is that, unlike the congruence induced by wsbb, it does not require quantifying over infinitely many contexts.

6.3.8. 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 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.

6.3.9. Bisimulation for Markov Decision Processes through Families of Functional Expressions

In [24], we transfered a notion of quantitative bisimilarity for labelled Markov processes to Markov decision processes with continuous state spaces. This notion takes the form of a pseudometric on the system states, cast in terms of the equivalence of a family of functional expressions evaluated on those states and interpreted as a real-valued modal logic. Our proof amounted to a slight modification of previous techniques used to prove

equivalence with a fixed-point pseudometric on the state-space of a labelled Markov process and making heavy use of the Kantorovich probability metric. Indeed, we again demonstrated equivalence with a fixed-point pseudometric defined on Markov decision processes; what is novel is that we recasted this proof in terms of integral probability metrics defined through the family of functional expressions, shifting emphasis back to properties of such families. The hope is that a judicious choice of family might lead to something more computationally tractable than bisimilarity whilst maintaining its pleasing theoretical guarantees. Moreover, we used a trick from descriptive set theory to extend our results to MDPs with bounded measurable reward functions, dropping a previous continuity constraint on rewards and Markov kernels.

7. Partnerships and Cooperations

7.1. National Initiatives

7.1.1. 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 & H2020 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-MSR joint lab

7.3.1.1. Privacy-Friendly Services and Apps

Title: Privacy-Friendly Services and Applications

Inria principal investigator: Catuscia Palamidessi

International Partners:

Cedric Fournet, Microsoft Research Lab, Cambridge, UK

Andy Gordon, Microsoft Research Lab, Cambridge, UK

Duration: 2014 - 2016

URL: http://www.msr-inria.fr/projects/privacy-friendly-services-and-apps/

Abstract: This is a project sponsored by Microsoft Research Lab, on methods to preserve privacy in web services and location-based services.

7.3.2. Inria Associate Teams

7.3.2.1. 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.

7.3.3. Inria International Partners

7.3.3.1. Informal International Partners

Moreno Falaschi, Professor, University of Siena, Italy

Mario Ferreira Alvim Junior, Assistant Professor, Federal University of Minas Gerais, Brazil

Annabelle Mciver, Associate Professor, Macquarie University, Australia

Charles Carroll Morgan, Professor, University of New South Wales, Australia

Carlos Olarte, Adjunct professor at Universidade Federal do Rio Grande do Norte, Brazil

Camilo Rueda, Professor, Universidad Javeriana Cali, Colombia

7.3.4. Participation In other International Programs

7.3.4.1. 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.

7.3.4.2. 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.

7.3.4.3. MUSICAL

Program: CNPq Science Without Borders.

Project title: Music and Spatial Interaction with Constraints, Algebra and Logic: Foundations and Applications.

Duration: Oct 2014- Oct 2016

URL: http://cic.puj.edu.co/~caolarte/musical/Musical/Welcome.html

Coordinator: Elaine Pimentel, Universidade Federal do Rio Grande do Norte (Brazil),

Other PI's and partner institutions: Camilo Rueda, PUJ Cali (Colombia). Carlos Olarte, Universidade Federal do Rio Grande do Norte (Brazil). Frank Valencia, CNRS-LIX and Inria Saclay (France). Gerard Assayag, IRCAM (France).

Abstract: This multi-disciplinary project aims to develop and integrate tools from logic and concurrency theory for the design and analysis of reactive systems and to their application to musical processes and multimedia systems.

7.4. International Research Visitors

7.4.1. Visits of International Scientists

Mauricio Cano, Masters Student, Universidad Javeriana Cali, Colombia, Nov 2014

Moreno Falaschi, Professor, University of Siena, Italy, from July 2014 until Aug 2014

Mario Ferreira Alvim Junior, Assistant Professor, Federal University of Minas Gerais, Brazil, Dec 2014

Maurizio Gabbrielli, Professor, University of Bologna, Italy, from July 2014 until Aug 2014

Daniel Gebler, PhD student, Free University of Amsterdam, The Netherlands, Jun 2014

Justin Hsu, PhD student, University of Pennsylvania, USA, Nov 2014

Annabelle McIver, Associate Professor, Macquarie University, Australia, Dec 2014

Hernan Claudio Melgratti, Associate Professor, University of Buenos Aires, Argentina, Apr 2014

Carroll Morgan, Professor, University of New South Wales and NICTA, Australia, Dec 2014

Carlos Olarte, Adjunct professor at Universidade Federal do Rio Grande do Norte, Brazil, from June 2014 until Jul 2014

Camilo Rueda, Professor, Universidad Javeriana Cali, Colombia, from Nov 2014 to Nov 2014

Geoffrey Smith, Professor, Florida International University, USA, Dec 2014

7.4.1.1. Internships

7.4.1.1.1. Raphaelle Crubillé

Duration: From Mar 2014 until Jul 2014

Subject: Formal modelling of RFID distance bounding protocols

Institution: ENS Lyon

7.4.2. Visits to International Teams

Konstantinos Chatzikokolakis and Catuscia Palamidessi visited the team of Annabelle McIver and Carroll Morgan at Macquarie University, Australia, July 2014.

Frank Valencia visited the team of Camilo Rueda (AVISPA) at Pontifical Universidad Javeriana Cali, from July 2014 until July 2014

8. Dissemination

8.1. Promoting Scientific Activities

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

8.1.1. Scientific events organisation

8.1.1.1. Member of the organizing committee

Catuscia Palamidessi is member of:

The Executive Committee of SIGLOG, the ACM Special Interest Group on Logic and Computation. Since 2014.

The Organizing Committee of LICS, the ACM/IEEE Symposium on Logic in Computer Science. Since 2010.

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 is member of:

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

8.1.2. Scientific events selection

8.1.2.1. Chair of conference program committee

Catuscia Palamidessi:

has served 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.

is serving as PC chair of LICS 2015. 30th Annual ACM/IEEE Symposium on Logic in Computer Science. Kyoto, Japan, 6-10 July 2015.

Frank Valencia:

is serving as PC chair (with Camilo Rueda and Martin Leucker as co-chairs) of ICTAC 2015: The 12th International Colloquium on Theoretical Aspects of Computing. Cali, Colombia Oct 29-31 2015. Co-located with 11th International Workshop on Developments in Computational Models DCM 2015.

is serving as chair of the organising committee (with Camilo Rueda and Martin Leucker as co-chairs) of ICTAC 2015: The 12th International Colloquium on Theoretical Aspects of Computing. Cali, Colombia Oct 29-31 2015.

8.1.2.2. Member of the conference program committee

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

ATVA 2015. The 13th International Symposium on Automated Technology for Verification and Analysis. Shanghai, China, 12-15 October 2015.

FORTE 2015. The 35th IFIP International Conference on Formal Techniques for Distributed Objects, Components and Systems. Inria Grenoble, France, 2-4 June 2015.

TPDP 2015. Theory and Practice of Differential Privacy London, UK, 18 April 2015.

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.

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

PETS 2015: The 15th Privacy Enhancing Technologies Symposium

FCS 2015: Workshop on Foundations of Computer Security

TPDP 2015: 1st workshop on the Theory and Practice of Differential Privacy

QAPL 2015: 13th Workshop on Quantitative Aspects of Programming Languages

ICISSP 2015: 1st International Conference on Information Systems Security and Privacy

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

APVP 2014: 5ème Atelier sur la Protection de la Vie Privée

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

PPDP 2014. The 16th International Symposium on Principles and Practice of Declarative Programming (PPDP 2014). Canterbury, UK, September 8-10, 2014.

ICLP DC 2014. Tenth ICLP Doctoral Consortium. Vienna, Austria, July 20, 2014?,

PPDP 2015. The 17th International Symposium on Principles and Practice of Declarative Programming (PPDP 2014). Siena, Italy, July 14-16, 2015.

ICLP DC 2015. 11th ICLP Doctoral Consortium. Cork, Ireland, 31 August 2015.

8.1.2.3. Reviewer

The members of the team reviewed several papers for international conferences and workshops.

8.1.3. Journal

8.1.3.1. Member of the editorial board

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.

Member of the Editorial Board of LIPIcs: Leibniz International Proceedings in Informatics, Schloss Dagstuhl –Leibniz Center for Informatics.

Konstantinos Chatzikokolakis is:

Editorial board member of the newly established Proceedings on Privacy Enhancing Technologies (PoPETs), a scholarly journal for timely research papers on privacy.

8.1.3.2. Reviewer

The members of the team reviewed several papers for international journals.

8.1.4. Other Editorial Activities

Catuscia Palamidessi is/has been:

Co-editor of the special issue of Logical Methods in Computer Science dedicated to selected papers of DisCoTec 2014.

Co-editor (with Franck van Breughel, Elham Kashefi and Jan Rutten) of *Horizons of the Mind. A Tribute to Prakash Panangaden*. *LNCS* 8464 Springer, 499 pages, 2014.

Co-editor (with Erika Ábrahám) of the proceedings of FORTE 2014. *LNCS* 8461 Springer, 301 pages, 2014.

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. Published on 12 December 2014.

Co-editor of the special issue of Mathematical Structures in Computer Science dedicated to the 17th International Workshop on Expressiveness in Concurrency. Published on 10 November 2014.

8.1.5. Other Activities

8.1.5.1. Invited talks

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

CONCUR 2014 and TGC 2014 (Joint keynote speaker) The 25th Conference on Concurrency Theory and the 9th International Symposium on Trustworthy Global Computing. Rome, Italy, September 2014.

NWPT 2014. The 26th Nordic Workshop on Programming Theory. Halmstad, Sweden October 2014.

8.1.5.2. Participation in other committees

Catuscia Palamidessi has been serving in the following committees:

Member of the Swedish Research Council Committee for Computer Science, 2014. 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 WWTF jury – selection committee for a tenure-track position at the Faculty of Informatics at the Vienna University of Technology. Fall 2014.

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.

8.1.5.3. 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 served as:

President of the Selection Committee for the LIX Postdoctoral Fellowship. May - July 2014.

8.2. Teaching - Supervision - Juries

8.2.1. Teaching

PhD : Catuscia Palamidessi has been teaching a course for PhD students, on Quantitative Information Flow and on Differential Privacy at the department of Computer Science of the University of Pisa, Italy. April 2014. Total 20 hours.

Master : Frank D. Valencia has been teaching the masters course "Computability Theory", 60 hours, at the Pontificia Universidad Javeriana de Cali, Colombia. July 27 - Nov 1, 2014.

8.2.2. Supervision

PhD (2011-14) Nicolás E. Bordenabe. Ecole Polytechnique. Grant Inria/DGA. Co-supervised by Catuscia Palamidessi and Konstantinos Chatzikokolakis.

PhD in progress (2014-) Michel Guzman. Ecole Polytechnique. Grant Inria CORDI-S. 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.

PhD (2011-14) Luis Fernando Pino Duque. Ecole Polytechnique. Grant Inria/DGA. 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-) Lili Xu. Ecole Polytechnique and Chinese academy of Science, Beijing, China. Co-supervised by Catuscia Palamidessi and Huimin Li.

8.2.3. Juries

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

Xihui Chen (UNI-Lu, L). PhD thesis reviewer and member of the committee board at the PhD defense. Title of the thesis: *Location Assurance and Privacy in Location-based Services*. Advised by Sjouke Mauw. Defended in June 2014.

Meilof Veeningen (TUE, Eindhoven, The Netherlands). PhD thesis reviewer and member of the committee board at the PhD defense. Title of the thesis: *Objective Privacy – Formal analysis of data minimisation in privacy-enhancing protocols*. Advised by Sandro Etalle. Defended in June 2014.

Fabrizio Biondi (ITU, Copenhagen, Denmark). PhD thesis reviewer and member of the committee board at the PhD defense. Title of the thesis: *Markovian Processes for Quantitative Information Leakage*. Advised by Andrzej Wasowski. Defended in May 2014.

Tri Mihn Ngo (University of Twente, Italy). PhD thesis reviewer and member of the committee board at the PhD defense. Title of the thesis: *Qualitative and Quantitative Information Flow Analysis for Multi-threaded Programs*. Advised by University of Twente. Defended in April 2014.

Francesca Pampaloni (IMT, Lucca, Italy). PhD thesis reviewer and member of the committee board at the PhD defense. Title of the thesis: *Quantitative Models of Information Flow: Tuning the Power of the Adversary*. Advised by Michele Boreale. Defended in March 2014.

Frank Valencia has been reviewer for the thesis of the following PhD students:

Laura Titolo (University of Udine, Italy). Title of the thesis: *An Abstract Interpretation Framework for Diagnosis and Verification of Timed Concurrent Constraint Languages*. Advised by Marco Comini. Defended in May 12, 2014.

8.2.4. 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.

8.3. Popularization

Catuscia Palamidessi has organized the round table "Security & Privacy : challenges of the future digital society" at the Forum STIC, University of Paris-Saclay, December 2014.

Konstantinos Chatzikokolakis has been one of the speakers at the above round table.

Konstantinos Chatzikokolakis has given the popularization talk "Protection de la vie privée et anonymat" at the Journée ISN, Académie de Créteil, March 2014.

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] A. ARISTIZÁBAL, F. BONCHI, C. PALAMIDESSI, L. PINO, D. VALENCIA. *Deriving Labels and Bisimilarity for Concurrent Constraint Programming*, in "FOSSACS 2011 : 14th International Conference on Foundations of Software Science and Computational Structures", Saarbrücken, Germany, M. HOFMANN (editor), Lecture Notes in Computer Science, Springer, March 2011, vol. 6604, pp. 138-152 [*DOI :* 10.1007/ISBN 978-3-642-19804-5], https://hal.archives-ouvertes.fr/hal-00546722

[5] 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

[6] 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]

[7] 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/

[8] 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/

[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] N. E. BORDENABE. *Measuring Privacy with Distinguishability Metrics: Definitions, Mechanisms and Application to Location Privacy*, École Polytechnique, September 2014, https://pastel.archives-ouvertes.fr/tel-01098088

[12] L. F. PINO DUQUE. *Efficient Verification and New Reasoning Techniques for Concurrent Constraint Programming*, Université Paris-Saclay, September 2014, https://hal.inria.fr/tel-01110625

Articles in International Peer-Reviewed Journals

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

[14] K. CHATZIKOKOLAKIS, C. PALAMIDESSI, C. BRAUN. *Compositional Methods for Information-Hiding*, in "Mathematical Structures in Computer Science", 2014, https://hal.inria.fr/hal-01006384

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

[16] T. GIVEN-WILSON, D. GORLA, B. JAY. *A Concurrent Pattern Calculus*, in "Logical Methods in Computer Science", 2014, https://hal.inria.fr/hal-00987578

[17] L. PINO, A. ARISTIZÁBAL, F. BONCHI, F. D. VALENCIA. *Weak CCP bisimilarity with strong procedures*, in "Science of Computer Programming", 2015, 42 p. , forthcoming [*DOI :* 10.1016/J.SCICO.2014.09.007], https://hal.inria.fr/hal-00976768

[18] L. F. PINO DUQUE, F. BONCHI, F. VALENCIA. *Efficient Algorithms for Program Equivalence for Confluent Concurrent Constraint Programming*, in "Science of Computer Programming", 2015, 41 p. , forthcoming, https://hal.archives-ouvertes.fr/hal-01098502

Invited Conferences

[19] K. CHATZIKOKOLAKIS, D. GEBLER, C. PALAMIDESSI, L. XU. *Generalized bisimulation metrics*, in "CONCUR - 25th Conference on Concurrency Theory", Rome, Italy, P. BALDAN, D. GORLA (editors), Springer, September 2014, vol. 8704, pp. 32-46 [*DOI :* 10.1007/978-3-662-44584-6], https://hal.inria.fr/hal-01011471

International Conferences with Proceedings

[20] S. ALVIM, K. CHATZIKOKOLAKIS, A. MCIVER, C. MORGAN, C. PALAMIDESSI, G. SMITH. *Additive and multiplicative notions of leakage, and their capacities*, in "Computer Security Foundations", Vienna, Austria, IEEE, July 2014, forthcoming, https://hal.inria.fr/hal-00989462

[21] N. E. BORDENABE, K. CHATZIKOKOLAKIS, C. PALAMIDESSI. *Optimal Geo-Indistinguishable Mechanisms for Location Privacy*, in "CCS - 21st ACM Conference on Computer and Communications Security", Scottsdale, Arizona, United States, M. YUNG, N. LI (editors), ACM Press, November 2014, This paper is the report version of a work that appeared in the proceedings of the 21st ACM Conference on Computer Security. Scottsdale, Arizona, USA, Nov. 2014 (CCS'14) [*DOI :* 10.1145/2660267.2660345], https://hal.inria.fr/hal-00950479

[22] K. CHATZIKOKOLAKIS, C. PALAMIDESSI, M. STRONATI. *A Predictive Differentially-Private Mechanism for Mobility Traces*, in "PETS 2014 - 14th Privacy Enhancing Technologies Symposium", Amsterdam, Netherlands, July 2014, https://hal.inria.fr/hal-01011260

[23] T. CHOTHIA, Y. KAWAMOTO, C. NOVAKOVIC. *LeakWatch: Estimating Information Leakage from Java Programs*, in "Computer Security - ESORICS 2014 - 19th European Symposium on Research in Computer Security, Proceedings, Part II", Wroclaw, Poland, Lecture Notes in Computer Science, September 2014, vol. 8713, pp. 219 - 236 [*DOI :* 10.1007/978-3-319-11212-1_13], https://hal.inria.fr/hal-01097603

[24] N. FERNS, S. KNIGHT, D. PRECUP. *Bisimulation for Markov Decision Processes through Families of Functional Expressions*, in "Horizons of the Mind. A Tribute to Prakash Panangaden (for his 60th birthday)", Oxford, United Kingdom, F. VAN BREUGEL, E. KASHEFI, C. PALAMIDESSI, J. RUTTEN (editors), Horizons of the Mind. A Tribute to Prakash Panangaden, Springer, May 2014, vol. 8464, pp. 319-342 [*DOI :* 10.1007/978-3-319-06880-0_17], https://hal.inria.fr/hal-01098566

[25] T. GIVEN-WILSON. *An Intensional Concurrent Faithful Encoding of Turing Machines*, in "7th Interaction and Concurrency Experience (ICE 2014)", Berlin, Germany, June 2014, https://hal.inria.fr/hal-00987594

[26] T. GIVEN-WILSON. *Expressiveness via Intensionality and Concurrency*, in "ICTAC 2014 - 11th International Colloquium on Theoretical Aspects of Computing", Bucharest, Romania, September 2014, https://hal.inria.fr/hal-00999082

[27] T. GIVEN-WILSON. *On the Expressiveness of Intensional Communication*, in "Combined 21th International Workshop on Expressiveness in Concurrency and 11th Workshop on Structural Operational Semantics", Rome, Italy, September 2014, https://hal.inria.fr/hal-01026301

[28] Y. KAWAMOTO, K. CHATZIKOKOLAKIS, C. PALAMIDESSI. *Compositionality Results for Quantitative Information Flow*, in "Proceedings of the 11th International Conference on Quantitative Evaluation of SysTems (QEST 2014)", Florence, Italy, G. NORMAN, W. H. SANDERS (editors), Lecture Notes in Computer Science, Springer, September 2014, vol. 8657, pp. 368-383 [*DOI :* 10.1007/978-3-319-10696-0], https://hal.inria.fr/hal-01006381

[29] L. F. PINO DUQUE, F. BONCHI, F. D. VALENCIA. *A Behavioral Congruence for Concurrent Constraint Programming with Non-deterministic Choice*, in "Proceedings of the 11th International Colloquium Theoretical Aspects of Computing (ICTAC 2014)", Bucarest, Romania, September 2014, vol. 8687, pp. 351-368 [*DOI :* 10.1007/978-3-319-10882-7_21], https://hal.inria.fr/hal-01006382

[30] L. XU, K. CHATZIKOKOLAKIS, H. LIN. *Metrics for Differential Privacy in Concurrent Systems*, in "FORTE - 34th IFIP International Conference on Formal Techniques for Distributed Objects, Components and Systems", Berlin, Germany, E. ABRAHAM, C. PALAMIDESSI (editors), Springer, June 2014, vol. 8461, pp. 199-215 [*DOI :* 10.1007/978-3-662-43613-4_13], https://hal.inria.fr/hal-00879140

Scientific Books (or Scientific Book chapters)

[31] E. ELSALAMOUNY, K. CHATZIKOKOLAKIS, C. PALAMIDESSI. *Generalized differential privacy: regions of priors that admit robust optimal mechanisms*, in "Horizons of the Mind. A Tribute to Prakash Panangaden", F. VAN BREUGEL, E. KASHEFI, C. PALAMIDESSI, J. RUTTEN (editors), Lecture Notes in Computer Science, Springer International Publishing, 2014, vol. 8464, pp. 292-318 [*DOI :* 10.1007/978-3-319-06880-0_16], https://hal.inria.fr/hal-01006380

References in notes

[32] 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

[33] 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