Scientific Report

First name / Family name	GIUNTI
Nationality	Italy
Name of the *Host Organisation*	INRIA
First Name / family name of the *Scientific Coordinator*	PALAMIDESSI
Period of the fellowship	01/02/2011 to 31/01/2012

I – SCIENTIFIC ACTIVITY DURING YOUR FELLOWSHIP

During the research period that I have spent at INRIA in the Comete team at LIX – Ecole Polytechnique Palaiseau, under the supervision of Catuscia Palamidessi I have worked on developing abstractions and reasoning techniques for establishing the security of concurrent programs and of their implementation obtained by refinement.

Non-determinism and concurrency put challenging questions in the secrecy analysis of protocols. A crucial point is related to the benignity of the scheduling of interactions. The implementations of concurrent models obtained by refinement rely indeed on scheduling policies that typically consist in a proper subset of the non-deterministic executions that could be fired in the model. A demonic scheduler [1] chooses the worst alternatives in order to disclose the secrecy of the computations; on contrast, an angelic scheduler helps the system trying to preserve the secrecy of computations. The benignity of the scheduler could be related to its physical locations, i.e. global versus local scheduling. The global (demonic) scheduler is the attacker controlling the network while the local (angelic) scheduler is the trusted coordinator of the internal choices of components.

I believe the boundary among global and local scheduling to be related to the process calculi notion of restricted channel, which is reminding to the scope mechanism of programming languages while is more powerful because of mobility. Indeed, many mathematical formalisms to describe security protocols are based on the pi calculus; in such models the restriction operator (*new*) plays a fundamental role as it allows to hide the use of a channel. In other words, the *subject of the communication* can be accessed only by the processes in the scope of the operator. The scope, however, can be enlarged dynamically by sending the channel to remote processes. For this reason the actual implementation of the *new* operator usually relies on non-dedicated channels that may be insecure due, for instance, to side-channel attacks. One natural approach to cope with this problem is to map the private communication within the scope of the *new* into open communications protected by cryptography. The compilation of a pi calculus process would generate a cryptographic protocol where each channel is mapped into a couple of spi calculus cryptographic keys: a public and a private one. However, the distribution of capabilities is problematic to deploy in a trivial way and eventually leads to break the *forward secrecy* of communications [2] whenever decryption keys are leaked. While a solution is available [3], the price to pay is to cope with a complex cryptographic protocol that relies on a set of trusted authorities acting as proxies.

Based on these considerations, in [4] we argue that the restriction operator of the pi calculus does not adequately ensure confidentiality. To tackle this problem, we enrich the pi calculus with an operator for confidentiality (*hide*), whose main effect is to restrict the access to the *object of the communication*, thus representing confidentiality in a natural way. The *hide* operator is meant for local communication, and it differs from *new* in that it forbids the extrusion of the name and hence has a static scope. To emphasize the difference between *hide* and *new*, we introduce a *spy process* that represents a side-channel attack on the non-dedicated channels. In practice, *spy* is able to detect whether there has been a communication on one of the channels not protected by a *hide*, but is not able to retrieve the content of the communication.

We develop the reduction semantics of the new formalism, which we call s*ecret pi calculus*, and its observational equivalence. Then, we present a proof method, based on a labelled transition system and on bisimulation, and we prove its soundness and full

abstraction with respect to observational equivalence. Lastly, we discuss some algebraic equalities and inequalities of the *secret pi calculus*, and we analyze some interesting example, notably an implementation of name matching, and a deployment of mandatory access control that takes inspiration from the D-Bus technology used in the KDE graphical environment.

There are several extensions that I am keen to investigate. One ongoing work does consider the presence of a scheduler which controls the channels protected by *new* while cannot access to the channels protected by *hide*. The aim is to develop a framework to compare implementations of systems obtained by refinement through scheduling policies that ignore the interactions arising on channels programmed with *hide*, which are non-deterministic. The considered notion of observational equivalence takes into account the (demonic) scheduler that acts as an enemy trying to leak secrets, in the secure protocols terminology. By means of bisimulation semantics we provide for a co-inductive proof technique which can be used to show the security of such systems in the presence of a demonic scheduler. Another approach I am interested in relies on a typed analysis of local communications, i.e. those protected by *hide*. This would lead to a partially typed system that appears as more suitable for heterogeneous networks where the components do not belong to the same authority. Indeed, in these environments we cannot assume all components to be typed by the same type checker. Differently, a typed analysis could be feasible for local resources, for instance by combining light-weight static and dynamic typing [5].

References in notes

[1] K.Chatzikokolakis, C.Palamidessi: Making random choices invisible to the scheduler. Inf. Comput. 208(6): 694-715 (2010).
[2] M.Abadi: Protection in Programming-Language Translations. ICALP 1998: 868-883
[3] M.Bugliesi, M.Giunti: Secure implementations of typed channel abstractions. POPL 2007: 251-262
[4] M.Giunti, C.Palamidessi, Frank D.Valencia. A Process-algebraic Approach to Secrecy in Untrusted Networks. Technical report, 2012
[5] M.Bugliesi, M.Giunti: Typed Processes in Untyped Contexts. TGC 2005: 19-32

II – PUBLICATION(S) DURING YOUR FELLOWSHIP

All reports are available online:http://www.lix.polytechnique.fr/~marco.giunti/

● *A Process-algebraic Approach to Secrecy in Untrusted Networks*. Co-authored by Catuscia Palamidessi and Frank D. Valencia. PENDING

ABSTRACT

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

based on bisimulation semantics.

● *Disentangling typed deadlocked sessions*. Co-authored by Antonio Ravara. PENDING
ABSTRACT
Most session typing systems guarantee type safety, but not deadlock-freedom. The system of Giunti and Vasconcelos accepts deadlocked processes, but their algorithm, however, does not accept resource-holding ones. Herein we explore that gap to unblock some of those deadlocked processes by following their behaviour as specified by the session types. We define a process translation mechanism, directed by typing environments, such that a well-typed process with some channels blocked due to wrong sequentialization, is transformed in a similar process, type-safe, but with no such deadlocks on those channels. The procedure is thus correct, but not complete: there are deadlocked processes that are not released by it.
We view this procedure as a tool to help the programmer to solve some bugs: the typing algorithm rejects a process because it is deadlocked; the tool suggests a fix by looking into the session type; the programmer decides if makes sense to accept the transformation.

● *Linearity, session types and the pi calculus*. Co-authored by Vasco T. Vasconcelos. PENDING
ABSTRACT
We present a reconstruction of session types in a conventional pi calculus. Our session types are qualified as linear or unrestricted. Linearly typed communication channels are guaranteed to occur in exactly one thread, possibly multiple times; afterwards they evolve to unrestricted channels. We equip types with a constructor that describes the two ends of a same communication channel, and propose a type system that is sound with respect to the reduction relation. We then introduce an algorithmic type system which we prove sound and complete with respect to original system.
We assess the expressivity of our typing system by providing three distinct encodings (from the pi calculus with polarized variables, from the pi calculus with accept and request primitives, and from the linear pi calculus) into our system. For each language we present operational and typing correspondences, showing that our system effectively subsumes the linear pi calculus as well as foregoing works on session types.
In the case of the linear pi calculus we also provide a completeness result, thus proving that linear pi is a sublanguage of ours.

● *Typed observational equivalence for sessions*. PENDING
ABSTRACT
We propose a behavioural theory to contrast processes described by session type abstractions. We introduce a notion of typed observational equivalence for a polyadic pi calculus with matching where the discerning capability of the observer is regulated by the type checker. In particular, type checking forces contexts to not interfere with a session shared by two participants. Behaviourally equivalent pi calculus processes exhibit the same observables in all type checked contexts. To avoid universal quantification, we rely on a proof technique based on bisimulation over typed labelled semantics.
By establishing the soundness and the completeness of bisimulation semantics with respect to observational equivalence, we provide a framework to reason about the behaviour of service-oriented protocols.

● *A type checking algorithm for qualified session types*. ACCEPTED in the 7th Workshop on Automated Specification and Verification of Web Systems, WWV 2011:96-114.
ABSTRACT
We present a type checking algorithm for establishing a session-based discipline in the pi calculus of Milner, Parrow and Walker. Our session types are qualified as linear or unrestricted. Linearly typed communication channels are guaranteed to occur in exactly one thread, possibly multiple times; afterwards they evolve as unrestricted channels. Session protocols are described by a type constructor that denotes the two ends of one and the same communication channel. We ensure the soundness of the algorithm by showing that processes consuming all linear resources are accepted by a typing system preserving typings during the computation and that type checking is consistent w.r.t. structural congruence.

III – ATTENDED SEMINARS, WORKSHOPS, CONFERENCES

- Chair in Information Technology and Digital Sciences (2010-2011) – Martin Abadi - 10 March 2011 - College de France, Paris - France
- SECSI Colloquium, 17-18 March 2011, ENS Cachan, France
- Behavioural Types Workshop - 19-21 April 2011 – New University of Lisbon, Portugal
- 6th Federated Conferences on Distributed Computing Techniques (DisCoTec) – 6-9 June 2011 - Reykjavik, Iceland *(Supported by the Comete project-INRIA)*
- Forum Digiteo, 18 October 2011, Ecole Polytechnique Palaiseau, France
- Comete/Parsifal Seminars (Ecole Polytechnique Palaiseau, France)
    - Algebraic type systems (A. Diaz-Caro)- 9 March 2011
    - Sound and Complete Axiomatization of Trace Semantics for Probabilistic Transition Systems (A. Silva) - 11 March 2011
    - Keeping track of your friends and enemies: privacy threats of new mobile technologies (M. Arapinis) – 19 September 2011
    - On the (in)security of widely-used contactless smartcards (F. Garcia) – 22 October 2011

IV – RESEARCH EXCHANGE PROGRAMME (REP)

1. SPARCIM – Spain – IMDEA Software – Gilles Barthe – 16-22 January 2012
2. PEG – Portugal - Technical University of Lisbon – Pedro Adao and Ana Maria Matos – 23-29 January 2012

During the week spent at IMDEA Software I have encountered many people working in the team led by Gilles Barthe; they presented me their last efforts in developing a tool, named CertiCrypt/EasyCrypt, to assist the construction and verification of cryptographic proofs. The meetings with Gilles Barthe, Cesar Kuntz and Federico Olmedo have been particularly fruitful. I believe indeed that some features of their verification tool could be used in a more general setting that is of my interest, e.g. for proving behavioral properties of process specifications by means of co-inductive techniques. I also discussed with Boris Kopf the possibility of enhancing the bisimulation proof method in order to cope with information flow

attacks.

At the Instituto Superior Tecnico of the Technical University of Lisbon I have discussed with Pedro Adao and Ana Maria Matos my last work done in collaboration with Catuscia Palamidessi and Frank D. Valencia concerning the deployment of programming abstractions for secret channels. Particularly, we have analyzed how these abstractions can be used to enforce a mandatory access control policy in message-passing communicating systems by taking inspiration from the D-Bus technology currently used in many graphical desktop environments. The theme is related to previous work of Pedro Adao on cryptographically sound implementations of communicating processes, which we extensively discussed. We envisioned a collaboration on developing reliable (semi-)automated proofs of security protocols, which is the central theme of the ComFormCrypt project led by Pedro Adao. Lastly I have presented recent work done during the ERCIM fellowship [4] to the audience of the SQIQ-IT Information Security Seminar.