PACE project

ANR 12IS02001

Deliverable D2015-123

September 2015

Introduction. We give below a presentation of the results obtained during the third year of PACE, from january to september 2015. Contrarily to what was planned, this document was produced in september 2015 rather than in december 2015. The main reason for doing so is to synchronise with the report due for ANR after 30 months. Both documents have been prepared together; the present document provides a more detailed than the ANR report.

A description of the goals and organisation of the ANR PACE project (*"beyond plain Processes: Analysis techniques, Coinduction and Expressiveness"*) can be found at
`http://perso.ens-lyon.fr/daniel.hirschkoff/pace/docs/Submission/PACE.pdf`.
The website of the PACE project can be found at `http://perso.ens-lyon.fr/daniel.hirschkoff/pace`.

Life of the project

0.1 Personnel

Funded by PACE. Two postdoctoral researchers have been hired by the PACE project, to start working in october 2015.

- Charles Grellois will work at INRIA Sophia (EPI Focus, in Bologna), on topics related to Task 1.
- Matias Lee will work at ENS Lyon, on topics related to Tasks 1 and 3.

Contributing to PACE.

- Paul Brunet has started a PhD in Lyon, under supervision of D. Pous. His work is on Task 3.
- Raphaëlle Crubillé has started in september 2015 a PhD under co-supervision of U. dal Lago (in Bologna) and T. Ehrhard (in Paris). She works on topics related to Task 1 of the project.
- Salim Perchy has starded a PhD in Saclay, under supervision of F. Valencia and Stefan Haar (LSV, Cachan). His work is on Task 2.2.
- Michell Guzman has started a PhD in Saclay, under supervision of C. Palamidessi and F. Valencia. His work is on Task 1.5, continuing the work of Luis Pino.

0.2 Visits, meetings

Contrarily to in 2013 and 2014, there was no general meeting of the PACE project in 2015. A meeting is planned for the beginning of 2016, in Shanghai (as envisaged in the submission document of PACE).

Visits between PACE partners

- Ugo Dal Lago visited ENSL and gave a talk at the Chocola seminar (September 2014).
- 12-28 January 2015: Ugo Dal Lago visited Fu's group in Shanghai, and collaborated with Yuxin Deng.
- 15-16 July : Sangiorgi has visited ENS Lyon, collaborating with A. Durier and D. Hirschkoff

Other visits

- 25 July-15 August 2015: Vignudelli has visited Prakash Panangaden's group in Montreal.
- Filippo Bonchi visited Duisburg University from 16/05/2015 to 22/05/2015 to work with Barbara Koenig and Daniela Petrisan about Behavioural Metrics.

- D. Pous invited Insa Stucke (Kiel University) in july 2015 and september 2015, to collaborate on scientific questions related to Task 3.
- Fabio Zanasi visited University of Southampto from 10/07/2013 to 31/07/2013 to work with Pawel Sobocinski on "Interacting Hopf Algebras".
- D. Pous and P. Brunet visited A. Silva and D. Petrisan in Nijmegen in june 2015.
- Filippo Bonchi visited Pisa University from 01/04/2015 to 01/08/2015 to work with Ugo Montanari, Roberto Bruni and Fabio Gadducci about Signal Flow Graphs.

We describe below the contributions obtained in PACE. A list of the publications of PACE, which includes the papers we mention, can be accessed at `http://perso.ens-lyon.fr/daniel.hirschkoff/pace/papers.html`.

1 Task 1: Advanced Coinductive Techniques

1.1 T1.1: Up-to techniques

Up-to techniques and equations. We have explored alternative techniques to the bisimulation proof method. The new techniques are based on unique solutions of special forms of inequations called contractions, and inspired by Milner's theorem on unique-solution of equations. The method is at least as powerful as the bisimulation proof method and its 'up-to context' enhancements. Moreover, it can be transferred onto other behavioural equivalences, possibly contextual and non-coinductive. This enables a coinductive reasoning style on such equivalences. Paper: Sangiorgi POPL 15

A Categorical Perspective on Up-To techniques. Together with Daniela Petrisan and Jurriaan Rot, F. Bonchi and D. Pous have developed a fibrational perspective on coinduction up-to that allows to modularly prove the soundness of differents sorts of (a) up-to techniques, (b) coinductive predicates (e.g., unary predicates like modal logic formulas or binary predicates like equivalence relations, preorders) and (c) state based systems (such as process calculi, weighted and nominal automata). From this perspective both up-to techinques and coinductive predicates are represented as endo-functors on a certain fiber. The soundness of a technique w.r.t. a predicate is proved by exhibiting a distributive law between the corresponding endo-functors. A proof up-to is just a bialgebra for such distributive law. This work has been presented at LICS-CSL 2014. In CONCUR 2015, we further extend this approach to deal with lax-bialgebras which, as we show in this work, are necessary to model weak semantics, like weak or dynamic bisimilarity.

1.2 T1.2: From equivalences to metrics

Bisimilarity Pseudometric. Regarding metrics, the bisimilarity pseudometric based on the Kantorovich lifting is one of the most popular metrics for probabilistic processes proposed in the literature. However, its application in verification is limited to linear properties. We proposed a generalization of this metric which allows to deal with a wider class of properties, such as those used in security and privacy, in the following work:

Konstantinos Chatzikokolakis, Daniel Gebler, Catuscia Palamidessi, and Lili Xu. Generalized bisimulation metrics. In Paolo Baldan and Daniele Gorla, editors, CONCUR - 25th Conference on Concurrency Theory, volume 8704 of Lecture Notes in Computer Science, pages 32?46, Rome, Italy, Septem ber 2014. Springer.

Behavioural Metrics, and Associated Up-to Techniques . Together with Paolo Baldan, Barbara Koenig and Henning Kerstan, F. Bonchi has developed a categorical framework that allows to automatically derive behavioural pseudo-metrics for a large variaty of state-based systems modeled as coalgebras. The approach is based on two alternative approaches to "lift" functors from Sets to Pseudo-Metrics Spaces. This work, which has been published at FSTTCS 2014, always gives rise to metrics wich are appropriate for Branching Time Semantics. In CALCO 2015, we extended it to deal with Behavioural Metrics for Linear Time Semantics. The switch from Branching Time to Linear Time Semantics is realized via the "generalized powerset construction" which requires to lift to Pseudo-Metrics Spaces not only functors, but also monads and distributive laws.

The categorical machinery used to deal behavioural metrics is a special case of the one exploited for up-to techniques, as described in the work in Task 1.1. Indeed, Pseudo-Metric Spaces form a fibration over Sets. Therefore, the use of our fibrational perspective seems the right way to distill up-to techniques for behavioural metrics. This is an ongoing work, as per august 2015, in collaboration between PACE members, Barbara Koenig and Daniela Petrisan.

Metrics for Higher-Order Languages. Comparing terms by way of equivalences, like context equivalence, can be too discriminating in a probabilistic scenario. This is the starting point of the work developed in Bologna by U. dal Lago on metrics for higher-order languages. The first step has been the study of trace and bisimulation metrics in an affine lambda calculus, in which copying is forbidden (CrubilléDalLago). Then, trace metrics have been shown to work well in a lambda-calculus capturing probabilistic polynomial time computation, paving the way for a new proof technique for computational indistinguishability (CappaiDalLago).

1.3 T1.3: Probabilistic and quantum higher-order languages

Probabilities and Higher-Order Languages. We have initiated the study of probabilities in higher-order languages. We have begun with the pure lambda-calculi, first call-by-name and then call-by-value.. We have considered applicative bisimilarity and context equivalence, with the objective of finding techniques for proviing congruence of the bisimilarity and then comparing the two equivalences. The main results include: congruence techniques for bisimilarity that follows Howe's method, enhancing it with on non-trivial "disentangling" properties for sets of real numbers; a detailed comparison between bisimilarity and context equivalence, both in call-by-name and in call-by-value, as well as a similar comparison for the respective preorders (similarity and the contextual preorder). Papers: Dal Lago et al ESOP 2014, POPL 2014.

We have studied the discriminating power offered by higher-order concurrent languages, and contrasted this power with those offered by higher-order sequential languages (à la lambda-calculus) and by first-order concurrent languages (à la CCS). The comparison is carried out by providing embeddings of first-order processes into the various languages, and then examining the resulting contextual equivalences induced on such processes. As first-order processes we have considered both ordinary Labeled Transition Systems (LTSs) and Reactive Probabilistic Labeled Transition Systems (RPLTSs). paper: Bernardo-Sangiorgi-Vignudelli LICS'14

On a related line of work, we have studied the possibilities of discriminanting reactive probabilistic processes using three forms of tests, which respectively allow only probabilities, only non-determinism, both probabilities and non-determinism. An important feature of these tests is that none of them has copying facilities, which are known to make the resulting testing equivalence coincide with probabilistic bisimulation. Hierarchies among all such equivalences have been established.

Moreover, number of facts are provided that lead us to conjecture that (i) may testing and must testing coincide on reactive probabilistic processes and (ii) nondeterministic and probabilistic tests reach the same discriminating power as probabilistic bisimilarity. paper: Bernardo-Sangiorgi-Vignudelli Quest'14

We have investigated the nature of probabilistic computation in the context of function algebras. More specifically, we have introduced the notion of a computable probabilistic function as a function from natural numbers to distributions of natural numbers which can be computed by a probabilistic Turing machine. This has later been characterized by a variation on the Kleene's algebra for functions. The same correspondence carries over to polytime computable functions, where the function algebra becomes one in the style of Leivant's tiered recursion. [DalLagoZuppiroli2014,DalLagoZuppiroliGabbrielli2014].

Quantum Lambda-calculi. Applicative bisimulation has been shown to work well in quantum lambda calculi. This has been done in two steps. First, it has been shown that purely linear quantum lambda calculi are amenable to coinductive techniques (DalLagoRioli), but that this does not give rise to a full abstraction result. Then, a more expressive lambda-calculus for quantum computation in which copying of classical values is possible has been shown to support state-based and distribution-based bisimulation, the latter being fully abstract (DengDalLagoFeng).

1.4 T1.4: Quantum processes

Yuxin Deng investigated bisimulation semantics of probabilisitic and quantum processes. In particular, his book "Semantics of Probabilistic Processes: An Operational Approach" was published by Springer

(www.springer.com/978-3-662-45197-7), which summarizes some recent progress in probabilistic concurrency theory.

1.5 T1.5: Spatial-epistemic CCP

Luis Pino, Filippo Bonchi and Frank Valencia have proved that weak saturated bisimilarity for ccp (concurrent constraint programming) fails to be a congruence in the presence of non-deterministic choice (even for guarded choice). They proposed a new co-inductive relation called weak full bisimilarity and then showed that it coincides with closure under arbitrary contexts of weak-saturated bisimilarity. The new relation does not require quantification over all possible contexts. This work was published as:

- A Behavioural Congruence for Concurrent Constraint Programming with Non-Deterministic Choice. ICTAC 2014 - 11th International Colloquium on Theoretical Aspects of Computing (2014).

We also showed that standard reduction of weak to strong bisimilarity does not work for ccp and provided a novel reduction that admits efficient procedures in:

- Weak CCP Bisimilarity with Strong Procedures. Science of Computer Programming. 2014.

More recently we provided new co-inductive techniques and associated efficient algorithms to check notions of behavioral equivalence for ccp in

- Efficient Algorithms for Program Equivalence for Confluent Concurrent Constraint Programming. Science of Computer Programming, 2015.

Moreover, in collaboration with Fabio Gadducci, Francesco Santini we provide a generalization of ccp label and bisimulation semantics to account for Soft ccp. This was published as:

- A Labelled Semantics for Soft Concurrent Constraint Programming. COORDINATION 2015: 133-149

Finally, in the following paper we used abstract interpretation semantics for ccp to exhibit a secrecy flaw in a security protocol.

- M. Falaschi, C. Olarte, C. Palamidessi. Abstract Interpretation of Temporal Concurrent Constraint Programs. Theory and Practice of Logic Programming CUP. Theory and Practice of Logic Programming, 15(3): 312-357 (2015)

2 Task 2: Expressiveness

2.1 T2.1: Absolute theory

Expressiveness in Process Calculi Regarding expressiveness in process calculi, in the following paper we introduced the 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. This generalized form of pattern matching is refereed to as intentionality. Many popular process calculi can be encoded in CPC; this allows for a gain in expressiveness, formalised through encodings.

- T. Given-Wilson; D. Gorla; B. Jay. A Concurrent Pattern Calculus. Logical Methods in Computer Science (2014).

Furthermore in the following paper we presented an approach to encoding Turing Machines into intensional process calculi that is faithful, reduction preserving, and structurally equivalent.

- T. Given-Wilson. An Intensional Concurrent Faithful Encoding of Turing Machines. 7th Interaction and Concurrency Experience (ICE 2014) (2014).

In the following paper we explored further the expressiveness of intensionality for concurrent and sequential computation. In particular, by means of possibility/impossibility of encodings, we showed in the following paper that intensionality alone can encode synchronism, arity, communication-medium, and pattern-matching, yet no combination of these without intensionality can encode any intensional language.

- T. Given-Wilson. Expressiveness via Intensionality and Concurrency. ICTAC 2014 - 11th International Colloquium on Theoretical Aspects of Computing (2014).

- Thomas Given-Wilson. On the Expressiveness of Intensional Communication. In Johannes Borgstrom and Silvia Crafa, editors, Combined 21th International Workshop on Expressiveness in Concurrency and 11th Workshop on Structural Operational Semantics, volume 160 of Electronic Proceedings in Theoretical Computer Science, pages 30?46, Rome, Italy, September 2014. Open Publishing Association.

Relating Denotational Models for Higher-Order Computation and Process Calculi Levy-Longo and Bohm Trees are the best known tree structures on the lambda-calculus, and correspond to forms of bisimulation on open lambda-terms. We have studied general conditions under which an encoding of the lambda-calculus into the pi-calculus is sound and complete with respect to such trees. We have applied these conditions to various encodings of the call-by-name lambda-calculus, showing how the two kinds of tree can be obtained by varying the behavioural equivalence adopted in the pi-calculus and/or the encoding. The conditions can can be adapted to other concurrency formalisms. Paper: Sangiorgi Xian, CONCUR 2014

Interacting Hopf Algebras . Streams are one of the most familiar examples of coinductive objects: mathematical series like the Fibonacci numbers are studied by undergraduates all over the worlds. F. Bonchi, Pawel Sobocinski and Fabio Zanasi have introduced an algebra for the specification, the implementation and the analysis of stream-processing circuits, called "the Calculus of Signal Flow Graphs". They have given a syntax, a denotational semantics, a sound and complete axiomatization, a full abstraction theorem and a realizability result. Interestingly enough, the axiomatization shows that the operators of this calculus forms two Hopf algebras that interact according to some distributive laws.

2.2 T2.2: Expressiveness in social networks

Frank Valencia and his PhD student Salim Perchy been working on three sub-tasks related to Task 2.2. The first task is the development a new family of algebraic structures to express spatial concepts such as mobility and epistemic notions such as beliefs, lies and opinions which are ubiquitous in social networks. The second task is the extension of a logic of beliefs with a new modality to expresss "utterences". This allows us to specify agents that interact by exchanging facts, beliefs and even intensional lies ("hoaxes"). The ability to express hoaxes, a common phenomena in social networks, is the main novelty of our logic. The third task is definition of the semantics of a concurrent programming language where processes interact by asking and posting inforimation, and moving across spatial hierarchies. The semantics domain for this language arises from the algebraic structure in the first task. The first two tasks have been recently published as

- An algebraic view of space/belief and extrusion/utterance for concurrency/epistemic logic. PPDP 2015: 161-172

2.3 T2.3: Applications to privacy, confidentiality and anonymity

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 with the inference of the secret information. In

- K. Chatzikokolakis, C. Palamidessi, C. Braun. Compositional Methods for Information-Hiding. Mathematical Structures in Computer Science, Cambridge University Press, 2014.

we considered a probabilistic process calculus to specify such systems, and we study how the operators affect the probability of error. (This work is also related to Task 1.3.)

Similarly, 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 the following work we identified optimal mechanisms for generalized differential privacy, i.e. mechanisms that maximize the utility for a given level of privacy.

- E. ElSalamouny. K. Chatzikokolakis, C. Palamidessi. Generalized differential privacy: regions of priors that admit robust optimal mechanisms. Horizons of the Mind. A Tribute to Prakash Panangaden Springer International Publishing (Ed.) (2014) 292-318.

In the paper:

- Nicolas E. Bordenabe, Konstantinos Chatzikokolakis, and Catuscia Palamidessi. Optimal Geo-Indistinguishable Mechanisms for Location Privacy. In Gail-Joon Ahn, Moti Yung, and Ninghui Li, editors, CCS - 21st ACM Conference on Computer and Communications Security, pages 251-262, Scottsdale, Arizona, United States, November 2014. Gail-Joon Ahn, ACM.

we showed that it is possible to combine the advantages of two approaches that have been proposed to limit and control the privacy loss: one is the geo-indistinguishability notion of Andrés et al (a formal notion of privacy that protects the user's exact location) and 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.

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. In the following paper we derived bounds on the g-leakage of the whole system in terms of the g-leakage of its components.

- Kawamoto, Yusuke and Chatzikokolakis, Konstantinos and Palamidessi, Catuscia. Compositionality Results for Quantitative Information Flow. Proceedings of the 11th International Conference on Quantitative Evaluation of SysTems (QEST 2014).

In the paper below 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 considered a notion of utility based on identity gain functions, which is closely related to min-entropy leakage, and we derive bounds for it. Finally, given some graph symmetries, we provide a mechanism that maximizes utility while preserving the required level of differential privacy.

- M. Alvim, M. Andrés, K. Chatzikokolakis, P. Degano, C. Palamidessi. On the information leakage of differentially-private mechanisms. Journal of Computer Security. 2014, to appear.

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 the paper below 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.

- Chatzikokolakis, Konstantinos and Palamidessi, Catuscia and Stronati, Marco. A Predictive Differentially-Private Mechanism for Mobility Traces. Proceedings of the 14th International Symposium on Privacy Enhancing Technologies (PETS 2014).

In the following paper we presented two mechanisms for achieving geo-indistinguishability, one generic to sanitize locations in any setting with reasonable utility, the other custom-built for a limited set of locations but providing optimal utility.

- Konstantinos Chatzikokolakis, Catuscia Palamidessi, and Marco Stronati. Geo-indistinguishability: A Principled Approach to Location Privacy. In Gautam Barua Raja Natarajan and Manas Ranjan Patra, editors, ICDCIT 2015 - Proceedings of the 11th International Conference on Distributed Computing and Internet Technology , volume 8956 of Lecture Notes in Computer Science, pages 49-72, Bhubaneswar, India, February 2015. Springer.

Our main contribution in the following paper is to present how scheduling and observation effect information leakage properties. In particular, we show that scheduling can hide some leaked information from perfect observers, while some scheduling may reveal secret information that is hidden to imperfect observers. In addition we present an algorithm to construct a scheduler that minimises the min-entropy leakage and min-capacity in the presence of any observer.

- Yusuke Kawamoto and Thomas Given-Wilson. Quantitative Information Flow for Scheduler-Dependent Systems. In The 13th International Workshop on Quantitative Aspects of Programming Languages and Systems (QAPL 2015), Electronic Proceedings in Theoretical Computer Science, page to appear, London, United Kingdom, April 2015. Open Publishing Association.

Finally, we overviewed quantitative approaches to privacy and open problems in:

- Catuscia Palamidessi. Quantitative Approaches to the Protection of Private Information: State of the Art and Some Open Challenges. Proc. of Principles of Security and Trust 2015: 3-7 (T2.3)

3 Task 3: Analysis techniques

The work by F. Bonchi and D. Pous on automata algorithms based on coinductive techniques was selected by the ACM SIGPLAN committee to appear in the "research highlights" section of Communications of the ACM. It appeared as the cover page of

> F. Bonchi, D. Pous, Hacking Nondeterminism with Induction and Coinduction. Research Highlights in Communications of the ACM, Vol. 58 No. 2, 2015.

together with a video interview.

The Relevance of Branching Bisimilarity. We have studied branching bisimilarity from the point of view of decidability and complexity properties. This kind of results is useful when trying to come up with algorithms for the verification of behavioural equivalences. A promising line of works has emanated from this approach, comprising notably papers at ICALP'13 and ICALP'14 by Y. Fu and co-authors, and a paper at LICS'15 by C. He and M. Huang. These works suggest that branching bisimilarity, which can be seen as a refinement of (coinductively defined) weak bisimilarity, appears to enjoy better properties, and makes it possible to provide better solutions to questions for which weak bisimilarity raises important technical difficulties.

The Representation of Labelled Transition Systems for Higher-Order Languages in Theorem Proving Environments. J.-M. Madiot, D. Pous, and D. Sangiorgi recently transported the abstract theory of up-to techniques onto languages whose bisimilarity and LTS go beyond those of first-order models (CONCUR'14). Their approach consists in exhibiting fully abstract translations of the more sophisticated LTSs and bisimilarities onto the first-order ones.

Algorithms for Enriched Kleene Algebras. D. Pous proposed algorithms for checking language equivalence of finite automata over a large alphabet, using symbolic automata where the transition function is compactly represented using (multi-terminal) binary decision diagrams (BDD). He applied such ideas to Kleene algebra with tests (KAT), the resulting work has been presented at POPL'15.

With his student P. Brunet, D. Pous showed that the equational theory of Kleene algebra with converse is PSPACE-complete (RAMICS'14, J.LAMP '15). They pursued their exploration of relation algebra fragments and they recently obtained a characterisation of the equational theory of Kleene allegories (Kleene algebra extended with converse and intersection), which lead to a new decidability result. This work has been presented at LICS'15.

Algorithmic Properties of Automata with Clocks. G. Li and co-authors investigates timed recursive models under different types of clocks, such as local clocks, global clocks and frozen clocks. They have proved that the reachability of a timed pushdown system with local clocks and frozen clocks and a unique global clock is decidable, while that with two global clocks is undecidable. Furthermore, they plan to take updatable clocks into account. They have proved that a timed automaton with one updatable clock and diagonal-free constraint is decidable on reachability, while that with diagonal constraint is undecidable. The future work will include the investigation of timed recursive model with one updatable clock and diagonal-free constraint.

Tasks and subtasks in the PACE project

Task 1: Advanced Coinductive Techniques

Task leader: Davide Sangiorgi / Deputy task leader: Xu Xian

- T1.1: Up-to techniques
- T1.2: From equivalences to metrics
- T1.3: Probabilistic and quantum higher-order languages
- T1.4: Quantum processes
- T1.5: Spatial-epistemic CCP

Task 2: Expressiveness

Task leader: Fu Yuxi / Deputy task leader: Catuscia Palamidessi

- T2.1: Absolute theory
- T2.2: Expressiveness in social networks
- T2.3: Applications to privacy, confidentiality and anonymity

Task 3: Analysis techniques

Task leader: Damien Pous / Deputy task leader: Deng Yuxin

- T3.1: Algorithms relying on up-to techniques
- T3.2: Up-to techniques in algorithms for metrics
- T3.3: Algorithms for quantum bisimulations
- T3.4: Minimization algorithms for symbolic bisimulation