Collection of Abstracts for the Workshop of REPAS (REliable and Privacy-Aware Software systems)

ANR-16-CE25-0011

13th and 14th of June 2017, Ljubljana

Francesco Gavazzo: An Abstract Account to Applicative Bisimulation Metrics. (Joint work with Raphaëlle Crubillé, Ugo Dal Lago.) Following our previous work on effectful applicative bisimilarity, we present some preliminary results concerning the study of applicative bisimulation metrics in an abstract setting. Concretely, we will investigate applicative bisimulation metrics in the context of (linear and classical) call-by-value λ-calculi with algebraic effects modeled through monads and operations. To this end, we first generalise the relational and categorical apparatus based on monads and relator to take into account quantitative information, thus employing basic notions and ideas from coalgebra and monoidal topology. We then try to single out those monads and relators for which applicative bisimilarity distance is a sound methodology for the so-called contextual distance. Moving from a qualitative to a quantitative setting allows us to develop a general framework in which it is possible to study both behavioural equivalences and metrics in a uniform abstract setting. However, such a generalisation comes with a price to pay, raising several problems (both from a mathematical and a computational perspective), some of which are still open.

Prakash Panangaden: Bicategories of Markov Processes. (Joint work with Florence Clerc, Harrison Humphrey.) We construct bicategories of Markov processes where the objects are input and output sets, the morphisms (one-cells) are Markov processes and the two-cells are simulations. This builds on the work of Baez, Fong and Pollard, who showed that a certain kind of finite-space continuous-time Markov chain (CTMC) can be viewed as morphisms in a category. This view allows a compositional description of their CTMCs. Our contribution is to develop a notion of simulation between processes and construct a bicategory where the two-cells are simulation morphisms. Our version is for processes that are essentially probabilistic transition systems with discrete time steps and which do not satisfy a detailed balance condition. We have also extended the theory to continuous space processes.

Barbara König: Behavioural Metrics via Functor Lifting - A Coalgebraic Approach. (Joint work with Paolo Baldan, Filippo Bonchi, Henning Kerstan and Daniela Petrisan.) Behavioural metrics provide us with a qualitative notion of bisimilarity, which allows to measure the behavioural distance of two states. The aim of this talk is to give a uniform account of bisimilarity metrics in a coalgebraic setting. In order to achieve this, we will lift functors from Set to pseudo-metric spaces. We consider two forms of such functor liftings: the Kantorovich and the Wasserstein lifting, which are inspired by the corresponding metrics on probability distributions. In addition, functor lifting has one parameter: an evaluation function, related to predicate liftings of coalgebraic modal logic. We will discuss the role of this parameter and the canonicity of the lifting. We will review metric transition systems by de Alfaro, Faella and Stoelinga and pseudometrics for probabilistic transition systems by van Breugel and Worrell and show how they arise as instances of our framework.

Daniela Petrisan: Is the Kantorovitch lifting canonical? (Joint work with Filippo Bonchi and Barbara Knig.) In this talk we discuss behavioural metrics obtained via functor liftings in a fibred setting. In particular, we discuss canonicity and universal properties of these liftings.

Giorgio Bacci: On the Metric-based Approximate Minimization of Markov Chains. We address the behavioral metric-based approximate minimization problem of Markov Chains (MCs), i.e., given a finite MC and a positive integer k, we are interested in finding a k-state MC of minimal distance to the original. By considering as metric the bisimilarity distance of Desharnais at al., we show that optimal approximations always exist; show that the problem can be solved as a bilinear program; and prove that its threshold problem is in PSPACE and NP-hard. Finally, we present an approach inspired by expectation maximization techniques that provides suboptimal solutions. Experiments suggest that our method gives a practical approach that outperforms the bilinear program implementation run on state-of-the-art bilinear solvers.

Bart Jacobs: Disintegration. Disintegration is a basic technique in probability for extracting a channel (conditional probability, Kleisli map) from a joint state. This talk describes the basics of disintegration, and also how it is used in many situations: going back and forth between a Bayesian network and a joint state, Bayesian belief updates, especially for point observations in continuous probability, naive Bayesian classifiers, etc. Disintegration has been implemented in the EfProb library; this will be used for demonstrations.

Ugo Dal Lago: Toward Higher-order Cryptography. (Joint work with Raphaëlle Crubillé.) We report on our ongoing effort to give a framework, based on game semantics, for the specification of (higher-order) complexity bounded probabilistic programs. The framework allows to specify many constructions in state-of-the art cryptography, but also new ones. A key notion turns out to be that of computational indistinguishability, which generalizes the one in classic cryptography.

James Worrell: The Expressiveness of Metric Temporal Logic. Several versions of Metric Temporal Logic have been used by researchers to express privacy and security policies on event traces and audit logs. In this tutorial we take a foundational approach to measuring the expressiveness of temporal logics, using monadic first-order logic as a yardstick. We highlight the notion of separation, due to Dov Gabbay, and its use in proving the expressiveness completeness of linear temporal logic over Dedekind-complete linear orders. We generalize this notion to the metric setting and show how it can be used to show the expressive completeness of a core fragment of Metric Temporal Logic that extends linear temporal logic by augmenting the Until and Since modalities with time constraints.

Vincent Danos: Consesus. I will present a family of probabilistic algorithms obtaining consensus on immutable and growing data structures with adversarial behaviours. We will see that growth alters fundamentally the speed security trade-off of flat consensus.

Catuscia Palamidessi: Verification of Differential Privacy in Concurrent Systems. Differential Privacy (DP) is one of the most successful approaches to privacy protection in statistical databases. It provides a formal privacy guarantee, ensuring that sensitive information relative to individuals cannot be easily inferred by disclosing answers to aggregate queries.

If two databases are adjacent, i.e. differ only for the value an individual?s data, then the query should not allow to tell them apart by more than a certain factor. This induces a bound also on the distinguishability of two generic databases, which is determined by their distance on the Hamming graph of the adjacency relation. Recently, we have proposed a generalized version of DP that can be applied to arbitrary metric domains, by expressing the indistinguishability requirement in terms of a bound on the given distance. In this talk, we consider the problem of verifying that two probabilistic concurrent processes that differ for the value of a secret satisfy (generalized) DP, i.e., they give raise to observable traces whose distance does not exceed the required bound. To this purpose, we consider an extension of the Bisimulation Metrics based on the Kantorovich distance. However, the standard Kantorovich lifting is not suitable for capturing the metric of distribution expressing the DP property. We therefore explore a generalized notion of Kantorovich lifting, suitable for arbitrary metric domains, and therefore also for the generalized DP. We show that the standard results extend smoothly to the generalized case, and that a bound on the generalized bisimulation distance is also a bound for the distance on traces, which guarantees the soundness of the method for proving DP. Finally, compare the efficiency of computing the Kantorovich lifting for DP with the one of the standard metric.

Marco Gaboardi: A semantic account of metric preservation in presence of probabilities. Program sensitivity measures how robust a program is to small changes in its input, and is a fundamental notion in domains ranging from differential privacy to cyber-physical systems. A natural way to formalize program sensitivity is in terms of metric preservation with respect to the metrics on the input and output spaces: requiring that an r-sensitive function maps inputs that are at distance d to outputs that are at distance at most r×d. I will present some recent work where we started the study of program sensitivity and metric preservation from a denotational point of view. I will introduced metric CPOs, a novel semantic structure for reasoning about computation on metric spaces, by endowing CPOs with a compatible notion of distance. This structure is useful for reasoning about metric properties of programs, and specifically about program sensitivity and metric preservation. I will then discuss how to extend this structure to account for metrics over probability distributions. In particular, I will discuss how these metrics can be used to guarantee programs differentially private.