

Preface to the Special Issue on Coordination Models and Languages (COORDINATION 2017)

Modern information systems rely increasingly on combining concurrent, distributed, mobile, adaptive, reconfigurable and heterogenous components. As a result, new models, architectures, languages and verification techniques are necessary to cope with the complexity induced by the demands of today's software development. In particular, many new driving applications are emerging out of the Internet of Things, workflow management, blockchain technology and service oriented approaches. Further challenges are arising concerning security, performance and privacy aspects, especially when considering coordination in large scale collective systems that involve humans as integrated part of such systems. Coordination languages have emerged as a promising approach, in that they provide abstractions that cleanly separate behaviour from communication, therefore increasing modularity, simplifying reasoning, and ultimately enhancing software development.

This special issue is dedicated to papers selected from the 19th International Conference on Coordination Models and Languages (COORDINATION 2017), which took place as part of the federated conferences on Distributed Computing Techniques (DisCoTec 2017), held in Neuchâtel, Switzerland, 19-21 June 2017. It contains extended versions of the best papers which have undergone additional rounds of detailed anonymous peer reviews. Eight papers have been selected out of this process. They report on theoretical and practical results on a number of coordination issues, which include topics such as replicated data types, optimal gradient-based algorithms for self-organisation in pervasive computing, attribute-based communication, theoretical foundations for web-services based on testing theory, information flow policies, behavioural contracts, stochastic analysis of business processes and session oriented programming.

The article *On the Semantics and Implementation of Replicated Data Types* by Fabio Gadducci, Hernan Melgratti and Christian Roldan addresses the problem of distributed data stores that maintain copies of the same item on multiple devices. It recasts the standard approach in terms of a denotational framework in which a data type is a function mapping visibility into admissible arbitrations. This characterisation accommodates underspecification and refinement and it enables a direct characterisation of correct implementations of a replicated data type.

The article *Optimal Single-Path Information Propagation in Gradient-based Algorithms* by Giorgio Audrito, Ferruccio Damiani and Mirko Viroli considers self-organisation through peer-to-peer interactions in scenarios such as pervasive computing and the Internet of Things. In that setting, being able to compute the estimated distances between, possibly mobile, devices is essential but very challenging. This paper presents "BIS (Bounded Information Speed) gradient", a fully distributed, optimal algorithm for this purpose and compares it with state-of-the-art approaches in the literature.

The article *AErlang: Empowering Erlang with Attribute-based Communication* by Rocco De Nicola, Tan Duong, Omar Inverso and Catia Trubiani provides a novel mechanism to dynamically select groups of communicating entities by relying on predicates over their attributes, so-called attribute based communication. In this paper the impact of this communication mechanism is studied by enriching the concurrency-oriented language Erlang with attribute-based primitives and applying it to several case studies.

The article *Full-abstraction for Client Testing Preorders* by Giovanni Bernardi and Adrian Francalanza contributes to the theoretical foundations for web-services, where processes are interpreted as servers, and tests as clients. A novel coinductive characterisation of client preorders

is provided, based on the notion of must-testing and a compliance relation, which is used to find a confirmative answer to whether these preorders are decidable.

The article *Many-to-Many Information Flow Policies* by Alberto Lluch Lafuente and Paolo Baldan proposes a language for expressing information flow policies requiring different security levels. It characterizes it through the study of a semantic foundation based on causal models of computation and explores its practical applicability by focusing on systems specified as safe Petri nets.

Behavioral contracts are abstract descriptions of expected communication patterns followed by either clients or servers during their interaction. The article *A Theory of Retractable and Speculative Contracts* by Franco Barbanera, Ivan Lanese and Ugo de'Liguoro studies two categories of behavioral contracts: so-called retractable contracts dealing with backtracking and so-called speculative contracts dealing with speculative execution. The authors show that the two types of contracts induce the same notion of compliance and, consequently, the same notion of subcontract relation, and that both notions can be decided in quadratic time.

The article *Stochastic Analysis of BPMN with Time in Rewriting Logic*, by Francisco Durán, Camilo Rocha and Gwen Salaün, employs executable specifications written in rewriting logic to model and analyze BPMN processes involving time and quantitative aspects. Their method allows for discrete-event simulation and automatic stochastic verification of properties, in particular through the use of Maude's statistical model checker PVeStA.

The article *Session-ocaml: a Session-based Library with Polarities and Lenses* by Keigo Imai, Nobuko Yoshida and Shoji Yuen proposes a novel library for session-typed concurrent and distributed programming in OCaml, with as two main ingredients, on the one hand, polarised session types to provide an alternative formulation of duality, and, on the other hand, parameterised monads with a data structure called 'slots' manipulated with lenses, to statically enforce session linearity including delegations.

The Guest Editors wish to thank all authors that contributed to this special issue for their submissions and careful revisions of their articles, and all reviewers for their detailed and constructive feedback suggesting valuable improvements. We are also grateful to the Editor-in-Chief Jan Bergstra and the members of the Editorial Board for giving us the opportunity to publish this special issue, and to the Editorial Assistant Bas van Vlijmen who guided us through a smooth reviewing process. We hope you will enjoy this selection of research articles and that they will be a source of inspiration for future work.

Jean-Marie Jacquet

Université de Namur, Namur, Belgium

E-mail address: Jean-Marie.Jacquet@unamur.be

URL: <http://staff.info.unamur.be/jmj>

Mieke Massink

Istituto di Scienza e Tecnologie dell'Informazione "A. Faedo", CNR, Pisa, Italy

E-mail address: Mieke.Massink@isti.cnr.it

URL: <http://www.isti.cnr.it/People/M.Massink>