

*Commenced Publication in 1973*

Founding and Former Series Editors:

Gerhard Goos, Juris Hartmanis, and Jan van Leeuwen

## Editorial Board

David Hutchison

*Lancaster University, Lancaster, UK*

Takeo Kanade

*Carnegie Mellon University, Pittsburgh, PA, USA*

Josef Kittler

*University of Surrey, Guildford, UK*

Jon M. Kleinberg

*Cornell University, Ithaca, NY, USA*

Friedemann Mattern

*ETH Zurich, Zürich, Switzerland*

John C. Mitchell

*Stanford University, Stanford, CA, USA*

Moni Naor

*Weizmann Institute of Science, Rehovot, Israel*

C. Pandu Rangan

*Indian Institute of Technology, Madras, India*

Bernhard Steffen

*TU Dortmund University, Dortmund, Germany*

Demetri Terzopoulos

*University of California, Los Angeles, CA, USA*

Doug Tygar

*University of California, Berkeley, CA, USA*

Gerhard Weikum

*Max Planck Institute for Informatics, Saarbrücken, Germany*

More information about this series at <http://www.springer.com/series/7411>

Márk Jelasity · Evangelia Kalyvianaki (Eds.)

# Distributed Applications and Interoperable Systems

16th IFIP WG 6.1 International Conference, DAIS 2016  
Held as Part of the 11th International Federated Conference  
on Distributed Computing Techniques, DisCoTec 2016  
Heraklion, Crete, Greece, June 6–9, 2016  
Proceedings

*Editors*  
Márk Jelasity  
University of Szeged  
Szeged  
Hungary

Evangelia Kalyvianaki  
City University London  
London  
UK

ISSN 0302-9743                      ISSN 1611-3349 (electronic)  
Lecture Notes in Computer Science  
ISBN 978-3-319-39576-0              ISBN 978-3-319-39577-7 (eBook)  
DOI 10.1007/978-3-319-39577-7

Library of Congress Control Number: 2016939991

LNCS Sublibrary: SL5 – Computer Communication Networks and Telecommunications

© IFIP International Federation for Information Processing 2016

This work is subject to copyright. All rights are reserved by the Publisher, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilms or in any other physical way, and transmission or information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed.

The use of general descriptive names, registered names, trademarks, service marks, etc. in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant protective laws and regulations and therefore free for general use.

The publisher, the authors and the editors are safe to assume that the advice and information in this book are believed to be true and accurate at the date of publication. Neither the publisher nor the authors or the editors give a warranty, express or implied, with respect to the material contained herein or for any errors or omissions that may have been made.

Printed on acid-free paper

This Springer imprint is published by Springer Nature  
The registered company is Springer International Publishing AG Switzerland

# Foreword

The 11th International Federated Conference on Distributed Computing Techniques (DisCoTec) took place at Aquila Atlantis Hotel in Heraklion, Greece, during June 6–9, 2016. It was organized by the Institute of Computer Science of the Foundation for Research and Technology Hellas and the University of Ioannina, Greece. The DisCoTec series is one of the major events sponsored by the International Federation for Information Processing (IFIP). It comprises three conferences:

- COORDINATION, the IFIP WG6.1 International Conference on Coordination Models and Languages
- DAIS, the IFIP WG6.1 International Conference on Distributed Applications and Interoperable Systems
- FORTE, the IFIP WG6.1 International Conference on Formal Techniques for Distributed Objects, Components and Systems

Together, these conferences cover a broad spectrum of distributed computing subjects, ranging from theoretical foundations and formal description techniques to systems research issues.

Each day of the federated event began with a plenary speaker nominated by one of the conferences. The three invited speakers were Tim Harris (Oracle Labs, UK), Catuscia Palamidessi (Inria, France), and Vijay Saraswat (IBM T.J. Watson Research Center, USA).

Associated with the federated event were also two satellite workshops, that took place on June 8–9, 2016:

- The 9th Workshop on Interaction and Concurrency Experience (ICE) with keynote lectures by Uwe Nestmann (Technische Universität Berlin, Germany) and Alexandra Silva (University College London, UK)
- The Final Public Workshop from the LeanBigData and CoherentPaaS projects

Sincere thanks go to the chairs and members of the Program and Steering Committees of the involved conferences and workshops for their highly appreciated efforts. Organizing DisCoTec 2016 was only possible thanks to the dedicated work of the Organizing Committee, including George Baryannis (Publicity Chair) and Vincenzo Gulisano (Workshops Chair), with excellent support from Nikos Antonopoulos and Alkis Polyraakis of PCO-Convin. Finally, many thanks go to IFIP WG6.1 for sponsoring this event, Springer *Lecture Notes in Computer Science* for their support and sponsorship, and to EasyChair for providing the refereeing infrastructure.

# Preface

This volume contains the proceedings of DAIS 2016, the 16th IFIP International Conference on Distributed Applications and Interoperable Systems, sponsored by the IFIP (International Federation for Information Processing) and organized by the IFIP Working Group 6.1.

DAIS was held during June 4–7, 2016, in Heraklion, Crete, Greece, as part of DisCoTec, the 11th International Federated Conference on Distributed Computing Techniques, together with FORTE (the 36th IFIP International Conference on Formal Techniques for Distributed Objects, Components and Systems) and COORDINATION (the 18th IFIP International Conference on Coordination Models and Languages).

There were 34 submissions for DAIS. Each submission was reviewed by at least three, and on average 3.8, Program Committee members. The committee decided to accept 13 full papers and three short papers, resulting in an acceptance rate of 38 % for full papers. Out of the 13 accepted full papers, five were chosen after a shepherding process, during which a committee member helped the authors implement the improvements required by the reviewers.

The accepted papers represent a compelling sample of the state of the art in the area of distributed applications and interoperable systems. Cloud computing and services received a large emphasis this year. The proceedings include contributions on resource management (optimizing energy consumption, adaptive optimization of Hadoop parameters, predicting resource usage, and an investigation of erasure coding libraries for redundancy) as well as techniques for cloud-based data processing and computing (shuffling methods, GPU access from virtual machines, and edge-cloud computing). In the area of decentralized systems, we have contributions on self-organizing DHTs, decentralized social networks, and content distribution systems. The area of Byzantine fault tolerance is also represented as well as the emerging area of complex event processing, where performance, load balancing, and language support are investigated.

The conference was made possible by the work and cooperation of many people working in several committees and organizations that are listed in these proceedings. In particular, we thank the Program Committee members for their commitment and thorough reviews and for their active participation in the discussion and shepherding phases, and all the external reviewers for their help in evaluating submissions. We would like to thank the publicity chair, Valerio Schiavoni, for his excellent work in publicizing DAIS 2016. We would also like to thank Tim Harris, our invited keynote speaker. Finally, we also thank the DisCoTec general chair, Kostas Magoutis, and the DAIS Steering Committee chair, Rui Oliveira, for their constant availability, support, and guidance.

April 2016

Márk Jelasity  
Evangelia Kalyvianaki

# Organization

## DisCoTec Steering Committee

Elie Najm (Chair)	Telecom-ParisTech, France
Rocco De Nicola	IMT Lucca, Italy
Kurt Geihs	University of Kassel, Germany
Farhad Arbab (Coordination)	CWI, The Netherlands
Rui Oliveira (DAIS)	University of Minho, Portugal
Jean-Bernard Stefani (FORTE)	Inria, France
Alain Girault	Inria, France
Uwe Nestmann	TU Berlin, Germany
Michele Loreti	University of Florence, Italy
Jim Dowling	KTH, Sweden
Marjan Sirjani	Reykjavik University, Iceland
Frank de Boer	CWI, The Netherlands
Lea Kutvonen	University of Helsinki, Finland
John Derrick	University of Sheffield, UK
Gianluigi Zavattaro	University of Bologna, Italy

## DAIS Steering Committee

Alysson Bessani	Universidade de Lisboa, Portugal
Sara Bouchenak	INSA Lyon, France
Jim Dowling	KTH Royal Institute of Technology, Sweden
Frank Eliassen	University of Oslo, Norway
Pascal Felber	Université de Neuchâtel, Switzerland
Karl Goeschka	Vienna University of Technology, Austria
Rüdiger Kapitza	Technical University of Braunschweig, Germany
Kostas Magoutis	FORTH-ICS and University of Ioannina, Greece
Rui Oliveira	Universidade do Minho, Portugal
Peter Pietzuch	Imperial College London, UK
Romain Rouvoy	University of Lille 1, France
François Taiani	Université de Rennes 1, France

## DAIS 2016 Program Committee

### Program Committee Chairs

Márk Jelasity	University of Szeged, Hungary
Evangelia Kalyvianaki	City University London, UK

**Publicity Chair**

Valerio Shiavonni                      Université de Neuchâtel, Switzerland

**Program Committee Members**

Luciana Arantes	Université Pierre et Marie Curie-Paris6, France
Carlos Baquero	HASLab, INESC TEC and University of Minho, Portugal
Sonia Ben Mokhtar	LIRIS CNRS, France
Alysson Bessani	Universidade de Lisboa, FCUL/LaSIGE, Portugal
Andrea Bondavalli	University of Florence, Italy
Sara Bouchenak	INSA Lyon, France
Jian-Nong Cao	Hong Kong Polytechnic University, Hong Kong, SAR China
Miguel Correia	IST/INESC-ID, Portugal
Paolo Costa	Microsoft Research Cambridge, UK
Wolfgang De Meuter	Vrije Universiteit Brussel, Belgium
Jim Dowling	KTH, Sweden
Frank Eliassen	University of Oslo, Norway
Ittay Eyal	Cornell University, USA
David Eyers	University of Otago, New Zeland
Pascal Felber	Université de Neuchâtel, Switzerland
Kurt Geihs	Universität Kassel, Germany
Karl M. Goeschka	FH Technikum Wien, Austria
Franz J. Hauck	University of Ulm, Germany
K.R. Jayaram	IBM Research, USA
Vana Kalogeraki	Athens University of Economics and Business, Greece
Rüdiger Kapitza	TU Braunschweig, Germany
Attila Kertesz	University of Szeged, Hungary
Benjamin Mandler	IBM, Israel
Evangelos Markatos	University of Crete, FORTH-ICS, Greece
Miguel Matos	HASLab, INESC TEC and University of Minho, Portugal
Rene Meier	Lucerne University of Applied Sciences and Arts, Switzerland
Alberto Montresor	University of Trento, Italy
Kiran-Kumar Muniswamy-Reddy	Harvard School of Engineering and Applied Sciences, USA
Marta Patino	Universidad Politecnica de Madrid, Spain
Peter Pietzuch	Imperial College London, UK
Hans P. Reiser	University of Passau, Germany
Altair Santin	Pontifical Catholic University of Paraná, Brazil
Spyros Voulgaris	VU University Amsterdam, The Netherlands



**Additional Reviewers**

Abreu, Vilmar  
Baraki, Harun  
Behl, Johannes  
Ceccarelli, Andrea  
Dar, Kashif Sana  
Gonçalves, Ricardo  
Huu Tran, Tam  
Kambona, Kennedy  
Köstler, Johannes  
Lollini, Paolo  
Maia, Francisco  
Marynowski, João Eugenio  
Mendes, Ricardo  
Obelheiro, Rafael  
Rakotondravony, Noelle

Renaux, Thierry  
Robu, Bogdan  
Rodriguez Avila, Humberto  
Saey, Mathijs  
Salaün, Gwen  
Sartakov, Vasily  
Shoker, Ali  
Swalens, Janwillem  
Taubmann, Benjamin  
Timo, Hönig  
Tomaras, Dimitris  
Vianello, Valerio  
Zacheilas, Nikos  
Zoppi, Tommaso

# Systems Challenges in Graph Analytics (DAIS 2016 Keynote)

Dr. Tim Harris

Oracle Labs, Cambridge, UK  
timothy.l.harris@oracle.com

**Abstract.** Graphs are at the core of many data processing problems, whether that is searching through billions of records for suspicious interactions, ranking the importance of web pages based on their connectivity, or identifying possible “missing” friends on a social network. This talk will discuss the challenges in building large, scalable, in-memory graph analytics systems. Many of these challenges come from the way that graph algorithms behave differently based on the structure of the input graph: a planar road network graph can produce a significantly different load on the machine’s memory system from a low-diameter social network graph. It can be necessary to select particular algorithms for these different cases, and to make contrasting decisions over how the machine’s resources are allocated. Finally, we face challenges simply from the scale at which we operate: making efficient use of the hardware in new SPARC machines with over 4000 threads.

**Speaker:** Tim Harris leads the Oracle Labs group in Cambridge, UK. His research interests span multiple layers of the stack, including parallel programming, VMM/OS/runtime-system interaction, and opportunities for specialized architecture support for particular workloads. He has also worked on the implementation of software transactional memory for multi-core computers, and the design of programming language features based on it. Tim has a BA and PhD in computer science from Cambridge University Computer Laboratory. He was on the faculty at the Computer Laboratory from 2000–2004 where he led the department’s research on concurrent data structures and contributed to the Xen virtual machine monitor project. He was at Microsoft Research from 2004, and then joined Oracle Labs to found the Cambridge office in 2012.

# Contents

Enhanced Energy Efficiency with the Actor Model on Heterogeneous Architectures . . . . .	1
<i>Yaroslav Hayduk, Anita Sobe, and Pascal Felber</i>	
Evaluating the Cost and Robustness of Self-organizing Distributed Hash Tables . . . . .	16
<i>Sveta Krasikova, Raziq C. Gómez, Heverson B. Ribeiro, Etienne Rivière, and Valerio Schiavoni</i>	
Mignon: A Fast Decentralized Content Consumption Estimation in Large-Scale Distributed Systems . . . . .	32
<i>Stéphane Delbruel, Davide Frey, and François Taïani</i>	
Privacy-Preserving Data Allocation in Decentralized Online Social Networks . . . . .	47
<i>Andrea De Salve, Paolo Mori, Laura Ricci, Raed Al-Aaridhi, and Kalman Graffi</i>	
An RDMA Middleware for Asynchronous Multi-stage Shuffling in Analytical Processing. . . . .	61
<i>Rui C. Gonçalves, José Pereira, and Ricardo Jiménez-Peris</i>	
Holistic Shuffler for the Parallel Processing of SQL Window Functions. . . . .	75
<i>Fábio Coelho, José Pereira, Ricardo Vilaça, and Rui Oliveira</i>	
Providing CUDA Acceleration to KVM Virtual Machines in InfiniBand Clusters with rCUDA . . . . .	82
<i>Ferran Pérez, Carlos Reaño, and Federico Silla</i>	
Benchmarking Wireless Protocols for Feasibility in Supporting Crowdsourced Mobile Computing . . . . .	96
<i>João Rodrigues, Joaquim Silva, Rolando Martins, Luís Lopes, Utsav Drolia, Priya Narasimhan, and Fernando Silva</i>	
BFT-Dep: Automatic Deployment of Byzantine Fault-Tolerant Services in PaaS Cloud . . . . .	109
<i>Bijun Li and Rüdiger Kapitza</i>	
<i>BFT-Bench</i> : Towards a Practical Evaluation of Robustness and Effectiveness of BFT Protocols . . . . .	115
<i>Divya Gupta, Lucas Perronne, and Sara Bouchenak</i>	

Self-Balancing Job Parallelism and Throughput in Hadoop. . . . . 129  
*Bo Zhang, Filip Křikava, Romain Rouvoy, and Lionel Seinturier*

Resource Usage Prediction in Distributed Key-Value Datastores . . . . . 144  
*Francisco Cruz, Francisco Maia, Miguel Matos, Rui Oliveira, João Paulo, José Pereira, and Ricardo Vilaça*

A Performance Evaluation of Erasure Coding Libraries for Cloud-Based Data Stores: (Practical Experience Report) . . . . . 160  
*Dorian Burihabwa, Pascal Felber, Hugues Mercier, and Valerio Schiavoni*

Dynamic Load Balancing Techniques for Distributed Complex Event Processing Systems . . . . . 174  
*Nikos Zacheilas, Nikolas Zygouras, Nikolaos Panagiotou, Vana Kalogeraki, and Dimitrios Gunopoulos*

PAN – Distributed Real-Time Complex Event Detection in Multiple Data Streams . . . . . 189  
*Lukas Probst, Ivan Giangreco, and Heiko Schuldt*

Bringing Complex Event Processing into Multitree Modelling of Sensors. . . . . 196  
*Alexandre Garnier, Jean-Marc Menaud, and Nicolas Montavont*

**Author Index** . . . . . 211