

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, Zurich, 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/7408>

Philipp Reinecke · Antinisca Di Marco (Eds.)

Computer Performance Engineering

14th European Workshop, EPEW 2017
Berlin, Germany, September 7–8, 2017
Proceedings

Editors
Philipp Reinecke
Bristol
UK

Antinisca Di Marco
University of L'Aquila
L'Aquila
Italy

ISSN 0302-9743 ISSN 1611-3349 (electronic)
Lecture Notes in Computer Science
ISBN 978-3-319-66582-5 ISBN 978-3-319-66583-2 (eBook)
DOI 10.1007/978-3-319-66583-2

Library of Congress Control Number: 2017949517

LNCS Sublibrary: SL2 – Programming and Software Engineering

© Springer International Publishing AG 2017

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. The publisher remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Printed on acid-free paper

This Springer imprint is published by Springer Nature
The registered company is Springer International Publishing AG
The registered company address is: Gewerbestrasse 11, 6330 Cham, Switzerland

Preface

This volume of LNCS contains the proceedings of the 14th European Performance Engineering Workshop, held in Berlin, Germany, September 7–8, 2017. EPEW was part of the week-long umbrella conference QONFEST, which co-located QEST, CONCUR, FORMATS, and EPEW, along with several workshops. This gave researchers the opportunity to explore and engage with a broad range of topics and colleagues across the space of performance, dependability, and security modelling, verification, evaluation, and engineering. We wish to express our gratitude for the support QONFEST received from the Freie Universität Berlin, the Technische Universität Berlin, the Ernst-Reuter-Gesellschaft, the DFG, and the Max-Planck-Gesellschaft.

The goal of the annual EPEW workshop series is to gather academic and industrial researchers working on all aspects of performance engineering. The papers presented at the workshop reflect the diversity of modern performance engineering, with topics ranging from the analysis of hybrid Petri nets and Markov decision processes, even under uncertainty; to performance, security and energy analysis of computer systems and networks; to machine-learning techniques for predictive analysis and testing. The domains of the application studies are diverse and at the cutting edge of current developments, ranging from cloud computing environments to cyber-physical systems and to communication protocols.

EPEW 2017 received submissions from 14 countries all over the world. There were 30 submissions. Each paper was peer reviewed by an average of four reviewers from the Program Committee (PC) on the basis of its relevance, novelty, and technical quality. After the collection of reviews, the PC members discussed the quality of the submissions for one week before getting the final decision. Based on the reviews and discussions, 18 high-quality contributions were selected for publication in the proceedings and presentation at the workshop.

This year, we were honored to have two keynote speakers: Prof. William Knottenbelt, from Imperial College London (UK), who works in applied quantitative analysis; and Antonino Sabetta, a senior researcher at the Security Research department of SAP Research (Sophie Antipolis, France), who works in the analysis and management of vulnerabilities of open-source components when embedded in large-scale enterprise applications.

We thank our keynote speakers, as well as all PC members and external reviewers for their terrific work in the review process. We also express our thanks to the Organizing Committee, especially to the two General Chairs, Uwe Nestmann (TU Berlin) and Katinka Wolter (FU Berlin) for their continuous and valuable help, the EasyChair team for their conference system, and Springer for their continued editorial support.

Above all, we would like to thank the authors of the papers for their contribution to this volume. We are sure that these contributions will be as useful and inspiring to the readers as they were to us.

September 2017

Philipp Reinecke
Antinisca Di Marco

Organization

EPEW Program Chairs

Antinisca Di Marco	University of L'Aquila, Italy
Philipp Reinecke	Bristol, UK

QONFEST General Chairs

Katinka Wolter	Freie Universität Berlin, Germany
Uwe Nestmann	Technische Universität Berlin, Germany

EPEW Program Committee

Davide Arcelli	Università de L'Aquila, Italy
Rena Bakhshi	Netherlands eScience Center, The Netherlands
Simonetta Balsamo	Università Ca' Foscari di Venezia, Italy
Marta Beltran	Universidad Rey Juan Carlos, Spain
Marco Bernardo	University of Urbino, Italy
Ana Busic	Inria and ENS, France
Laura Carnevali	University of Florence, Italy
Giuliano Casale	Imperial College London, UK
Dieter Fiems	Ghent University, Belgium
Jean-Michel Fourneau	Université de Versailles St Quentin, France
Stephen Gilmore	University of Edinburgh, UK
Boudewijn Haverkort	University of Twente, The Netherlands
András Horváth	University of Turin, Italy
Gábor Horváth	Budapest University of Technology and Economics, Hungary
Alain Jean-Marie	CNRS University of Montpellier, France
William Knottenbelt	Imperial College London, UK
Samuel Kounev	University of Würzburg, Germany
Vasilis Koutras	University of the Aegean, Greece
Lasse Leskelä	Aalto University, Finland
Catalina M. Lladó	Universitat Illes Balears, Spain
Andrea Marin	University Ca' Foscari Venice, Italy
Raffaella Mirandola	Politecnico di Milano, Italy
Marco Paolieri	University of Southern California, USA
Roberto Pietrantuono	University of Naples Federico II, Italy
Agapios Platis	University of the Aegean, Greece
Anne Remke	WWU Münster, Germany
Markus Siegle	Universität der Bundeswehr, München, Germany

Miklos Telek	Budapest University of Technology and Economics, Hungary
Nigel Thomas	Newcastle University, UK
Catia Trubiani	Gran Sasso Science Institute, Italy
Petr Tuma	Charles University, Czech Republic
Aad Van Moorsel	Newcastle University, UK
Maaïke Verloop	CNRS Toulouse, France
Joris Walraevens	Ghent University, Belgium
Qiushi Wang	Nanyang Technological University, China
Huaming Wu	Tianjin University, China
Armin Zimmermann	Technische Universität Ilmenau, Germany

EPEW Additional Reviewers

Alnafessah, Ahmad	Iffländer, Lukas
Baltas, Ioannis	Meszaros, Andras
Herbst, Nikolas	Pekergin N., Nihal
Horvath, Illes	Pilch, Carina
Hüls, Jannik	von Kistowski, Jóakim

Abstracts of Invited Talks

Cryptocurrency and Blockchain Technology: Challenges and Opportunities

William J. Knottenbelt

Imperial College Centre for Cryptocurrency Research and Engineering,
Imperial College London, London, UK
wjkn@imperial.ac.uk

The meteoric rise of blockchain-enabled cryptocurrencies, and Bitcoin [2] and Ethereum [1] in particular, has received global attention, not least from governments, entrepreneurs and researchers. Cryptocurrencies, of which there are now more than 800¹, provide an attractive alternative to traditional fiat currencies via a distributed, trustless and self-governing framework which not only enables low-friction financial transactions around the globe but also preserves the freedom and privacy of spending inherent in cash transactions.

Cryptocurrency and blockchain technology brings with it a host of new challenges from the quantitative modelling perspective. Indeed, a range of issues including performance, security, energy use, incentives and scalability are poorly understood, as are the inherent trade offs between them, despite these being critical barriers to mass adoption. What analyses are carried out often do not take into account problems posed by the lack of diversity that emerges from a natural tendency towards dominant concentrations of computational and other power. These can arise from something as simple as the majority of network participants flocking to deploy the most energy-efficient cryptocurrency mining hardware. Indeed it is estimated that up to 70% of the computational power assuring the integrity of the Bitcoin network is provided by a single model of a hardware device. This device was recently found to have a backdoor that could be used by the manufacturer to shut the device down².

This talk will cover some of the challenges and opportunities posed in this context, with a special emphasis on the performance evaluation and quantitative modelling perspectives. It turns out that classical performance evaluation techniques, especially Markovian analysis and queueing theory, are readily applicable to the study of cryptocurrencies and blockchains. Further, a judicious combination of analytical modelling, simulation and benchmarking techniques can be effectively applied to yield insights. Building on [3], we will illustrate this in the context of a study of a queue-based Ethereum mining pool [4] whose superficially fair reward scheme turns out not only to penalise more powerful miners, but also to incentivise a number of attacks which can

W.J. Knottenbelt—The content of the talk is the result of joint work with A. Zamyatin, K. Wolter, C. Mulligan, P. Harrison, S. Werner and I. Stewart, amongst others.

¹ Source: <http://coinmarketcap.com>. Accessed 5 July 2017.

² Source: <http://antbleed.com>. Accessed 5 July 2017.

increase rewards, including the donation of mining power to other participants in certain circumstances. Examples of such attacks observed in the real world will be presented.

The talk will conclude by outlining student-led spinout activity and ongoing directions of research in the Imperial College Centre for Cryptocurrency Research and Engineering. The former includes Gradbase³, a qualification verification startup, Aventus⁴, a blockchain-based ticketing company and Kotiva Technologies⁵, who are seeking to use blockchain technology to increase the integrity of supply chains. The latter includes work being supported by industrial partners such as Blockchain.com and Outlier Ventures, as well as grants sponsored by government-related bodies such as Innovate UK.

Biography

William Knottenbelt is Professor of Applied Quantitative Analysis and Director of Industrial Liaison in the Department of Computing at Imperial College London, where he became a Lecturer in 2000. He is a founder of the Imperial Blockchain Forum, is co-Director of the Centre for Cryptocurrency Research and Engineering and is Director of the Data Economy Lab in Imperial's Data Science Institute. He serves on the editorial board of the cryptocurrency/blockchain journal *Ledger*, is an editor of *Performance Evaluation Journal*, and has served as general or program chair of numerous conferences and workshops related to quantitative modelling and analysis. A keen supporter of student-led innovation, he is the Innovation Fellow for the Department of Computing and serves on the Entrepreneur First Science Partners panel. In June 2017, he presented his Inaugural Lecture entitled “Memoirs of the Memoryless: A Markovian Meander from Disk Drives to Digital Money”, which is available online⁶.

References

1. Vitalik B.: Ethereum: A next-generation smart contract and decentralized application platform (2014). <https://github.com/ethereum/wiki/wiki/White-Paper>. Accessed 18 June 2017
2. Satoshi, N.: Bitcoin: A peer-to-peer electronic cash system (2008). <https://bitcoin.org/bitcoin.pdf>. Accessed 18 June 2017
3. Meni, R.: Analysis of bitcoin pooled mining reward systems. arXiv preprint [arXiv:1112.4980](https://arxiv.org/abs/1112.4980) (2011)
4. Zamyatin, A., Wolter, K., Werner, S., Mulligan, C.E.A., Harrison, P.G., Knottenbelt, W.J.: Swimming with fishes and sharks: beneath the surface of queue-based Ethereum mining pools. In: Proceedings of the 25th IEEE International Symposium on Modeling, Analysis and Simulation of Computer and Telecommunications Systems (MASCOTS 2017), September 2017

³ See <http://gradba.se>.

⁴ See <http://aventus.io>.

⁵ See <http://kotiva.tech>.

⁶ See <https://youtu.be/TTQOwyXXKHw>.

Open-Source Libraries Included in Enterprise Applications: Workhorses or Trojan Horses?

Antonino Sabetta

SAP Labs, France

`antonino.sabetta@sap.com`

The adoption of open-source software (OSS) components in the software industry has grown at a spectacular pace over the last decade. By some estimates [3], the average commercial software product contains 100 distinct open source components whose code weights as much as 35% of the overall application size¹.

At the same time, new vulnerabilities affecting open-source software (OSS) are reported on a daily basis, sometimes hitting the headlines of mainstream media (as it happened, for example, with Heartbleed² and ShellShock³).

The relevance of this problem has been well documented by now [1, 2] and establishing effective vulnerability management practices for OSS is broadly understood as a priority in the software industry.

Despite the deceiving simplicity of the existing solutions (especially of the most obvious: *updating to a recent, non-vulnerable version*), OSS libraries with known vulnerabilities are found to be used for quite some time after a fixed version has been released [3].

As a matter of fact, updating a library to a more recent release is quite straightforward at *development* time. However, things become considerably more difficult when vulnerable OSS libraries are part of large enterprise systems that are already in *operation* and serve business-critical functions. Any change (including corrections) may cause costly system downtime and comes with the risk that new unforeseen issues could arise.

For this reason, it is extremely important to properly assess whether an application requires an urgent patch to update an OSS dependency, or whether the update could be scheduled for the next regular release cycle. Just the presence of a vulnerable dependency is not enough to justify a urgent update, with its high costs and even higher risks. The real question is whether a given vulnerability is indeed *exploitable* given the particular way the dependency is used.

Unfortunately, assessing the exploitability and the potential impact of a vulnerability found in a dependency is difficult, expensive, and error-prone. Vulnerabilities are

A. Sabetta—The content of the talk is the result of joint work with Serena E. Ponta and Henrik Plate, SAP Labs France.

¹ The same study reports that for applications developed for internal use, the proportion is as high as 75%.

² <http://heartbleed.com/>.

³ <https://shellshocker.net/>.

documented in advisories that consist of short, high-level, textual descriptions expressed in natural language, whereas a reliable assessment demands much lower-level, detailed, technical information.

The consequences of a wrong assessment can be expensive. If an exploitable vulnerability is not identified as such, users remain exposed to attackers. When, on the contrary, a correction is produced for a non-exploitable vulnerability, the effort of developing, testing, and deploying the correction is spent in vain.

This talk summarizes the key elements of our research on how to make the assessment of OSS vulnerabilities more efficient and systematic [4]. Our approach aims to automatically produce concrete evidence (when it can be found) supporting the case for urgent patching. Such evidence consists of concrete call sequences (traces) that start from application methods and reach the vulnerable methods of a dependency. We complement *potential* traces obtained through static analysis with *actual* observations of runtime executions collected through dynamic instrumentation. Our approach relies on the availability of detailed (code-level) vulnerability information, which we extract by mining software repositories with the support of machine learning. The initial research prototype that we implemented to validate our approach evolved over time into an enterprise-grade OSS vulnerability analysis toolkit (internally known as *Vulas*), which is used regularly in hundreds of development (and maintenance) projects across our company.

Biography

Antonino Sabetta is a senior researcher at the Security Research department of SAP. The main focus of Antonino's recent work is the analysis and management of vulnerabilities of open-source components embedded in large-scale enterprise applications. In particular, Antonino is interested in the application of machine-learning to the mining of open-source software repositories and the automation of the vulnerability management workflow.

Before moving to SAP in 2010, Antonino was a researcher at CNR, Pisa, Italy. He earned his PhD in Computer Science and Automation Engineering from the University of Rome Tor Vergata, Italy in 2007. From the same university he had received in 2003 his “Laurea *cum Laude*” degree in Computer Engineering.

References

1. Arce, I., et al.: Avoiding the top-10 software security design flaws. Technical report, IEEE Center for Secure Design. IEEE Computer Society (2014)
2. OWASP Foundation. OWASP Top 10 – 2013 (2013). https://www.owasp.org/index.php/Top_10_2013-Top_10
3. Pittenger, M.: Open source security analysis: The state of open source security in commercial applications. Technical report, Black Duck Software (2016)
4. Plate, H., Ponta, S.E., Sabetta, A.: Impact assessment for vulnerabilities in open-source software libraries. In: Proceedings of the IEEE International Conference on Software Maintenance and Evolution (ICSME) (2015)

Contents

Advances in Markov Models

Analysis of Markov Decision Processes Under Parameter Uncertainty	3
<i>Peter Buchholz, Iryna Dohndorf, and Dimitri Scheftelowitsch</i>	
Bounded Aggregation for Continuous Time Markov Decision Processes	19
<i>Peter Buchholz, Iryna Dohndorf, Alexander Frank, and Dimitri Scheftelowitsch</i>	
Interactive Markovian Equivalence	33
<i>Arpit Sharma</i>	

Advances in Quantitative Analysis

Delay Analysis of Resequencing Buffer in Markov Environment with HOQ-FIFO-LIFO Policy	53
<i>Rostislav Razumchik and Miklós Telek</i>	
Analysis of Timed Properties Using the Jump-Diffusion Approximation. . . .	69
<i>Paolo Ballarini, Marco Beccuti, Enrico Bibbona, Andras Horvath, Roberta Sirovich, and Jeremy Sproston</i>	
Stability Analysis of a Multiclass Retrial System with Coupled Orbit Queues	85
<i>Evsey Morozov and Ioannis Dimitriou</i>	

Model Checking

Model Checking the STL Time-Bounded Until on Hybrid Petri Nets Using Nef Polyhedra	101
<i>Adrian Godde and Anne Remke</i>	
A New Approach to Predicting Reliable Project Runtimes via Probabilistic Model Checking	117
<i>Ulrich Vogl and Markus Siegle</i>	

Cyber-Physical Systems

Learning-Based Testing of Cyber-Physical Systems-of-Systems: A Platooning Study	135
<i>Karl Meinke</i>	

An Inspection-Based Compositional Approach to the Quantitative Evaluation of Assembly Lines	152
<i>Marco Biagi, Laura Carnevali, Tommaso Papini, Kumiko Tadano, and Enrico Vicario</i>	

Performance, Energy and Security

Machine Learning Models for Predicting Timely Virtual Machine Live Migration	169
<i>Osama Alrajeh, Matthew Forshaw, and Nigel Thomas</i>	
Model-Based Simulation in Möbius: An Efficient Approach Targeting Loosely Interconnected Components	184
<i>Giulio Masetti, Silvano Chiaradonna, and Felicita Di Giandomenico</i>	
Analysis of Performance and Energy Consumption in the Cloud	199
<i>Mehdi Kandi, Farah Aït-Salaht, Hind Castel-Taleb, and Emmanuel Hyon</i>	
Deriving Power Models for Architecture-Level Energy Efficiency Analyses	214
<i>Christian Stier, Dominik Werle, and Anne Koziolk</i>	
ADaCS: A Tool for Analysing Data Collection Strategies	230
<i>John C. Mace, Nipun Thekkummal, Charles Morisset, and Aad Van Moorsel</i>	

Case Studies

Improving ZooKeeper Atomic Broadcast Performance by Coin Tossing	249
<i>Ibrahim EL-Sanosi and Paul Ezhilchelvan</i>	
Modelling and Analysis of Commit Protocols with PEPA	266
<i>Said Naser Said Kamil and Nigel Thomas</i>	
Stochastic Models for Solar Power	282
<i>Dimitra Politaki and Sara Alouf</i>	
Author Index	299