Report from Dagstuhl Seminar 13341

Verifiably Secure Process-Aware Information Systems

Edited by

Rafael Accorsi[1], Jason Crampton[2], Michael Huth[3], and Stefanie Rinderle-Ma[4]

1 **Universität Freiburg, DE,** rafael.accorsi@iig.uni-freiburg.de
2 **RHUL – London, GB,** jason.crampton@rhul.ac.uk
3 **Imperial College London, GB,** M.Huth@imperial.ac.uk
4 **Universität Wien, AT**

Abstract

From August 18–23, 2013, the Dagstuhl Seminar "Verifiably Secure Process-aware Information Systems" was held in Schloss Dagstuhl – Leibniz Center for Informatics. During this seminar, participants presented their current research and discussed open problems in the arising field of securing information systems driven by processes. The executive summary and abstracts of the talks given during the seminar are put together in this paper.

Seminar 18.–23. August, 2013 – www.dagstuhl.de/13341
1998 ACM Subject Classification C.2.2 Protocol Verification, C.4 Reliability, availability, and serviceability, D.2.0 Protection Mechanisms, D.2.4 Software/Program Verification, D.2.11 Software Architectures, D.4.5 Reliability, D.4.6 Security and Protection, F.3.2 Process models, H.4.1 Workflow management, I.2.8 Scheduling, J.1 Administrative Data Processing, K.5.2 Regulation
Keywords and phrases Business Processes, Information Security, Compliance, Risk-Aware Processes, Service Compositions
Digital Object Identifier 10.4230/DagRep.3.8.73

1 Executive Summary

Rafael Accorsi
Jason Crampton
Michael Huth
Stefanie Rinderle-Ma

Business processes play a major role in many commercial software systems and are of considerable interest to the research communities in Software Engineering, and Information and System Security. A process-aware information system provides support for the specification, execution, monitoring and auditing of intra- as well as cross-organizational business processes.

Designing and enacting secure business processes is as tricky as "Programming Satan's Computer", as Ross Anderson and Roger Needham observed in a paper with that title. Recent fraud disasters show how subtle secure process engineering and control can be. The Swiss bank UBS suffered from a rogue trader scandal in 2011, which led to a loss of a then-estimated US$2 billion, was possible because the risk of trades could be disguised by using "forward-settling" Exchange-traded Funds (ETF) cash positions. Specifically, processes that implemented ETF transactions in Europe do not issue confirmations until after settlement

Verifiably Secure Process-Aware Information Systems, *Dagstuhl Reports*, Vol. 3, Issue 8, pp. 73–86
Editors: Rafael Accorsi, Jason Crampton, Michael Huth, and Stefanie Rinderle-Ma

Dagstuhl Reports
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

has taken place. The exploitation of this process allows a party in a transaction to receive payment for a trade before the transaction has been confirmed. While the cash proceeds in this scheme cannot be simply retrieved, the seller may still show the cash on their books and possibly use it in further transactions. Eventually, the mechanics of this attack allowed for a carrousel of transactions, thereby creating an ever growing snowball. Similar analyses, usually based upon insider threats, can also be made for fraud cases such as the well-documented Société Générale case, but also for the WorldCom and Parmalat cases.

Addressing these problems requires, on the one hand, strong security and compliance guarantees. On the other hand, these guarantees must be substantiated by formal methods ensuring a verifiably secure business process enactment. It should be noted that these concerns are not confined to the financial service sector or to insider threats. For example, the planned unification of Eurpean data protection law into a sole Data Protection Regulation law is likely to change the statuatory duties of the private sector. Under this plan, companies will be legally required to report any breaches of this regulation and may be liable to penalties in the range of 2–5% of their global annual turnover. European industries seem to be ill-prepared to ensure that their information systems and processes will comply with the security requirements of that upcoming regulation, and the threat of substantial fines means that there is an urgent need to create more resilient systems and processes, which calls for more research within the thematic scope of this seminar.

At the interface of security requirements, business needs, and compliance methodologies we can ask many practically relevant research questions, and their answers are bound to have significant impact in academia and industry alike. Relatively little work has been done, however, on adapting or creating new formal methods with which one can check that processes are compliant with rules, preserve demanded privacy constraints, *and* enforce desired security policies at the same time.

One main purpose of the seminar was to present the state of the art in research within the three communities of Security, Verification, and Process-Aware Information Systems to all three communities in an accessible manner and with a view of identifying important research topics at the intersections of these communities. In addition, that exercise was also meant to explore what strategic activities could help in promoting research at the junction of these communities. This agenda was persued through a mix of keynotes, technical presentations, break-out groups under the WorldCafe method, sessions with free-style discussions, and tool demonstrations.

We now highlight some of the key questions and findings can emerged during that week of work – we refer to the online archive of presentation slides, papers, and abstracts for more detailed discussions and findings. Three action items that seemed of particular importance to the participants were:

1. The need for a classification of security properties that are revelant for process-aware information systems, and an understanding of what formal methods might be able to analyse such properties.
2. The need for a set of concrete examples of business processes that are annotated with security considerations or constraints. These might be examples from the real world that have been sufficiently sanatized and anonymized.
3. The need for a review/survey article on the state of the art in formal methods, written for non-experts and ideally for an audience that deals with security, privacy, audit or business processes.

It was also asked what makes formal methods and tools “practical” in this problem space; their was concern about the scalability of these methods, but also about the considerable

effort it would take to transfer foundational tools to real application domains – were such somewhat routine but important transfer may not be supported by standing funding models. Concerns were also voiced about the current research in security and privacy, which tends to ignore recent innovations in process composition, such as choreographies.

Another point of considerable interest made concerned the organization of research in this problem space. At the moment, researchers work on problems within their areas and when they begin to collaborate with people from another area this is then more of a bottom-up process where techniques and tools across areas are combined to see what problems one could now solve. It was remarked that it may often be more effective to take a top-down approach in which key problems of the inter-area domain are first formulated and then researchers from the areas get together and try to come up with solutions that draw from their own tool boxes but that may also invent new tools for the problem at hand.

There was also a lot of discussion about what is so distinctive about *process-aware* information systems, and whether these differences to conventional information systems offer perhaps also opportunities. For example, it was discussed whether there is value in validating such systems at a high level of abstraction without considering how such processes get implemented in IT infrastructures and abstraction layers. The participants had mixed views on such merits but it was felt that validation at that level would be easier to realize and that the identification of weaknesses or vulnerabilities at that layer would no doubt be of value.

Another problem mentioned was the need to support legacy systems, and that this need would not go away. Faced with this, it appears that formal validation techniques will have to be able to reason about composed systems in which some parts only have a somewhat well defined interface, but whose internal behavior cannot be guaranteed or predicted to a good degree.

Finally, it was also noted that some of the research problems that suggest themselves to the specialists may not be issues in the field. For example, we may want trusted system composition across organizations but there may not be the need to formally validate such trust since contractual or other legal mechanisms may be in place that incentivize parties to honor that trust, and that give parties a means of seeking damages in case that trust has been violated. On the other hand, such legal mechanisms may not be adequate in the upcoming Internet of Things were 2-party, end-to-end composition will be the exception and not the norm.

2 Table of Contents

3 Overview of Talks

3.1 Deriving RBAC models from process models and logs

Anne Baumgraß (Hasso-Plattner-Institut – Potsdam, DE)

In process-aware information systems (PAIS) permissions need to be tailored, both to allow legitimate users to perform their specific tasks and to avoid fraud and abuse. Role-based access control (RBAC) is a de facto standard to model and specify access control policies. Although RBAC provides a number of advantages for the management of access control policies, the definition of a specific RBAC model is a complex and time-consuming task for security experts. In addition, a constant evolution of business processes as well as its corresponding permissions result in structures that have changed over time and do not necessarily represent a tailored (desired) RBAC configuration of an organization. With the aim to obtain the current RBAC configurations and further support the definition of desired RBAC models,this talk presents an approach that is able to derive candidate RBAC models from process models managed in PAIS as well as log files recorded by PAIS.

3.2 Better enforce than verify! How to ensure compliance of business processes at runtime

Nataliia Bielova (INRIA Rennes – Bretagne Atlantique, FR)

It is well known that compliance is important for business processes. This talk will advocate the use of runtime enforcement techniques to ensure compliance of business processes. Differently from verification, that is concerned with detection of non-compliant executions of a process, runtime enforcement is aimed at correcting business process while it runs to guarantee the desired behaviour. First, we present the case study – a business process of drug dispensation from the Hospital San Raffaele, and an expected behaviour of this process. Second, we show how different runtime enforcement mechanisms can ensure compliance at runtime. Third, we explain and compare the theoretical models of such mechanisms and formal guarantees they provide.

3.3 Security in the Context of Business Processes: Thoughts from a System Vendor's Perspective

Achim D. Brucker (SAP Research – Karlsruhe, DE)

Enterprise systems in general and process aware systems in particular are storing and processing the most critical assets of a company. To protect these assets, such systems need to implement a multitude of security properties. Moreover, such systems need often to

comply to various compliance regulations. In this keynote, we present process-level security requirements as well as discuss the gap between the ideal world of process-aware information systems and the real world. We conclude our presentation by discussing several research challenges in the area of verifiable secure process aware information systems.

3.4 On the Parameterized Complexity of the Workflow Satisfiability Problem

Jason Crampton (Royal Holloway University of London, UK)

A workflow specification defines a set of steps and the order in which those steps must be executed. Security requirements may impose constraints on which groups of users are permitted to perform subsets of those steps. A workflow specification is said to be satisfiable if there exists an assignment of users to workflow steps that satisfies all the constraints. An algorithm for determining whether such an assignment exists is important, both as a static analysis tool for workflow specifications, and for the construction of run-time reference monitors for workflow management systems. Finding such an assignment is a hard problem in general, but recent work using the theory of parameterized complexity suggests that efficient algorithms exist under reasonable assumptions about workflow specifications. We improve the complexity bounds for the workflow satisfiability problem. We also generalize and extend the types of constraints that may be defined in a workflow specification and prove that the satisfiability problem remains fixed-parameter tractable for such constraints.

3.5 The Dos and Don'ts of Business Process Compliance

Guido Governatori (NICTA, AU)

The aim of business process compliance is to ensure that the processes of a business satisfy the relevant legal requirements In the first part of this contribution we propose the semantics of legal requirements in terms of the possible ways in which a process can be executed. In particular we provide a classification of the various types of obligations (and related notions) and what they mean in terms of a business process. In the second part we examine how various logics and logical formalisms address the legal requirement, and their suitability to represent legal reasoning with a particular focus on regulatory compliance for business processes.

3.6 Policy-Based Numerical Aggregation of Trust Evidence

Michael Huth (Imperial College, UK)

The decision to trust and so to allow an action may be informed by many, heterogeneous forms of evidence. Often, such evidence is quantitative rather than qualitative, e.g. the age of some software, the location of a device, the amount of a financial transaction, etc. We present a language Peal, for "pluggable evidence aggregation language," in which the aggregation of such evidence can be specified so that aggregated values can be compared with thresholds. These comparisons are conditions that may be plugged into existing policy languages for access control (e.g., XACML) to enrich them with considerations of risk, trust, cost, etc.. We show how such conditions can generate logical formulas that "compile away" any reference to numerical values yet precisely capture said conditions. Then we give a demonstration of a tool we developed, in which we can subject such conditions to important verification tasks, e.g. vacuity checking and sensitivity analysis of thresholds, by analyzing the aforementioned compilations with the SMT solver Z3.

3.7 Formal Methods with Industry?

Fuyuki Ishikawa (National Institute of Informatics – Tokyo, JP)

In accordance with the active discussions during this seminar, this talk presents our activities with the Japanese industry on formal methods (FM). FM has recently attracted strong interests in the Japanese industry. This talk first presents various kinds of guidelines in active FM promotion to "non-FM people". Our educational program Top SE, an educational course for the industry, also has included intensive support for FM. The program has kept increasing attendees (now 40- per year) from the industry for its 1-year educational course. This talk presents experiences in the program, how FM are taught with real(-like) problems and what the top-level engineers tackle after they learn basics of FM (in the process-aware systems for example).

3.8 Policy Auditing over Incomplete Logs: Theory, Implementation and Applications

Limin Jia (Carnegie Mellon University, US)

We present the design, implementation and evaluation of an algorithm that checks audit logs for compliance with privacy and security policies. The algorithm, which we name reduce, addresses two fundamental challenges in compliance checking that arise in practice. First, reduce operates on policies that quantify over data with possibly infinite domains. Second, audit logs are inherently incomplete (they may not contain sufficient information to determine

whether a policy is violated or not), reduce proceeds iteratively: in each iteration, it provably checks as much of the policy as possible over the current log and outputs a residual policy that can only be checked when the log is extended with additional information. We implement reduce and use it to check simulated audit logs for compliance with the HIPAA Privacy Rule. Our experimental results demonstrate that the algorithm is fast enough to be used in practice.

3.9 Checking System Compliance by Slicing and Monitoring Logs

Felix Klaedtke (ETH Zürich, CH)

Many kinds of digitally stored data should only be used in restricted ways. The intended usage may be stipulated by government regulations, corporate privacy policies, preferences of the data owners, etc. Such policies cover not only who may access which data, but also how the data may or must not be used after access. An example of such a usage restriction is that "collected data must be deleted after 30 days and not accessed or forwarded to third parties." A promising approach to detect violations with respect to such policies in IT systems utilizes monitoring techniques.

In this talk, I will present a scalable solution for compliance checking based on monitoring the agents' actions, where policies are specified in an expressive temporal logic and the system actions are logged. In particular, our solution utilizes the MapReduce framework to parallelize the process of monitoring the logged actions. I will sketch the theoretical framework underpinning our solution, i.e., the sound and complete reorganization of the logged actions into parts, which we call slices, and that can be analyzed independently of each other, and orthogonal methods for generating such slices, and means for combining these methods. Finally, I will report on a real-world case study, which demonstrates the feasibility and the scalability of our monitoring solution.

3.10 Information Flow Security for Business Processes – just one click away

Andreas Lehmann (Vattenfall Europe Netzservice GmbH, DE)

Introducing Anica as automated non-interference check assistant and Seda to do non-interference checks in business processes in a user friendly way.

3.11 Correct PAIS – A Naive Academic View

Niels Lohmann (University of Rostock, DE)

In my talk, I wanted to give an impression how academia copes with the verification of process-aware information systems. To achieve correctness, we can try to proof that a model satisfies a specification. As an example, I summarized work to verfiy information flow control for business processes. By transforming the non-interference property into a reachability problem, we can use existing algorithms and tools.

In a different school of thought, correctness by construction aims at generating a satisfying model out of the specification. To exemplify this, I sketched how models compliant to legal regulations and role-based access control can be synthesized automatically.

Both examples are supported by tools which scale well on industrial business process models. This encourages the hope of integrating verificiation techniques into practical tools in a spellchecking-style.

3.12 Data-Aware Business Processes: Formalization and Reasoning Support

Marco Montali (Free University of Bozen-Bolzano, IT)

The need for overcoming the traditional dichotomy between data and processes has been extensively advocated both by industry and academia. However, the "data-process engineering divide" still affects the majority of contemporary process-aware information systems. In this talk, we will review our recent research activities, whose common denominator is to demonstrate that making business processes data-aware paves the way towards a better understanding of process-aware information systems. Despite undecidability issues, we will show that it is possible to isolate interesting classes of data-aware processes for which reasoning support is indeed feasible. We will mainly focus on formal verification, synthesis and data quality concerns, giving also a glimpse on monitoring and mining.

3.13 Automatic Analysis and Certification of Policy Safety

Charles Morisset (Newcastle University, UK)

Attribute-based Access Control (ABAC) extends traditional Access Control by considering an access request as a set of pairs attribute name-value, making it particularly useful in the context of open and distributed systems. However, ABAC enables attribute hiding attacks, which allow an attacker to gain some access by withholding information. In this talk, we revisit in the context of the language PTaCL the notion of safety introduced by Tschantz and Krishnamurthi, and we present the tool ATRAP (Automated Term Rewriting

for Authorisation Policies), which allows for the automated analysis of safety for PTaCL policies, by producing counter-examples or generating Isabelle proofs of safety.

3.14 Trading Efficiency/Decidability for Expressiveness: Architecture Modeling with Fitzroy

Michael Norrish (NICTA – Canberra, AU)

The Fitzroy system provides a rich language for specifying component-based architectures. In emphasising data types such as lists and records, as well as full-blown arithmetic, we embrace undecidability, hoping that our reasoning support will work sufficiently well in pragmatically interesting cases. Our primary application is now modelling the architectures of systems built on top of the seL4 microkernel. Given that kernel's features, we expect many systems to be particularly concerned with security properties, and we have performed preliminary taint analyses to derive information flow properties of some sample system designs.

3.15 Process Mining: Discovering Process Maps from Data

Anne Rozinat (Fluxicon Process Laboratories, NL)

Most organizations have complex processes that are invisible, thus hard to manage or improve. Each stakeholder sees only part of the process. Manual discovery through workshops, interviews, and review of existing documentation is costly and time-consuming, and rarely reflects actual process complexity. Process mining closes this gap by making the real process visible. Our process mining software Disco leverages existing IT data to generate a complete, accurate picture of the process, with actionable insight. Disco automatically analyzes actual process flows, highlights bottlenecks, shows all variants, and allows animated "replay" of the process flow, all done interactively, driven by process questions.

3.16 Interval-based Process Monitoring for Uncertain Event Streams

Matthias Weidlich (Technion, IL)

Run-time monitoring of process execution is an important feature of process-aware information systems to satisfy security and compliance requirements. In this talk, we focus on two aspects of run-time monitoring, i.e., interval-based event semantics and uncertainty management. We argue that the lifecycle of business activities suggests to consider segmented interval events as the underlying model for reasoning. Further, we present three probabilistic methods for reasoning about constraints defined for segments of interval events for uncertain event streams for which the time of occurrence is unknown for certain events.

3.17 Preserving Demanded Privacy Constraints

Edgar Weippl (Secure Business Austria Research, AT)

The collection, processing, and selling of personal data is an integral part of today's electronic markets. However, the exchange of sensitive information between companies is limited by two major issues. Firstly, regulatory compliance with laws such as SOX requires anonymization of personal data prior to transmission to other parties. Secondly, transmission always implicates some loss of control over the data as further transmission is possible without knowledge of the data owner. In this paper, we introduce a novel approach that aims at solving both concerns in one single step, thus, our k-anonymity-based algorithm is at the same time able to anonymize and fingerprint microdata such as database records. Furthermore, we show that both the anonymization strategy as well as the fingerprint are collusion-resistant, which means that a group of data receivers is not able to subvert neither of the properties by combining their data sets.

3.18 Architecting with Architectural Design Decisions in the Context of Verifiably Secure Process-aware Information Systems

Uwe Zdun (University of Vienna, AT)

In the recent years, the software architecture community has proposed to use architectural design decisions (ADDs) for capturing the design rationale and recording the architectural knowledge (AK). While this approach helps to address key problems in architecting software systems like the gradual vaporization of AK over time, there are a number of important factors that hinder the widespread application of the approach. Examples are the considerable effort required for documenting ADDs in many existing approaches, the lack of integration with architectural design methods, and the limited understanding of how architects decide and understand software designs. The talk will introduce the topic and discuss key open challenges in the area of architecting with ADDs and present results from recent research in this area, especially with regard to and in the context of verifiably secure process-aware information systems.

3.19 Predictive Security Analysis @ Runtime

Maria Zhdanova (Fraunhofer SIT, DE)

Security Information and Event Management (SIEM) is a key technology to obtain a holistic view of an organization's security state and identify (emerging) security threats for timely and adequate response to security incidents. Today's service infrastructure paradigms such as IaaS and PaaS raise new challenges for SIEM solutions in terms of multi-level/multi-domain

security event processing and predictive security monitoring. The EU FP7 IP project MASSIF (Management of Security information and events in Service InFrastructures) aimed to create a next-generation SIEM architecture addresses these challenges in relation to four industrial scenarios: Olympic games, mobile money transfer service, critical infrastructure process control (dam), managed enterprise service infrastructures. Security monitoring in MASSIF SIEM is performed by the Predictive Security Analyzer (PSA) that implements predictive security analysis at runtime in order to identify misuse patterns in event streams. Predictive security analysis at runtime is a novel method for the evaluation of security-related events and their interpretation with respect to the known process control flow and given security properties. Based on real-time events from the process execution environment, a formal process process (APA) and security model (monitor automaton) the PSA allows to predict the close-future process behavior and detect possible violations of security requirements even before a critical situation occurs. One of the misuse cases that demonstrates the applicability of the PSA refers to the insider threat in a hydroelectric power plant.

Participants

- Rafael Accorsi
Universität Freiburg, DE
- Lujo Bauer
Carnegie Mellon University, US
- Anne Baumgraß
Hasso-Plattner-Institut – Potsdam, DE
- Nataliia Bielova
INRIA Rennes – Bretagne Atlantique, FR
- Achim D. Brucker
SAP Research – Karlsruhe, DE
- David Cohen
Royal Holloway University of London, GB
- Jason Crampton
Royal Holloway University of London, GB
- Christopher Dearlove
BAE Systems – Chelmsford, GB
- Guido Governatori
NICTA – St. Lucia, AU
- Christian Günther
Fluxicon Process Lab., NL
- Gregory Z. Gutin
Royal Holloway University of London, GB
- Michael Huth
Imperial College London, GB
- Fuyuki Ishikawa
National Institute of Informatics – Tokyo, JP
- Limin Jia
Carnegie Mellon University, US
- Mark Jones
Royal Holloway University of London, GB
- Günter Karjoth
Hochschule Luzern, CH
- Felix Klaedtke
ETH Zürich, CH
- Agnes Koschmider
KIT – Karlsruhe Institute of Technology, DE
- Jim Huan Pu Kuo
Imperial College London, GB
- Andreas Lehmann
Vattenfall Europe Netzservice GmbH, DE
- Niels Lohmann
Universität Rostock, DE
- Raimundas Matulevicius
University of Tartu, EE
- Marco Montali
Free Univ. of Bozen-Bolzano, IT
- Charles Morisset
Newcastle University, GB
- Alessandro Mosca
Free Univ. of Bozen-Bolzano, IT
- Michael Norrish
NICTA – Canberra, AU
- Andreas Oberweis
KIT – Karlsruhe Institute of Technology, DE
- Alexander Paar
TWT GmbH Science and Innovation, DE
- Stefanie Rinderle-Ma
Universität Wien, AT
- Anne Rozinat
Fluxicon Process Lab., NL
- Thomas Stocker
Universität Freiburg, DE
- Mark Strembeck
Wirtschaftsuniversität Wien, AT
- Meike Ullrich
KIT – Karlsruhe Institute of Technology, DE
- Matthias Weidlich
Technion, IL
- Edgar Weippl
Secure Business Austria Research, AT
- Nicola Zannone
TU Eindhoven, NL
- Uwe Zdun
Universität Wien, AT
- Maria Zhdanova
Fraunhofer SIT – Darmstadt, DE