Theories and Intricacies of Information Security Problems

Anne V. D. M. Kayem, Christoph Meinel (Eds.)

Technische Berichte Nr. 63

des Hasso-Plattner-Instituts für
Softwaresystemtechnik
an der Universität Potsdam

Universität Potsdam

HPI Hasso Plattner Institut
IT Systems Engineering | Universität Potsdam

Technische Berichte des Hasso-Plattner-Instituts für
Softwaresystemtechnik an der Universität Potsdam

Technische Berichte des Hasso-Plattner-Instituts für
Softwaresystemtechnik an der Universität Potsdam | 63

Anne V. D. M. Kayem | Christoph Meinel (Eds.)

Theories and Intricacies of Information Security Problems

Universitätsverlag Potsdam

Bibliografische Information der Deutschen Nationalbibliothek
Die Deutsche Nationalbibliothek verzeichnet diese Publikation in der Deutschen Nationalbibliografie; detaillierte bibliografische Daten sind im Internet über http://dnb.de/ abrufbar.

Universitätsverlag Potsdam 2013
http://verlag.ub.uni-potsdam.de/

Am Neuen Palais 10, 14469 Potsdam
Tel.: +49 (0)331 977 2533 / Fax: 2292
E-Mail: verlag@uni-potsdam.de

Die Schriftenreihe **Technische Berichte des Hasso-Plattner-Instituts für Softwaresystemtechnik an der Universität Potsdam** wird herausgegeben von den Professoren des Hasso-Plattner-Instituts für Softwaresystemtechnik an der Universität Potsdam.

ISSN (print) 1613-5652
ISSN (online) 2191-1665

Das Manuskript ist urheberrechtlich geschützt.

Online veröffentlicht auf dem Publikationsserver der Universität Potsdam
URL http://pub.ub.uni-potsdam.de/volltexte/2013/6045/
URN urn:nbn:de:kobv:517-opus-60455
http://nbn-resolving.de/urn:nbn:de:kobv:517-opus-60455

Zugleich gedruckt erschienen im Universitätsverlag Potsdam:
ISBN 978-3-86956-204-9

FOREWORD

Welcome to the First Workshop on Theories and Intricacies of Information Security Workshop (INTRICATE-SEC"12) and to Johannesburg, South Africa. In recent years, information security systems have become increasingly popular because of hardware and software innovations. Powerful processors make time intensive computations faster but have also made security violations much easier. Attacks on computing systems have shifted from being simple attacks perpetrated by socially maladjusted teenagers to a concerted criminal activity. It therefore makes sense to explore alternative or unconventional approaches to solving problems of information security including the support of a variety of user requirements as well as provider requirements. The fact that computer systems are changing rapidly and extending their boundaries in environments such as cloud computing and service oriented architectures introduces more intricacies in security measures.

INTRICATE-SEC"12 provides a meeting place for researchers to contribute papers, by invitation only, and discuss cutting edge research on the theories and intricacies of provisioning security. The workshop has attracted 5 submissions from different but complimentary topics of Information Security. In ordering the papers we have ensured that there is a logical flow from one to the next, the aim being to emphasise our workshop's theme.

We would like to express our sincere gratitude to the Information Security South Africa Conference (ISSA12) organizing committee and the Hasso-Plattner-Institut for sponsoring and hosting the workshop. More importantly, we thank all the authors for their submissions and contributions.

Program Committee Co-Chairs
Anne V.D.M. Kayem, University of Cape Town, South Africa
Christoph Meinel, Hasso-Plattner-Institut, University of Potsdam, Germany

SECURITY ENHANCEMENT WITH FOREGROUND TRUST,COMFORT, AND TEN COMMANDMENTS FOR REAL PEOPLE

STEPHEN MARSH, ANIRBAN BASU, AND NATASHA DWYER

Abstract. Security as an enabling paradigm has not succeeded half as well as we might have hoped. Systems are broken or breakable, and users (people) have something of a lack of faith, understanding, or patience with security measures that exist. Whilst secure systems and solutions are the backbone of a working interconnected system of systems, they are not people-oriented, and they are oftentimes arcane enough to have an air of 'security theatre' about them. We can also assume that they will continue to grow in both complexity and application if we continue as we are in our arms race.

To answer what we perceive to be a problem here, we are working on the integration of socio-psychological notions of trust into computational systems where it makes sense (both human- and system-facing). This work includes the development of our Device Comfort paradigm and architecture, wherein mobile devices and nodes in infrastructures have a embedded notion of comfort that they can use to reason about their use, behaviour, and users. This notion, contextually integrated with the environment the device is in, aids in decision making with regard to, for instance, information flow, security posture, and user-oriented advice. Most importantly, the notion embeds trust reasoning and communication into the device, with which the user can be aided to understand situation, risk, and actions by device, infrastructure, and themselves - which we call Foreground Trust, after Dwyer. We conjecture that comfort and foreground trust both enhance security for devices and increase the understanding of security for the user, through use of human-comprehensible and anthropomorphic concepts. In this paper, we discuss some security problems, address the misnomer of trusted computing, and present an overview of comfort and foreground trust. Finally, we briely present our ten commandments for trust-reasoning models such as those contained within Device Comfort, in the hope that they are of some use in security also.

1. Introduction

We live in a complex world, one in which decisions about security of many kinds occupy an important place. For as long as there have been people, these decisions have been important. The difference between ten, a hundred, or a thousand years ago and now is the tools we use. Like any power tool, computers, either desktop, laptop, or in our pockets, help us to do things more quickly. They also help put us into difficult situations more quickly. Information power tools, which is what computers are, have potentially exposed their users' information - private, heretofore shared only with a chosen few, to the many. Protecting this information, as well as the tools themselves and the access they have to others' tools and information, is

Key words and phrases. Trust, Device Comfort, Foreground Trust, Human Computer Interaction.

the task of information security. Attacking the information, for pleasure or profit, is the ‘task’ of the ‘adversary.’

We have reached an unfortunate stage in the evolution of information systems, however - we conjecture that, if a system is not compromised it most certainly can be, and that the attacks are coming faster and most importantly in unique ways - while we still have the worms and viruses of the past, we now have to contend with social engineering and targeted attacks. To do this, we pour more and more intellectual capital into defences against the adversaries. But to what end? Systems now not compromised can be, and many are, with or without our knowledge. Couple this with the increasing complexity of the defences themselves, which ultimately results in more frustration at the very least on the part of the users we are trying to defend, and we arrive at a challenging confusion: the system is broken. An arms race has been carrying on for many years, and the only loser is the person who wants to get her job done, or play.

Enhanced security mechanisms, better passwords, different kinds of passwords, more complex login procedures, more demands on users, or abrogation of responsibility are not the answers at least, they’re not the answers that make for security with users in mind. We find a little solace in usable security, but ultimately we feel that we should look for a more human-centric approach to security. To achieve this, we approach the problem from the point of view of human social norms, in particular, ion our work, trust and comfort, and their darker siblings distrust an discomfort. In fact, these are topics we have been dealing with for some time in different areas, including agent systems and critical infrastructures. They lend themselves particularly well to human-oriented security because they are in essence human-oriented security - and they have worked for humans for millennia. The paradigm that most interests us in this instance is what Dwyer [4] calls Foreground Trust - in essence the ability of technological devices to present information to users in order to allow them to make security-focused (trust-focused) decisions. Our most recent work in this area has been concerned with integrating comfort and trust reasoning techniques into mobile devices, which we call Device Comfort.

This paper discusses the Device Comfort paradigm as a security tool for both information and personal security from a high level perspective (interested readers can find more information in other work [11]). As well, we will discuss a set of ‘commandments’ for trust and comfort facing the user, commandments that we feel can benefit any security technology where humans are a concern (and this is, of course, all of them). The next section introduces and briefly discusses Foreground Trust, before we proceed to a slightly more in depth discussion of Device Comfort as a form of Foreground Trust in section 3. Section 4 takes the form of a discourse on presenting information to the user, and is an extension of work we did in [9]. After a discussion and a brief look at ongoing work in Device Comfort in section 5 we conclude in section 6.

2. Foreground Trust and Related Work

Trust Enablement was presented in [4], where it was applied to the task of the system to connect people through technology by giving them the tools and information they need to make trusting decisions regarding other people. Extended to Foreground Trust, it was further expounded upon in [13]. The basic premise of Foreground Trust is this: if people have enough information to make trusting decisions, they will do so. We acknowledge that there is, of course, rather a lot of trust placed in the person in this instance, which we will address with regard to information security in the next section. However, the premise is inherently sound - people 'get' trust, they understand it in a very deep sense, and it is a tool that has evolved over millennia to allow humans to make decisions or handle complexity in the face of risk [2, 6, 3, 7, 12].

Extending the concept of Foreground Trust to a human-technology relationship is a necessary next step in evolving a security solution that is human-oriented. In some settings this should be closely integrated with societal behaviours and norms. Work by Murayama et al on the Japanese concept of Anshin is closely related work in this area [5, 14]. In [1] we find behavioural history based reputations to inform security decisions in networks, and also the outlook that one entity's view of the network is its own when it comes to security; and that this 'personal' view (i.e. local reputation of other entities) is more important than the global view.

The benefits of this approach are manifest: humans, as has already been noted, understand trust, and they understand how trust decisions are made. Integrating these decisions into the interface between human and 'security' is, we conjecture, a sensible approach to allow the human to understand the security risks and resultant posture of the systems they are using to get their work done. Ultimately, the Device Comfort paradigm is an extension of trust reasoning into mobile devices that are inherently human-facing.

3. Device Comfort

Device Comfort extends trust reasoning technologies by allowing mobile devices to reason with and about computational trust [7] and to communicate these reasonings to the owner of the device in a human-oriented manner. The Device Comfort paradigm goes further in that it allows the device itself to make trust-based decisions, comfort-based decisions, and policy-based decisions independent of the user, and adjust its security posture accordingly. We have written extensively about Device Comfort elsewhere [8, 10, 11] and so only briefly discuss the notion here as it applies to a non-traditional security measure.

Device Comfort was initially envisaged as a tool to help teens using smartphones make more sensible decisions about what they are using the phones for (in particular, with regard to the phenomenon of Sexting). However, we quickly became aware that the paradigm had applicability in many different uses of mobile devices for many different users, and have altered our outlook accordingly.

The premise of Device Comfort is quite simple: to Advise, Encourage, and Warn (as for a constitutional monarch, in fact) and ultimately be able to proscribe actions for the users of the mobile device (call it AEWP). Indeed, we'd go as far as claiming that this is what all security methodologies and tools should be doing. It does this by using the strengths of the device as a sensing mechanism, as well as having inbuilt security policies. Device comfort is a dynamic phenomenon the more sense-capabilities a device has the more we can integrate into comfort. Currently we see Device Comfort as a measure based on reasoning about the following:

- The user's identity
- Enhanced trust reasoning about the user, and the ongoing relationship with respect to trust that the device has in the user (and/or owner);

- The current task (for instance, making a call, sending text, sending pictures, email, etc.)
- The current location (which virtually all mobile devices can determine with some accuracy)
- A Comfort Policy-base (provided by the owner of the device, as well as the owners of any information the device stores or can access, basically presented to the device, and thus the user, on access.)

A more formal exposition is given in [11]. We see Comfort as a dynamic, context-based reasoning mechanism. The internal architecture of the technology is based on sensing tools communicating with a comfort agent via tuple spaces, and the agent communicating with the user through a sensible interface. We have very carefully considered how the mechanism becomes human-oriented through the interface, both as what we call 'annoying technology' [11] and as an embodiment of the AEWP concept. We aim, through the interface to give the user second thoughts, encourage anticipatory regret if possible, and learn from what the device can tell them about the situations they find themselves in - and in this way, to learn about security for themselves and their information.

3.1. **The Human Security Posture.** Because Device Comfort was devised as a tool to allow potentially less-aware (or less concerned) users to make sensible risk-based decisions, it is unique in that it is human-oriented as well as seeking to allow humans to understand their context not only for information security, but also for personal security. It is possible to envisage situations where the device is 'uncomfortable' because of aspects of its physical context location, electronic neighbourhood (what devices are present), and so on, as well as its electronic context, including policies. In these instances, it is useful to present these aspects of context to the user as something that they should be concerned about and want to change, because they themselves may be at risk, and not just the device of the information on it (and so: "leave the area" is a valid comfort response, for instance, as is "don't send that message from here"). In essence, the tool has a use in the formation of 'second thoughts' for users in risky situations [13]. Whilst we have not as yet conducted experiments involving this aspect of Device Comfort, we hope to be able to address it in the near future.

4. The Ten Commandments of Foreground Trust for Security

In [9] we presented a discussion of the complexity of trust models that are, as all should be, human-focused. The arguments there are similar to those in this paper: too much complexity is not helpful, for instance. We did, at that time, have eight commandments for trust models. In this paper, we extend this to ten (there should always be ten, after all) with the addition of two commandments which were discussed as a result of that previous work. We also extend the commandments to take into account security models and techniques, in the hope that they can be more generally applied. We make no claim to uniqueness in these commandments - indeed we would be heartily surprised if they were not in some shape or form discussed elsewhere in the usable security world. But we do claim that they are useful to bear in mind in the trust world too, an they are most certainly applicable in any case. The commandments, and a brief discussion, follow.

(1) The model is for people.

(2) The model should be understandable, not just by mathematics professors, but by the people who are expected to use and make decisions with or from it.

(3) Allow for monitoring and intervention. Humans weigh trust and risk in ways that cannot be fully predicted. A human needs to be able to make the judgement, especially when the model is in doubt.

(4) The model should not fail silently, but should prompt for and expect input on failure or uncertainty.

(5) The model should allow for a deep level of configuration. Trust and security models should not assume what is 'best' for the user. Only the user can make that call.

(6) The model should allow for querying: a user may want to know more about a system or a context.

(7) The model should cater for different time priorities. In some cases, a trust/security decision does need to be made quickly. But in other cases, a speedy response is not necessary

(8) The model should allow for incompleteness. Many models aim to provide a definitive answer. Human life is rarely like that. A more appropriate approach is to keep the case open; allowing for new developments, users to change their minds, and for situations to be re-visited.

(9) Trust (and security) is an ongoing relationship that changes over time. Do not assume that the context in which the user is situated today will be identical tomorrow.

(10) It is important to acknowledge risk up front (which is what all trust, including Foreground Trust, does).

5. Discussion and Ongoing Work

The commandments translate into design principles guiding the creation of a Comfort Device that negotiates trust, security and control on behalf of and in conjunction with the user. Knowledge about how the user would want to operate such a device is necessary. We are working on a series of research questions to inform our design perspective. For instance, we argue that there are some trust decisions that the user is more interested in than others. But how do we distinguish between these interactions? Catering for interactions will differ depending on users needs and interests. Other questions include how users differ amongst themselves between what is the priority in an interaction.

The intricacies of an interface are also relevant for a successful project. What graphic style is suitable? With what sort of language does the user want to be informed of risk in the beginning of an interaction? Formal? Colloquial? Although these questions might seem to be those that need to be addressed at the later stage of delivery or even trivial, in actuality such considerations shape how users approach issues of security and trust. Are users welcomed to consider these complicated notions on their own terms or alienated?

It should be clear that, in all cases systems such as those outlined here have a need for a solid foundation. Security, in the sense of a secure system (as secure as we may be able to make it, and at as low a level as possible), is necessary, but not sufficient to bring about the needed interactions and relationships that the system must have with the user, and vice versa. Ultimately, trust and comfort serve to make security more flexible. The major difference between traditional security and trust/comfort is the way in which risk is handled: security, including trusted computing, aims to minimize or eliminate risk if possible, resulting in a so-called 'trusted' system. Ironically, given that trust is founded upon risk, this process would if possible reach the situation where trust was in fact not needed. Trust and comfort acknowledge, accept, and manage risk in context. We feel that this is a more flexible approach because it accepts the potential for all systems to be compromised and to try to work to the best ability anyway, whilst not failing silent (see the commandments). While we are in a position of needing security at some level, we feel that we can enhance security both by this acceptance and by the raising of awareness the trust relationship can effect with the user.

As well as the questions above, we are working on integration of trust and comfort reasoning into more complex settings including critical infrastructure interdependencies and management, and human-oriented processes.

6. Conclusions

Trust and Comfort are flexible, awareness-enhancing approaches to risk in context. They encourage flexible adaptations to risk that are human oriented and seek to enhance understanding of security in the people using them. Our work in this area has been exploring formal models for comfort based on our extant trust models, the integration of comfort into mobile devices, the design of comfort enabled user interfaces, and the extension of comfort in different infrastructures and contexts.

Of necessity, space for this paper is short. There is, however, much to say and much to be done on the topics of Foreground Trust, Enablement, and Comfort for security, and we hope that this short paper has shed enough light on the topic and its considerations to interest the reader.

References

[1] A. Basu. A Reputation Framework for Behavioural History. PhD thesis, University of Sussex, UK, January 2010.

[2] S. Bok. Lying: Moral Choice in Public and Private Life. Pantheon Books, New York, 1978.

[3] M. Dibben. Exploring Interpersonal Trust in the Entrepreneurial Venture. London: MacMillan, 2000.

[4] N. Dwyer. Traces of Digital Trust: An Interactive Design Perspective. PhD thesis, School of Communication and three Arts, Faculty of Arts, Education and Human Development, Victoria University, 2011.

[5] N. Hikage, Y. Murayama, and C. Hauser. Exploratory survey on an evaluation model for a sense of security. In H. Venter, M. Eloff, L. Labuschagne, J. Eloff, and R. von Solms, editors, IFIP Volume 232, New Approaches for Security, Privacy and Trust in Complex Environments, volume 232, pages 121-132, Springer, 2007.

[6] N. Luhmann. Trust and Power. Wiley, Chichester, 1979.

[7] S. Marsh. Formalising Trust as a Computational Concept. PhD thesis, University of Stirling, 1994. Available via www.stephenmarsh.ca.

[8] S. Marsh. Comfort zones: Location dependent trust and regret management for mobile devices. In In Proceedings TruLoco 2010: at IFIPTM 2010, Morioka Japan., 2010.

[9] S. Marsh, A. Basu, and N. Dwyer. Rendering unto Caesar the things that are Caesar's: Complex trust models and human understanding. In T. Dimitrakos, R. Moona, D. Patel, and D. H. McKnight, editors, Proceedings Trust Management VI: IFIPTM Conference on Trust Management, pages 191-200. Springer (IFIP AICT), 2012.

[10] S. Marsh and P. Briggs. Defining and investigating device comfort. In Proceedings of IFIPTM 2010: Short Papers, 2010.

[11] S. Marsh, P. Briggs, K. El-Khatib, B. Esfandiari, and J. A. Stewart. Defining and investigating device comfort. Journal of Information Processing, 19:231-252, 2011.

[12] S. Marsh and M. R. Dibben. Trust, untrust, distrust and mistrust: an exploration of the dark(er) side. In Peter Herrmann, Valerie Issarny, and Simon Shiu, editors, Trust Management: Proceedings of iTrust 2005. Springer Verlag, Lecture Notes in Computer Science, LNCS 3477, 2005.

[13] S. Marsh, S. Noël, T. Storer, Y. Wang, P. Briggs, L. Robart, J. Stewart, B. Esfandiari, K. El-Khatib, M. Vefa Bicakci, M. Cuong Dao, M. Cohen, and D. Da Silva. Non-standards for trust: Foreground trust and second thoughts for mobile security. In Proceedings STM 2011. Springer, 2012.

[14] Y. Murayama and Y. Fujihara. Issues on Anshin and its factors. In Zheng Yan, editor, Trust Modeling and Management in Digital Environments: From Social Concept to System Development, pages 441-452. IGI Global, 2010.

Acknowledgement 1. *The authors wish to thank the many researchers and colleagues who have helped on the Device Comfort project, including Tim Storer, Babak Esfandiari, Mehmet Vefa Bicacki, John Stewart, Sylvie Noël and Lewis Robart.*

(S. MARSH)
University of Ontario Institute of Technology,
Faculty of Business and Information Technology
2000 Simcoe Street North, Oshawa ON. Canada
E-mail address, S. Marsh: stephen.marsh@uoit.ca
URL: http://www.stephenmarsh.ca/

(A. BASU)
Tokai University,
2-3-23 Takanawa, Minato-ku,
108-8619, Japan
E-mail address, A. BASU: abasu@cs.dm.u-tokai.ac.jp, a.basu@sussex.ac.uk

(N. DWYER)
Victoria University,
Footscray Park Campus, Ballarat Road,
Footscray, 3011, Australia,
E-mail address, N. DWYERM: natasha.dwyer@vu.edu.au

INFORMATION SECURITY INNOVATION: PERSONALISATION OF SECURITY SERVICES IN A MOBILE CLOUD INFRASTRUCTURE

JAN H.P. ELOFF, MARIKI M. ELOFF, MADELEINE A .BIHINA-BELLA,
DONOVAN ISHERWOOD, MOSES DLAMINI, AND ERNEST K. NGASSAM

Abstract. The increasing demand for on-line and real-time interaction with IT infrastructures by endusers are facilitated by the proliferation of user centric devices such as laptops, iPods, iPads, and smartphones. This trend is furthermore propounded by the plethora of apps down loadable to end-user devices mostly within mobile-cum-cloud environments. It is clear that there are many evidences of innovation with regard to end-user devices and apps. Unfortunately little, if any, information security innovation took place over the past number of years with regard to the consumption of security services by end-users. This creates the need for innovative security solutions that are humancentric and flexible. This paper presents a framework for consuming loosely coupled (but interoperable) cloud-based security services by a variety of end-users, in an efficient and flexible manner using their mobile devices.

1. Introduction

The increasing demand for cost-effective always on connectivity on all types of end-user computing devices (e.g. desktop computer, laptop, MP3 player, tablet, smartphone) results in the need for new business models (mobile, cloud, services, platforms) that increase the level of exposure to a companys assets. This creates new security challenges for networked businesses as a number of 3rd-party services and infrastructures within complex ecosystems are integrated.

For instance, many actors are involved in the service provisioning ranging from the customer, the service provider, the content provider, the network provider, the cloud provider and the electronic or mobile payment provider. Each of these actors has an entry point to the service and therefore is a potential security risk. Investigating a security breach thus requires the collection of data from all these different sources. In addition, the existence of various mechanisms to access the network (e.g. wired, wireless, 3G, modem, VPN) creates many access points that can be exploited for unauthorised access to and misuse of the companys information. Detecting such events requires the continuous exchange of information between all service elements and network devices [1]. Furthermore, entities involved in the service provisioning can have conflicting security policies that need to be aligned to the companys policy.

In this collaborative environment, security risks shifts from the IT system as a whole to the services it offers to a multitude of independent users and to the data

Key words and phrases. security; innovation; mobile computing; cloud computing; mobile-cloud; security services; cloud architectures; service provisioning; service consumption; access to services.

that travel across systems (e.g. in cloud computing applications hosted on public infrastructures). For example, applications hosted on public cloud infrastructures are not only open to the general public but are also open to malicious individuals. Such applications become a public good and are susceptible to excessive and malicious use. Malicious or disgruntled individuals may decide to flood such applications with targeted distributed denial of service (DDoS) attacks so that the general public could not have access to them.

Maintaining a secure configuration in such heterogeneous IT landscapes is complex as security requirements are multi-lateral and diverse. This creates the need for innovative security solutions that are humancentric, flexible and also robust. Potential avenues for innovation within the information security domain include, amongst others, the following:

(1) The definition of data-centric policies that travel with the services as well as the data
(2) the usage of privacy-preserving computing [2] to ensure the privacy of all parties involved
(3) access control policies and mechanisms that take care of conflict management [3] between the members of an ecosystem
(4) the possible aggregation of different access control approaches such as usage and optimistic based access control [4].
(5) simple and basic authentication services on mobile devices
(6) forensic tools for mobile-cum-cloud environments [5] services utilization, using mobile devices.

This is an opportunity to capitalize on the advantages offered by cloud computing for accessing value-added business services, by end-users. In general, end-users are not concerned by the complexity of the technical infrastructure required to set up cloud-based services for large consumption but rather the intended business outcome offered by exposed services.

This paper presents an innovative framework for accessing loosely coupled (but interoperable) cloud-based security services by a variety of end-users, in a secure, effective and flexible manner, anywhere and anytime, using their mobile devices.

The remaining part of this paper is structured as follows: Section II provides some background information discussing the concept of innovation within the domain of information security. Furthermore, background information is provided on the current state-of-the-art in information security services and existing approaches to services oriented architectures. In section III, a generic SiYP (Security-In-Your-Pocket) platform is presented from a Service Oriented Architecture (SOA) point of view. Section IV presents a conclusion and future work.

2. Background and Related Work

From the previous section it is clear that various technologies, tools and devices will form part of this proposed framework in order to provide the required flexibility, efficiency and security. In this section the different terms and technologies will be discussed as well as how they relate to each other. IDC [6] identified the importance of leveraging the strengths of both innovation and security in order to gain a competitive advantage over companies who do not do it. However, no extensive research was found that addresses how information security can strengthen innovation and vice versa. In this section these terms will be defined as well as their interrelationship. Some well-known terminologies are associated with information security such as information security services, privacy and trust. However, it is important to understand how, for example, trust, relates to service oriented architecture and innovation.

2.1. **Innovation.** For any organisation to gain and maintain a competitive advantage and be an economic leader, being innovative is of utmost importance. Innovation should not be confused with invention, which is only the idea or model for a new or improved product, process, device or system, whereas innovation is bringing this idea to market as a real product, process, system or device that is part of the economic system [7]. Innovation is only accomplished with the first commercial transaction involving this new invention.

Schumpeter said, in as far back as 1934, that Economic change revolves around innovation, entrepreneurship and market power [8].

2.2. **Information security services.** According to the ISO 7498-2 standard, produced by The International Standards Organisation (ISO) information security can be defined in terms of the five security services, namely identification & authentication, authorisation, confidentiality, integrity and nonrepudiation [9]. These services are required to ensure that information are protected and secured at all times, whether in storage of any nature, during transmission or usage. A definition for each of these services follows:

The identification and authentication of any subject who wants to access any computer system is the first step towards enforcing information security. A subject requesting access needs to present a user-id that uniquely identifies it. On presentation of such a user-id, the user-id should be verified to ensure that it does, in fact, belong to the subject who presented it.

The next step is to determine if the authenticated subject has the right to access the computer facilities in question. In terms of the authorisation process, control is, therefore, exerted over the access rights of all authenticated subjects.

All information must be strictly accessible to authorised parties only. Protecting the confidentiality of information, therefore, gives the assurance that only authorised parties will have access to the information in question.

Information should not only be kept confidential, but its integrity should also be guaranteed. Only authorised parties should be able to change the content of protected information. In other words, unauthorised changes to information must be prevented, ensuring that the information can be deemed accurate and complete.

The last step is to ensure that no action is performed to affect information security, for example, changing some of the content of information that could be denied at a later stage. This process is referred to as non-repudiation. It may be argued that the security services are not applicable in current computing applications; however, various authors have proven, through research, that these services are essential, especially in mobile and cloud computing [10,11].Subsection text.

2.2.1. *Privacy.* Privacy, which is closely related to information security, is defined as the right of an individual or a group to isolate information about themselves from others. This ability allows individuals to reveal themselves selectively. The Oxford English Dictionary refers to privacy as the condition of free from public attention, undisturbed, or the freedom from interference or intrusion [12].

2.3. **Trust.** Trust is in principle a human action. A person may trust another to behave in a certain manner, where trust is based on past experience, recommendation or the reputation of the other person. The Oxford English Dictionary [12] defines trust as the confidence in or reliance on some quality or attribute of a person or thing, or the truth of a statement. This definition of trust does not take all aspects of trust into consideration nor does it satisfy the requirements of trust as required in the Web 2.0 environment [13]. Trust is bi-directional with mainly two parties involved, namely the truster and the trustee. The truster is the party who trusts, confides, or relies on the other party; he/she is the one who believes or credits; the one that gives credit, a creditor. The trustee is, on the other hand, the party who is trusted, or to whom something is entrusted; a person in whom confidence is put [12]. Jsang et al [14] goes further and distinguishes between reliability trust and decision trust. Reliability trust implies that trustee will act as expected, while decision trust refers to the situation where the truster depends on the trustee, even though some risks may be involved.

2.4. **Service oriented architecture and cloud computing.** Service oriented architecture (SOA) developed from older concepts such as distributed computing and modular programming into cloud computing. Cloud computing can be seen as a service-oriented architecture (SOA) exploring almost every computing component including, but not limited to distributed computing, grid computing, utility computing, on-demand, open source, Peer-to-Peer and Web 2.0 [15]. It is a natural next step from the grid model to a supply and demand utility model. In minimizing potential security and trust issues as well as adhering to governance issues facing Cloud computing, prerequisite control measures are required to ensure that a concrete Cloud computing Service Level Agreement (SLA) is put in place and maintained when dealing with mobile applications.

2.5. **Mobile computing.** Mobile computing refers to any computing device that possesses processing and storage capabilities and that can connect to other computing devices, preferably through wireless connections. Devices include, but are not limited to cellular phones, tablets, PDAs, laptops, notebooks etc.

2.6. **Security in mobile and cloud computing.** Trust is a key element of security in cloud computing. If one party is not trustworthy, it is clear that this partys security, even if claimed to be strong, is not security at all [16]. One of the most important protocols in ensuring transparency and security within Cloud computing is the SLA. The SLA is the only legal agreement between the service provider and client and its importance should not be under estimated [17]. The only means that the cloud provider can gain the trust of clients is through the SLA, therefore the SLA has to be standardised. The following are the main aspects as a guideline for SLA:

(1) Services to be delivered and performance
(2) Tracking and Reporting
(3) Problem Management
(4) Legal Compliance
(5) Resolution of Disputes Customer Duties
(6) Security Responsibility
(7) Confidential Information Termination

However, ensuring mobile and cloud security is still a serious challenge as identified and addressed by various research studies [18-20].

From the above discussions it is clear that for the proposed framework it will be important that the existing security services should be embedded in the SOA as well as in the mobile-cloud infrastructure. The user will play an important part in these security services. The mobile applications, being loosely coupled, propose unique security challenges.

3. A GENERIC SERVICE MODEL FOR THE SECURITY IYP PLATFORM

As eluded in previous sections, the intended purpose of a user centric security platform is to supply the user with the ability to bind available apps to be consumed as service. This requires a range of security services for ensuring that all transactions performed during the consumption of cloud-based services are secured and guaranteed, not only for good quality of services, but also to adhere to prescribed service level agreements amongst interoperable services.

In general, the SiYP platform is regarded as a container of apps grouped together and supported by a range of loosely coupled, security-based, services deployed in the cloud. Security-based services will be implemented and deployed in the cloud for consumption based on the various constraints required for ensuring that any app consumed by end-users adhered to the prescribed security standard. However, the approach adopted in this new paradigm provides the flexibility for the end-user to wrap the business app with the minimum necessary security app (and therefore service) required for an effective consumption without any security breach. For illustration purpose, an end user who would like his device to be secured based on its location while using a given business app would configure the business app such that the location service provided in the platform is activated. This approach provides interested users, to only rely on the consumption of security services that are necessary while consuming a given business app through the cloud. Equally, the end-user would have the ability to enforce the traceability of all its operations while consuming the service, provided that he has activated the secured security app in the platform, dedicated to such a role. Of course, one may argue that security enforcement would require that all security-based services are activated [9]. However with the limited computational resources available on mobile devices as well as the complexity that could arise if all those services are invoked, it makes perfect sense to adopt the approach of delivering those services ondemand, to be consumed by end-users on an as-needed basis.

It follows that the overall structure of the platform would reflect the manner in which apps are usually grouped together based on the kind of services being rendered by any given app. In the app world, each app as presented to end-users contains the necessary functionality that reflects the service to be provided to the consumer. As such, there is a range of services at the lower level of the hierarchy that are invoked and aggregated in order to meet end-user's expectation. Therefore two different apps may share many services although their end-result appears different to the end-users. This is a true reflection of the principles of loosely coupling and interoperability that make up a service model following modern Service Engineering (and therefore SOA) paradigm [21]. As such the conceptualisation of our proposed

SiYP platform would equally be based on current trends, as reflected in the state-of-the-art of Services Engineering [21].

We perceive the SiYP service model, therefore, as a four-tier architecture. As a naming convention, all the apps are called “My” followed by “App name” and the term security is implied. For instance My Device indicates that this is the app for the user device security. The description of each layer of the architecture as well as components thereof follows.

3.1. **The Presentation Layer.** This layer represents the entry-point to the security innovation platform by end-users using the necessary computational medium for the consumption of any given business app. For instance, the Business in Your Pocket (BiYP) interface available on the end-user’s device would enable him/her to consume business services in the cloud anywhere and anytime. In order to do so, the flexibility of the platform would allow the user to personalise a range of security apps required for the secure consumption of all business services attached to the app. Discussion on those business services are beyond the scope of this work. Since the end user might require that security, with respect to the device being used for service consumption, be enforced, he/she then has to “wrap-on” the business app, the appropriate security app available from the lower layer in the hierarchy. This security app also invokes a range of security services from lower layers as well as required security services for interaction with third party API’s that are not part of the SiYP platform. The presentation layer is therefore the environment used by consumers to personalise their required security services in order to ensure that security is enforced during the consumption of a given business app.

3.2. **The Application Layer.** This layer provides a pool of security apps that can be added to and removed from a business service based on the user requirements. Examples of some possible security apps are described below. They are My Device, My Failure Prevention, My Cloud and My Social Network.

3.2.1. *My Device app.* This application enables users to select the level of security for the specific device they are using to connect to the corporate network. Security levels will differ from one computing device to another based on the device features and usage profile of the user. To this effect, my Device will use services such as Location, Connectivity and Logging. For instance, one user may use his laptop and smartphone mainly in the office but his iPad and iPod mainly out of the office. The laptop is mostly used for processing power intensive usage such as creating and editing documents, running and installing applications while the iPad is used primarily to read documents as well as to read and write emails. The iPod may be used for Wi- Fi access to his corporate emails. This information could be recorded in his corporate user profile along with identifiers of all the devices he/she uses. This can allow single sign-on with his corporate user credentials on all devices. The security level for a device will be defined in terms of various services such as Location, Connectivity, and Logging.

The Location service recognizes whether the user is in a trusted familiar environment (e.g. home, office, regular coffee shop) and adjusts the security features accordingly. In an unknown location, security will be tighter (e.g. 2- step authentication) and geo-location data will only be sent to authorized recipients.

The Logging service keeps a record of the users activity and the data stored on the device that is not linked to any application (e.g. multimedia file, software and hardware identifiers). This is used to specify the level of security for the user data according to the selected device (e.g. encryption of transmitted data). This is also applicable to the SIM card data. If one SIM card is transferred from one device to another, security policies applicable to this device can be enforced on the SIM card as well.

The Connectivity service specifies the security level required based on the connection channel used (e.g. 3G, Wi-Fi or VPN). For instance, accessing the corporate network from a public non-encrypted Wi-Fi access point will require tighter controls than access through a more secure VPN.

3.2.2. *My Failure Prevention app.* This application will monitor the service consumed to detect unsafe events and situations that can lead to a failure. The application will use the Logging service to log such events and any associated data for a possible future root cause analysis and to prevent the recurrence of such events. An example of such an event, referred to a near miss, is the near exhaustion of critical resources (e.g. memory or battery power). The user will have the choice to specify a near miss threshold that indicates how close the near miss is to cause a failure to generate an alert [22].

This app can also be used to prevent misuse of the business service by creating a service profile. The service profile specifies how the business service is normally used by the average user. It can thus help detect suspicious activity that deviates significantly from the normal usage pattern and which can indicate that the service is being misused. The profile will provide information such as the average spending and the usual time and duration of the service usage.

3.2.3. *My Cloud app.* My Cloud app uses the cloud infrastructure to deliver security services to customers and other cloud service providers. This is referred to as Security as a Service (SecaaS) [23]. SecaaS is defined as the on-demand provisioning of cloud-based security services either to hosted cloud infrastructure and software or to clients onpremise private systems [23]. My Cloud leverages a cloud-based model to deliver security service-based offerings which include end point protection, on-demand transparency-enhancing technology, identity as a service, risk-based authentication etc.

Cloud-based end point protection leverages the SaaS model to provide cost effective anti-virus and anti-spyware services to all types of end point devices (desktop, laptops, tablets and smartphones) that each user employs when connecting to corporate resources. This service also maintains and manages all updates to each of these end points.

On-demand transparency-enhancing technology leverages the auditing and logging capability of cloud based systems. It provides customers with a clear indication of all the data centers where their data is hosted and/or replicated and who is accessing it and for what purposes. The customers can query their cloud service providers for the exact location of where their data sits at any particular time. And they could also query the cloud service provider for all access to their resources within a given time period, for example over the past 12 months. Providing the customer with such transparency can help ensure that customer data can only be stored in data centers in approved geographical locations. The cloud security app will help solve most of the jurisdictional compliance challenges.

Nowadays, organizations are faced with a workforce that is highly mobile. This type of a workforce constantly require access to corporate resources from wherever they sit (on-premise or hosted), at all times, from any location, using a plethora of devices [24]. This raises concerns on how to keep corporate resources available to only those who must have access. How to ensure that the right people, systems, and end point devices have the necessary access privileges to both on-premise and hosted cloud services. Cloud computing through Identity as a Service (IDaaS) presents a platform that does exactly that. It decouples the provisioning, de-provisioning, maintaining and managing users from the apps. This app provides digital identity as a service. IDaaS provides an API that authenticates users from different directories (e.g. Active Directory, LDAP or Web Service) on cloud-based apps and supports cross-domain authentication.

3.2.4. *My Social Network app.* My Social Network app monitors and manages all social media accounts of a user from one place. This app offers social network Single Sign-On portal. You sign-in on one Social Network and gain access to everything you need. For a first time user, he/she will register once and then choose from a wide pool of all the social networks that he/she wishes to gain access to. The app will then use the provided registration information to register the user on all the other social network sites. For those users who are already registered with different Social Networks and have different credentials, this app integrates all their accounts and put them under one umbrella called My Social Network app. The user can use this app to log on supplying it with his/her credentials for one Social Network where he/she is registered e.g. Facebook and then gain access to all his/her social media accounts.

This app use federated identity to authenticate the user only once and then share the credentials in assertions across all registered Social Networks. It provides the user in real-time with all the RSS feed updates happening in all of the social networks where he/she is registered, all in one place. This app ensures that the users are no longer required to manage multiple user accounts for multiple Social Networks. When the user logs out of the My Social Network app it automatically logs him/her out in all of the social networks.

Furthermore, this app uses a mobile devices built-in GPS to determine a users location at any given time. After picking the users location, it then scans the friend lists on each of his/her registered social media accounts to see if any one of them is in the same area. It then notifies the user of all friends who are within his/her vicinity giving their distance, time it will take to reach them, a detailed route plan and a good restaurant where they could meet and have coffee after a days long work.

3.2.5. *The Services Layer.* In this layer, loosely coupled (but interoperable) security services required for consumption by apps in the application layer are available. These services should be designed in such a way that they can be consumed, depending on the needs of the user within the context of selecting an appropriate business app. For example it will not make sense for a security service such as authentication to be used by the myDevice app in that it is supporting an application as opposed to a device. Therefore, although services in this layer are loosely coupled, interoperability will only be applicable amongst those services that are deemed to be interoperable according to security standards at both hardware and software levels. We briefly describe a selected number of services in this layer along the following lines.

(1) 3rd Party Security API (Gateway) This service is a specialized service required for interaction with a 3rd party application. In order to facilitate integration with other third party applications foreign to the platform, a range of API would be necessary for facilitating such interactions. This justifies the presence of this specialized service that is made available for ensuring that there is no breach in the security of the application being invoked outside the platform. Equally, third party applications would also want to prevent any security breach in their system, hence the importance of this specialized service in the platform. Furthermore, the monitoring of operation in the platform might be necessary if such a service has not yet been implemented in the platform. This is equally true when some forensic tools are required for enforcing the security in the platform should a breach arise. Hence the importance of the 3rd party gateway that is used for interaction with other security services not yet available in the platform.

(2) Trust The Trust service is an important security service that has the function of providing a truster with information that will assist in determining whether the trustee can be trusted. Such information would be derived from past experiences as well as the trustees reputation according the general public or community, based on the trust model used. This service would make use of existing data, feedback from users, 3rd party data, and social position to determine the trustworthiness of stakeholders in the iYP ecosystem. This Trust service would be consumed by apps such as MyBusiness and MySocialNetwork, where personal relationships and business collaborations are established. For example, a user of the BiYP application, wrapped with the MyBusiness security app, would like to collaborate with another user, for business purposes. The collaboration could be in the form of bulk-buying, selling products on credit, or customer recommendations. The MyBusiness app would be able to provide trust ratings and recommendations to the parties involved through the use of the Trust service, therefore facilitating trustworthy collaborations. Similarly, the MySocialNetwork security app could consume the Trust services functionality to facilitate trusting personal relationships.

3.2.6. *The 3rd Party Layer.* This layer is the extension of the API services gateways mentioned above. But the layer is physically situated in the service providers environment (e.g. suppliers, banks, etc.) and has the purpose of uploading and/or receiving information from the specialized API services in the third layer, for further processing. In summary, the service model forms the basis for the decision of an appropriate generic cloud-based technical infrastructure required for implementing and deploying the SiYP platform. However, discussions on the architectural model are beyond the scope of this paper.

4. CONCLUSION AND FUTURE WORK

The Security-in-Your-Pocket (SiYP) platform is innovative in the sense that it provides a user-centric approach towards the personalization of security services with specific reference to mobility within cloud infrastructures. SiYP, as discussed in this paper, also highlights the acute need for new security services that will support apps such as MyBusiness, MyDevice, and the like. Interesting examples of such services, amongst others, include a Transparency service and a Trust service. Future work will focus on the construction and usability aspects of the proposed SiYP platform.

References

[1] M.A. Bihina Bella, M.S. Olivier, J.H.P. Eloff. A fraud management system architecture for next-generation networks . Forensic Science International 2009; 185: 51-58.

[2] J. Wang, Y. Zhao, S. Jiang, J. Le. Providing Privacy Preserving in cloud computing 2009, pp. 213-216.

[3] F. Cuppens, N. Cuppens-Boulahia, M.B. Ghorbel. High level conflict management strategies in advanced access control models 2006, Vol. 186, pp. 3-26.

[4] K. Padayachee. An aspect-oriented approach towards enhancing Optimistic Access control with Usage Control 2010.

[5] K. Ruan, J. Carthy, T. Kechadi, M. Crosbie. Cloud Forensics. Advances in Digital Forensics VII 2011, IFIP AICT 361 pp. 35-46.

[6] C.A. Christiansen. Innovation and Security: Collaborative or Combative 2008, International Data Corporation (IDC) White Paper.

[7] S. Roth. New for whom? Initial images from the social dimension of innovation. International Journal of Innovation and Sustainable Development 2009, Vol. 4, pp. 231-252.

[8] J.A. Schumpeter. Capitalism, Socialism, and Democracy . Routedge, 1943.

[9] ISO 7498-2. Information processing systems Open systems Interconnection Basic Reference Model Part 2: Security architecture 1989.

[10] J. Chetty, M. Coetzee. Information Security for Service Oriented Computing: Ally or Antagonist 2011, pp. 460-465.

[11] Y. Huang, X. Ma, D. Li. Research and Application of Enterprise Search Based on Database Security Services 2010, pp. 238-241.

[12] The Oxford English Dictionary 2012.

[13] T. OReilly. What Is Web 2.0, Design Patterns and Business Models for the Next Generation of Software 2007.

[14] A. Jsang, R. Ismail, C. Boyd. A Survey of Trust and Reputation Systems for Online Service Provision. Decision Support Systems 2007, vol. 34, pp. 618-644.

[15] C. Weinhardt, A. Anandasivam, B. Blau, J. Stosser. Business Models in the Service World. IT Professional 2009, vol. 11, pp. 28-33.

[16] M. Masnick. Innovation In Security: It's All About Trust 2011.

[17] R. K. Balachandra, P.V. Ramakrishna, A. Rakshit. Cloud Security Issues 2009, pp. 517-520.

[18] N. Leavitt. Mobile Security: Finally a Serious Problem?. Computer 2011, vol. 44, pp. 11-14.

[19] X. Lin. Survey on cloud based mobile security and a new framework for improvement 2011, pp. 710-715.

[20] J. Oberheide, F. Jahanian. When mobile is harder than fixed (and vice versa): demystifying security challenges in mobile environments, 2010, pp. 43-48.

[21] H. Bullingera, K. Fhnrichb, T. Meirena. Service engineeringmethodical development of new service products. International Journal of Production Economics 2003, vol. 85, pp. 275-287.

[22] M.A. Bihina Bella, M. S. Olivier, J. H. P. Eloff. Near miss Detection for Software Failure Prevention (In press).

[23] Cloud Security Alliance. SecaaS: defining categories of Services 2011.

[24] M. T. Dlamini, H. S. Venter, J. H. P. Eloff, Y. Mitha. Authentication in the Cloud: A Risk-based Approach (In press).

Acknowledgement 1. *The support of SAP Research Pretoria and Meraka CSIR towards this research is hereby acknowledged. Opinions expressed and conclusions arrived at are those of the authors and not necessarily to be attributed to the companies mentioned in this acknowledgement.*

(J.H.P. Eloff, M.A. Bihina Bella, D. Isherwood, M. Dlamini, E.K. Ngassam)
SAP Research,
Internet Application and Services Africa,
Pretoria, South Africa
E-mail address, J.H.P. Eloff: jan.eloff@sap.com

(Mariki M. Eloff)
Institute of Corporate Citizenship,
University of South Africa,
South Africa
E-mail address, M.M. Eloff: eloffmm@unisa.ac.za

(M.A. Bihina Bella, M. Dlamini)
Faculty of ICT,
Tshwane University of Technology,
Pretoria, South Africa
E-mail address, M.A. Bihina Bella: madeleine.bihina.bella@sap.com

(D. Isherwood)
Academy of Computer Science and Software Engineering,
University of Johannesburg,
South Africa
E-mail address, D. Isherwood: donovan.isherwood@sap.com

(M. Dlamini)
Department of Computer Science,
University of Pretoria,
South Africa
E-mail address, M. Dlamini: moses.dlamini@sap.com

(E.K. Ngassam)
CSIR, Meraka, Pretoria, South Africa,
Faculty of ICT, Tshwane University of Technology,
Pretoria, South Africa
E-mail address, E.K. Ngassam: ernest.ngassam@sap.com

STANDING YOUR GROUND: CURRENT AND FUTURE CHALLENGES IN CYBER DEFENSE

BARRY V.W. IRWIN

Abstract. This paper explores the challenges facing those involved in cyber defense at an national, organisational and individual level. As the global economy grows more dependent on the Internet and connected infrastructure so the risk and impact of attack grows. A long standing response to attacks of various kinds conducted on the Internet has been to filter traffic, but not to respond to these. In some cases reactive action is taken, but even where attribution is possible, prosecution is rare. In recent months several countries have stated their policy of military response where they feel that their national infrastructure is threatened. The risk to organizations, civilian populations, and individuals is discussed in the case of such militant response or retaliation. The paper further considerations aspects such as reputation, neutrality, and the concept of Internet 'kill switches'.

1. Introduction

The realisation of cyber warfare, and cyber attacks has increased since the significant events of May 2007, in what has become termed the Estonian Incident. Denial of Service attacks affecting nation states have subsequently occurred on a number of occasions. In recent months the discovery of the Flame malware has raised more questions in terms of the offensive capability of nation states, with this software exhibiting previously unprecedented sophistication. At the opposing end of the spectrum individuals and organisations often bear the brunt of malicious online activity. This paper considers the challenges that emerging Internet states need to consider as part of their future development and planning. An integral part of this is to consider the approach taken when individuals and organisations detect and need to decide on an appropriate response to such attacks.

Section 2 briefly explores the evolution of Security on the Internet, considering the substantial changes that have taken place. This is followed in Section 3 by a discussion of current challenges faced in cyber defense. The future challenges from the perspective of the author are explored in Section 4.

2. Evolution of Security

The Internet has transitioned from its largely academic roots, where the principles of openness and data sharing were paramount to a global network, vital to global commerce and communication. Possibly as a factor of the rapid growth, legal systems globally have failed to keep up, and the ease with which actions may be performed on or against system located in areas geographically distant to the instigator, have made the enforcement of traditional laws difficult. The question of

Key words and phrases. Trust, Device Comfort, Foreground Trust, Human Computer Interaction.

legal jurisdiction in such cross border incidents greatly complicates the resolution from a law enforcement perspective. The long standing principle employed by many network administrators has been to *protect the border* and implement appropriate technologies to manage the communications of what has become an increasingly hostile, exterior network to those hosts within their organisation. This thin veil of security, aptly described by Bill Cheswick as “a sort of crunchy shell around a soft, chewy center” [3], has probably never been more apt. The initial approach taken with the widespread adoption of firewall systems was to block the *known bad* (blacklisting), which has transitioned to largely block everything except *what we need* (white listing). In both approaches, the majority of organisations still focus their defensive efforts outwards. The act of blocking itself takes many forms. Considering the case of the traditional IP Firewall, a block is often differentiated between two states:

Deny: in which IP datagrams matching the rule are discarded, and

Block: in which prior to the discard a protocol relevant error message is sent, indicating this has occurred.

In the case of blocking, the RFC specifications for the ICMP protocol provide for fairly specific signalling as to why a packet may have been dropped [9] (which may not have been only due to security reasons). A similar approach is taken when dealing with email, where spam is either discarded (preferably pre-acceptance) or a notification of the message having been quarantined or tagged is sent.

In recent years many organisations have started to consider the implications of focusing defensive resources inwards on their networks. While in many ways this has been primarily driven by technologies such as DLP (Data Loss/Leak Prevention), a secondary driver has been seen to be the mitigation of risk and liability.Considering the scenario of an administrator detecting repeated port scanning activity from a source, the response has usually consisted of a number of distinct phases.

Atribution: The source(s) are identified, initially as an IP address, and then resolved using various methods to a source organisation. This organisation may be an endpoint, or in the case of consumer Internet access, an Internet service provider.

Action: The identified sources are blocked, usually at the firewall or routing infrastructure at the point where the administrator has control.

Complaint: A notification of the malicious activity is sent to the responsible organisation, usually as an email directed to the abuse@ email address, which is required to exist as per RFC 2142 [5]. This is not always successful, as the addresses often do not exist[1] or no response is received. An escalation path is typically followed with the provider of network access to that organisation.

In severe cases such as Denial of Service (DoS) attacks an organisation may need to liaise with their upstream network provider(s) in order to perform mitigation actions as close to the ingress point of the traffic as possible. In some circumstances, Network Providers may de-announce targeted network address blocks, resulting in all traffic destined for these ranges of addresses being dropped. An alternate approach is selective null-routing of the traffic on specific network devices, which causes it to be dropped. In either of these scenarios, the targeted organisation can end up suffering a more substantial outage than was initially caused by an attack.

[1] A list of sites not complying to the norm is maintained at http://www.rfc-ignorant.org/

This raises the question of responsibility. Network Providers usually have clauses dealing with disconnection or dropping of traffic built into their operation agreements and Acceptable use policies (AUP) with their clients. When these are considered, what has traditionally been dealt with as a technical network operational issue can now be seen to be a potentially significant element in organisational risk management. One of the most prevalent forms of abuse complaints is the *Cease and Desist* type, most commonly filed by firms acting on behalf of media companies and groups such as the MPAA[2] and RIAA[3] against claims copyright violators. This is a request that the recipient cease form a particular action or potentially face legal action. The validity of these claims in areas outside of the jurisdiction of the organisations issuing them has become a hotly debated topic. While much of the legal groundwork and debates has been around piracy and copyright violation, the lessons learned can be well applied to shutting down malicious activity form a cyber defense point of view.

One of the most contentious topics raised within the Information Security community in recent years, is that of strike-back or aggressive self defence of ones assets against an aggressor. Parallels have been dawn between existing law dealing with one right to defend ones own property and person in the physical world, but no definitive legislation has been passed in this area clearly defining the so-called rules of engagement. This has been variously described as going a step beyond just blocking traffic, but rather attempting to shut down the source of the malicious activity. To a large extent this is still a legal grey-area. Legitimate take downs have been conducted, but these have consumed substantial resources, and been executed under judicial warrants. Should individuals, organisations, and nations have the right, and ability, to defend their networks and systems from attack?

3. Current Cyber Defense Issues

In recent months a number of states the USA [1], but also notably Russia and China [10], have stated that they will respond with physical force using their military to perceived major threats against their cyber infrastructures. Some thought into this this possibility by Waxman [15], who highlights the point that Military attacks are considered illegal under international law, with exceptions for self-defense or when authorized by the U.N. Security Council. While military action in response to an electronic threat may seem extreme, it provides a useful entry into the discussion of cyber defense and the problems being faced by a range of partied online today.

This section opens with a consideration of the context in which the three primary classes of entity on the internet experience cyber defense. This is followed by a brief discussion around the problem of attribution, incident response, and ongoing defensive measures to support operational needs.

3.1. **Context.** The context in which current cyber defense issues needs to be considered can be broken down into three broad groups. These groups are important to consider both as groupings of increasing vulnerability, but also as groups which can have positive impacts on the overall state of online safety.

[2] Motion Picture Association of America. http://www.mpaa.org/

[3] Recording Industry Association of America. http://www.riaa.com/

3.1.1. *Individual.* Individual users represent the most significant portion of the Internet user base, by some orders of magnitude, yet have a disproportionally insignificant ability to influence activities and policy. Conversely it is this body that has the most opportunity to influence the realities of the day-to-day running of the Internet. This group of users is also as a whole the most exposed to, and under resourced with regards to skills, knowledge and finance to deal with malicious activity. This context may also be vulnerable to either recruitment or co-option into performing malicious activity, whether motivated by political[4], socio-economic[5], or criminal ideals[6].

3.1.2. *Organisational.* Organisations range in size from SME operations to multinational organisation, with operational capacities exceeding some nation states. At the smaller end of this spectrum, organisations are often unaware that they have been attacked or are party to an attack on others. An example of this is evidenced by the prevalence of phishing activity hosted on the websites of these organisations, which often goes undetected until reported to the hosting provider. Unless staff are highly motivated, incidents are often ignored or go unreported. Larger organisations have the resources, from a legal, technical and financial standpoint to be able to deal with substantial activity, and are often compelled to given the risk of damage to brand and/or reputation in the event of incidents, such as data disclosure, defacements or outages.

These bodies consist of individuals working towards a common goal, but at the same time the organisation, particularly at the larger end of the spectrum, is seen as a faceless entity. Employee dissatisfaction or disillusionment, may well account for the significant proportion of incidents reported by organisations originating from inside the network. Employees may also neglect their responsibilities for cyber security as it may be seen as the organisations problem.

3.1.3. *National.* The interests of a nation state are the most significant. A nation state may choose to act when the situation with a particular incident targeting the state as a whole, or significant elements such as critical infrastructure or financial services becomes critical.

This decision may be due to a (perceived) threat to National infrastructure, and could be regarded as valid in terms of self-defense under International Law. This could be Internet infrastructure residing within the national boundaries, or other so-called critical infrastructure such as electrical power, water or general telecommunications. The resources that a nation state are able to bear are substantial, ranging from seizing administrative control of network ingress points to direct military intervention. Given these resources, consideration must be given as to what an appropriate response to a given threat is.

3.2. **Attribution.** One of the biggest challenges facing cyber security professionals is that of attribution. This is a well established problem, where systems involved in cybercrime, or other malicious activity, are rarely those which can be attributed to the perpetrators of such actions. This is further complicated by situations where

[4] Such as the numerous attacks during the Arab Spring unrest in Early 2011

[5] This would include groups such as the Occupy movement, and those targeting financial institutions as part of the response to the 2008 financial crisis.

[6] Prime examples of this would include the hackivist vigilante groups of Anonymous and LulzSec and their various offshoots.

IP addresses are shared due to technologies such as network address translation (NAT) or Proxy servers. On the case of the former, its is often near impossible to apportion responsibility to a party behind the NAT gateway.

While many parties involved in Information Security are primarily concerned with the attribution of responsibility. These same parties are often hard pressed to be able to perform the attribution process within their own organisation when receiving complaints from other parties. A particular area of concern is wireless networks, where there is a history of malicious activity being performed using either open wireless networks or having intruded on the networks. The same can be said of shared computing resources such as Internet café's, computer laboratories or consoles at schools or universities. Legislation such as South Africa's Electronic Communications and Transactions Act [12], and similar legislation elsewhere, places a burden of record keeping on organisations for attribution. Recent media attention has focussed on the anti-piracy legislation enacted in France and the United Kingdom, where Internet Service Providers have been tasked to attribute claims of copyright infringement to their users, and act against repeat offenders.

A larger problem exists with the rise of the use of networks of compromised systems in order to execute nefarious activities online, as explored by Clarke and Landau [4]. The question of where the attribution should lie in the event of an individuals system being part of a botnet denial of service attack needs to be considered. Should the individual be held responsible (when the software is almost certainly installed on their system unbeknownst to them)? Alternately, should the parties running the botnet be held liable? What about the party who rented the botnet in order to execute the attack? Current legal frameworks have extreme difficulty when dealing with issues such as this.

Blind filtering of traffic, or any kind of aggressive or intrusive response to these kind of hosts is likely to do more harm than good, in terms of reputational damage to the organisation perpetuating it, and to the often innocent and unaware targets [11]. Consider the extreme of where an actor makes use of a botnet either of its own creation or obtained via rental of compromised systems to attack another party. Attribution efforts are likely to lead to the endpoints, who may be guilty of nothing more than practising poor digital hygiene. It may be near on impossible to locate the operators or executors of the attack.

3.3. **Response.** In an ideal world perpetrators of malicious acts would be brought to book. Given the complications of attribution many organisations cease their response at the point where services have been restored and the threat mitigated. Follow-up and further investigation is usually only warranted in cases where substantial financial loss, or other damage has occurred. One of the areas where a positive response has been seen is in the areas of combating phishing and spamming operations. When dealing with these, the majority of providers acting fairly swiftly in the take-down of these operations.

Pursuing further actions in response to an incidents, particularly those taking place cross-border, requires both substantial motivation, the involvement of law enforcement, and financial resources. The high profile cases involving the shut down of the Maraposa [13] and DNSchanger [6] botnets are examples of where this has been successful.

3.4. **Operational Defence.** Organisms as, part of there ongoing biological process, process stimuli and produce a response. Considering malicious activity as a stimulus the idea that needs to be discussed is what an appropriate response to this stimulus is? The first point to consider with regard to cyber defense is as to whether the stimulus is recognised, with much malicious activity going undetected for extended periods of time. The chances of this occurring in a timely manner increase with the relative 'size' of the target, but other factors also influence this. Once recognised, the choice as to whether to respond needs to be made. Assuming a response is decided on there are a number of options which can be followed. The gamut of options runs form ignoring the incident as nuisance traffic, to actively filtering, to engaging with law enforcement and pursuing legal avenues at the other.

From an operational perspective the primary concern is to return the organisation to normality. In practical terms, in the case of malicious traffic coming over the network, this means blocking incoming network traffic, thereby mitigating the threat. This can be done at either their own boundary, or in conjunction with upstream providers.

From this point complaints and notification of the activity can be passed top parties identified by the attribution process. Remedial action is generally likely to occur other than in cases involving phishing and spamming, which have a raise awareness level, and are generally recognised to be criminal activities. The general understanding of a port-scan or brute for attempt against a service is less clear cut.

4. Future Challenges

Considering the issues discussed above, five distinct areas deemed to pose particular challenges in terms of ongoing and future cyber defense. These are discussed below.

4.1. **Aggression.** The first of these areas to consider is what constitutes aggression. Being able to quantify the degree of the attack is key in being able to then respond with appropriate force (being one of the tenets of most common understandings of the concept of self-defense in the physical world) [7]. Two of the most common activities observed by hosts on the Internet are generic portscanning and bruteforcing of common services - particularly the Secure Shell (SSH). While the former can be seen as the physical equivalent of checking if a door is unlocked, the latter is an active attempt to achieve an intrusion on a system. While portscanning is largely an annoyance it could be potentially part of a more sophisticated attack. The counterargument to this is that a sophisticated targeted attack is likely to be more subtle, and hence less likely to be detected. Common responses to the bruteforcing issue (also increasingly common with VoIP services) is to filter the source host after a specified number of connections.

The matter that needs to be considered is how should one respond to threats on a larger scale. Waxman [15] discusses the concept of both for and its appropriateness in the context of the United Nations Charter.

4.2. **Reputation.** The reputation of parties transacting on the Internet is especially important. In the event of organisations or net-blocks persisting with malicious behaviour. They could become subject to filtering or blackhole routing by other parties online. If the problem is severe enough upstream providers may need to resort in termination of their uplink.

While this approach may be appropriate for end users or organisations on the 'edge' of the network, a number of problems become apparent when one considers the impact that such accounts could have if applied higher up the connectivity chain. The potential exists for organisations and individuals to be significantly impacted should an ISP servicing them consistently fail to respond to malicious traffic exiting their network. This has already been seen in the realm of anti-spam solutions where large blocks of IP address space have been tagged as spam sources and hence had e-mail emanating from them tagged as spam, or rejected. Countries seen as havens for criminal activity could face sanction from network providers,and other states, particularly those with the capability to provide network transit. In the past, large portions of Korean and Nigerian address space have been filtered by US and European providers, due to spam and scam emails emanating from these countries. African nations are particularly at risk given the recent influx of Internet connectivity, and the relatively low skill levels and awareness around cyber security.

4.3. **Kill-switches.** Following from the discussion above, consideration needs to be given to the issue of does one county have the right to kill traffic to other countries, particularly when providing transit for links from submarine cables to landlocked neighbours. An example of a situation where this could have been considered was during the cyber attacks on Georgia, related to the geostrategic conflict involving South Ossetia in 2008 [8]. Much of this traffic was routed though terrestrial fibre-optic links running though Turkey. The net result was that Turkish internet users were negatively affected. Related to this is the concept of proverbial kill-switches, which have been proposed by a number of governments with the intention of being able to isolate the countries network in the event of a major threat. The practicalities of implementing such a system aside, the number of parties potentially affected should such a system be realised. When one considers highly connected countries such as the Netherlands, United Kingdom, and USA and the impact on the global network should they theoretically sever their communication paths, the negative impact and collateral damage against innocent parties both internal and (possibly unwittingly) using these Internet Exchange points for transit would be significant.

4.4. **Neutrality & Mutual Aid.** The concepts of a neutral state and military mutual-aid need to be reassessed as the world becomes increasingly connected. Consider the case of a country who has declared its neutrality in a dispute between two others. Traffic for one of the disaffected parties transits though the neutral state. How would interfering with the traffic in transit be interpreted.

Related to this is the concept of traditional military mutual aid agreements, such as the NATO bloc. How will these treaties be bought to bear in the future should member states be threatened not by physical hostilities, but rather against their telecommunications networks and other critical infrastructure? This issues were raised within NATO following the Estonian incidents of 2007, and revisited in the aftermath of the Attacks on Georgia [14].

4.5. **International co-operation.** As legal systems mature, and law enforcement skills are developed we are likely to see increased collaboration in dealing with malicious activity. This has to some extent been seen with takedown operations against spam operations and botnet control nodes; although these have been largely driven by the private sector, with Microsoft in particular having taken a leading role.

5. Conclusion

This paper has presented a number of challenges faced by information security practitioners. These areas are not clear cut and will require substantial discussion and debate in the near future. The debate around the impact of these identified areas will need to be considered at all levels. Each of the identified challenges will impact differently for the varying classes of consumers.

The evolution of the global network is likely to consider and become even more pervasive, and critical to our social and economic lives, as individuals, organisations and nations. Part of this evolution needs to be in terms of ensuring its ongoing integrity, from technical, economic, and legal perspectives. Probably the single most important challenge facing all stakeholders within cyber defense is that of raising the general awareness at grassroots level. If individuals take responsibility for the security of endpoint systems, a reduction in rates of compromise will follow, along with a drop in the malicious activity emanating from such systems. The action does not need to be complex, but rather ensuring good practice. Vendors such as Microsoft, Adobe, Google, and Mozilla have made great strides in increasing the ease with which software updates are obtained. Much as a motor vehicle requires checking of the wear of tyres and functioning of lights and indicators in order to be considered roadworthy, possibly the same should apply to endpoint systems?

References

[1] D. Alexander. U.S. reserves right to meet cyber attack with force. Online, November 2011. Last Accessed: 2012-06-30.

[2] S.W. Brenner and Leo L Clarke. Civilians in cyberwarfare: Conscripts. Vanderbilt Journal of Transnational Law, 43(4):1011-1076, October 2010.

[3] B. Cheswick. The design of a secure internet gateway. In Proc. Summer USENIX Conference, pages 233-237, 1990.

[4] D.D. Clark and S. Landau. The problem isn't attribution: it's multi-stage attacks. In Proceedings of the Re-Architecting the Internet Workshop, ReARCH '10, pages 11:1-11:6, New York, NY, USA, 2010. ACM.

[5] D. Crocker. Mailbox Names for Common Services, Roles and Functions. RFC 2142 (Proposed Standard), May 1997.

[6] Federal Bureau of Investigation. International cyber ring that infected millions of computers dismantled, November 2011. Last Accessed: 2012-06-24.

[7] M. Hoisington. Cyberwarfare and the use of force giving rise to the right of self-defense. Boston College International and Comparative Law Review 32, 32(1):439, Winter 2009.

[8] D. Hollis. Cyberwar case study: Georgia 2008. Small wars Journal, January 2011.

[9] Internet Assigned Numbers Authority (IANA). ICMP type numbers. [online], 13 February 2008. Accessed 2008-12-12.

[10] G. Patterson Manson. Cyberwar: The united states and china prepare for the next generation of conflict. Comparative Strategy, 30(2):121-133, 2011.

[11] J. McMahan. Self-defense and the problem of the innocent attacker. Ethics, 104(2):252-290, January 1994.

[12] Republic of South Africa. Electronic communications and transactions act, 2002 no. 25 of 2002. Online, July 2002.

[13] Matt Thompson. Mariposa botnet analysis. Technical report, Defence Intelligence., October 2010. Last accessed: 2012-06-30.

[14] E. Tikk, K. Kaska, K. Rnnimeri, M. Kert, A.-M. Talihrm, and L. Vihul. Cyber attacks against georgia: Legal lessons identified. Online, November 2008.

[15] M.C. Waxman. Cyber-Attacks and the Use of Force: Back to the Future of Article 2(4). Yale Journal of international Law, 36(2), 2010.

(B. Irwin)
Security and Networks Research Group (SNRG),
Department of Computer Science,
Rhodes University, South Africa
E-mail address, B.V.W. Irwin: `b.irwin@ru.ac.za`
URL: `http://www.cs.ru.ac.za/`

IS IT PRIVACY OR IS IT ACCESS CONTROL?

SYLVIA L. OSBORN

Abstract. With the widespread use of on-line systems, there is an increasing focus on maintaining the privacy of individuals and information about them. This is often referred to as a need for privacy protection. I will briefly examine definitions of privacy in this context, roughly delineating between keeping facts private and statistical privacy which deals with what can be inferred from data sets. Many of the mechanisms used to implement what is commonly thought of as access control are the same ones used to protect privacy. This presentation will explore when this is not the case, and in general the interplay between privacy and access control on the one hand, and, on the other hand, the separation of these models from mechanisms for their implementation.

1. Introduction

The right to privacy is enshrined in international and national covenants and charters on human rights. Concern for the privacy of on-line data began with the introduction of computing systems. By 1980 the OECD published its guide- lines dealing with the privacy of information and transborder ow of information [8]. In the database community, the Hippocratic database paper is considered the seminal paper in introducing privacy concerns to the database community [2]. Meanwhile, access control has always been a part of computer systems.

We begin by examining definitions and dimensions of privacy preservation, continue with an introduction to Sandhu's OM-AM framework, consider the available mechanisms for implementing access-related models, and then comment on how all these ideas together. We also briefly discuss the user. Our hope is that if there are gaps in the effective protection of information, this analysis might help to show where the gaps are.

2. Privacy vs Access Control in Computer Systems

In this section, we review some definitions of access control and privacy, in order to crystalize their similarities and differences. Because the discussion of access control is shorter, we proceed with it first, followed by some definitions of privacy, and finally highlight their similarities and differences.

2.1. **Access Control.**

Access control deals with controlling who has what kind of access to various resources. The resources can be physical (i.e. a computer system) or strictly deal with data. The data can describe documents, inventory, shipping requisitions for a large company, allocation of university courses to classrooms, the destination of an aircraft carrier, etc. In other words, although a lot of data concerns individuals, there is also a lot of other data dealing with other things. There are three well-known access control models. In the first, Discretionary Access Control (DAC), data is owned by the individual computer user (e.g. personal files in Unix);

in Mandatory Access Control (MAC), control is centralized and it is assumed that the enterprise owns (and labels) all the data. The third is Role-based Access Control (RBAC), where permissions are grouped into roles and roles are assigned as a unit to users. RBAC has been shown to be able to simulate both MAC and DAC [10].

The basic components of an RBAC system are users (U) or subjects, permissions (P) which are pairs ($o; a$) where o represents an object to be protected and a, an access mode on this object. Roles (R) consist of a set of permissions, represented by a permission-role assignment (PRA). Users' membership in roles is represented by a user-role assignment (URA). Roles can be arranged in a hierarchy such that a senior role inherits the permissions of its junior(s), and members of a senior role are also members of its juniors.

2.2. **Privacy.**

Privacy, on the other hand, typically infers that the data in question relates to human beings, or possibly to corporations. It is related to the right to privacy which is enshrined in international and national covenants and charters on human rights. The Merriam-webster dictionary defines privacy as "freedom from unauthorized intrusion" [13]. The classic version of the Hippocratic oath contains the following[1]:

> *What I may see or hear in the course of the treatment or even outside of the treatment in regard to the life of men, which on no account one must spread abroad, I will keep to myself, holding such things shameful to be spoken about.*

A definition given in a previous ISSA paper [10] is:

> *Privacy is the faculty and right that a person has to define, preserve and control the boundaries that limit the extent to which the rest of society can interact with or intrude upon. At the same time, he or she retains full control over information generated by, and related to, him or her.*

Here we begin to see one of the issues: when a data provider gives their information to, say, a company with whom they do business, they no longer have direct control over the data. The question of ownership, if that is a term we want to use, becomes clouded.

An interesting examination of the dimensions concerning privacy from a technical point of view has been given by Barker et al. [3]. They discuss four orthogonal aspects of data privacy, three of which are shown in Figure 1. Following Purpose along the x-axis, privacy protection decreases (purpose becomes more general) as one moves further from the origin. The first point refers to data given to a service for a Single use. Next comes Reuse Same, which allows multiple uses of the provided data for the original purpose. The third point, Reuse Selected, represents multiple uses of the data by the data collector for related purposes, e.g. in a medical situation, the information is provided to the health care professional for medical reasons, and some of it is released to the insurer. The Reuse Any point allows the data to be used in the future for unforeseen purposes, such as for example, medical research. Finally the Any point allows any reuse of the data, and should probably not be encouraged.

The y-axis in Figure 1 deals with visibility, i.e. who can see the data. The first non-origin point is labelled *Owner*. In general this would refer to the data provider whose information is being discussed. It might also refer to data stored in some

[1] From www.medterms.com

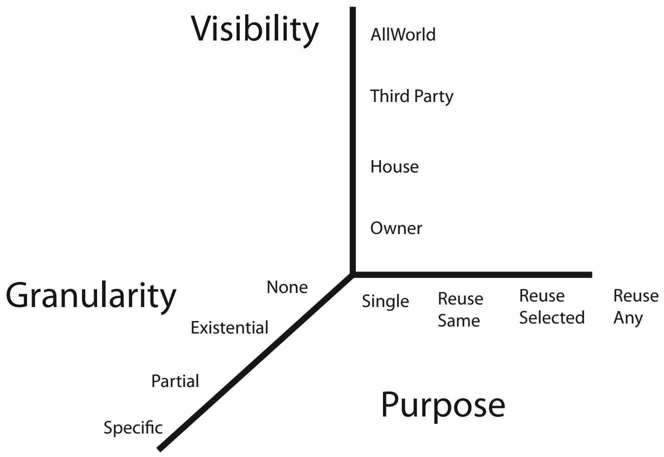

FIGURE 1. Data Privacy Taxonomy from [3]

cloud service to which the cloud-providing company has no access only the person storing data there has access. The next point is House, where the entity storing the data, which might be a company with whom the data provider is doing business, or a company like Google who might use some of the data to provide more/better service, may assume ownership and treat the data in some way not immediately obvious to the data provider. Third Party refers to data users who are authorized by the house to see the data, usually under some agreement with the house concerning the nature of the use of the data. The final point, All World, refers to data which is publicly exposed for anyone with access. An example would be information posted on a publicly available web page.

The z-axis shows different granularities with which information can be released. Moving out from the origin, just the Existence of something may be revealed; the next point, Partial would be the release of a range, like $100,000 - 150,000 for salary; and the least privacy occurs with the release of a Specific piece of information in its raw form.

The 4th dimension given in [3] is Retention. To comply with various rules and legislation, data may have an expiry date. The granularity axis in the previous discussion leads us to another aspect of privacy. Discrete facts or raw data can be revealed, as would be the case for the Specific point in the taxonomy. Or to turn it around, access to the individual attribute values of a person can be protected as an example of privacy protection.

Often the databases which contain these facts are also to be used for statistical queries. Mechanisms for hiding individuals' information for these statistical applications include k-anonymity [12] and differential privacy [6]. Informally, k-anonymity means that in data over which statistical queries are run, each individual is indistinguishable from $k - 1$ others, even if the released data is combined with publicly available information. Differential privacy describes a statistical data- base release in which the absence or presence of a single element will not affect the output of a query in a significant way.

An enhancement of treatment of data with a purpose for access has been described in [4]. They introduce the idea of a purpose tree, where a more general purpose (e.g., marketing) appears as an ancestor of a more specific purpose (e.g., telemarketing) . They also distinguish between intended purpose and access purpose: the requester of access to private data must provide an access purpose, and the provider of the private data gives an intended purpose for which they are willing to allow access. E.g. a data provider can allow their address to be used for the shipping purpose but not for marketing. Propagation in the purpose tree takes place, so that an access request with a specific purpose is allowed if the data provider has indicated a more general (an ancestor in the purpose tree) intended purpose.

In [1], we investigated the integration of such purposes with RBAC. We showed that purpose becomes an additional component of some permissions. It does not need to be present for all data access, e.g. for inventory it would not be necessary but it would for access to the customer's address. In RBAC, such permissions are inherited by senior roles, so that it somewhat muddies the model. Role activation needs to be accompanied by an access purpose, and has to filter out permissions if their intended purpose does not comply with the access purpose for a given role activation.

2.3. **Similarities and Differences.** From the preceding discussion it should be clear that access control in its traditional form is not adequate to protect privacy. At least two similarities are:

(1) Facts (attribute values) are to be made available (read) or not to other users
(2) (similarity with DAC) the decisions about what is visible to whom are made by individuals, not by some security administrator The differences include the following:
(3) when privacy is of concern, the resource being protected is data concerning an individual or an enterprise.
(4) there may be an additional issue of the purpose for the access which can be part of the privacy discussion
(5) privacy becomes an obligation when the data provider (the person, let's say, who is being described by the data) is no longer the sole owner of the information, but has given the data to another organization whose obligation it now is to protect their privacy.
(6) something not yet discussed: There are other issues dealing with data which the data provider might never have owned (or at least was not aware they owned), i.e. are not strictly related to data explicitly given by the data provider, such as tracking clicks on a web page, using the GPS on a mobile device to track location, and using face recognition software for tagging photos.

3. The OM-AM model

In a 2000 paper, Ravi Sandhu introduced the OM-AM model [11]. We will briefly outline this framework, and then discuss mechanisms which have been used and are available for implementing privacy protection and access control.

3.1. **The Model.** Software engineering practice tells us that to implement a system, we should start with more abstract requirements and move towards implementation. The OM-AM model highlights this distinction as it applies to access control systems. The Objective and Model (the OM) describe what the requirements are, and the Architecture and Mechanism (the AM) give the how for application of the requirements.

In the previous section, we have studied the OM part for privacy concerns. The objective for privacy preservation can be briefly summarized as information hiding/sharing in a very discretionary fashion. The addition of purpose for use of individual attribute values complicates the model. We will now look very briefly at the AM - the architectures and the mechanisms.

3.2. **Architecture and Mechanisms.** It is possible to look at a separate architecture/mechanism for each component of a model, as Sandhu did in [11]. However, this discussion will be more general.

For Access Control in a centralized computer system, the architecture commonly used is that of a security kernel or reference monitor. The reference monitor interrupts and validates each access request. It may check an access matrix, or access control list to verify that the user can perform the requested operation. In a distributed environment, this checking can be performed by a service if a service-oriented architecture is being used.

One very general solution is the architecture and details of the XACML framework, in particular the data flow architecture given in [7]; this framework provides a set of standards (involving PDP, PEP, etc.) for implementing a broad range of policies whose details are expressed in XACML. The architecture is an example of a way of representing the reference monitor, accompanied by a set of standards. Here the access control model is expressed as a set of policies encoded in XACML. Examples of XACML policies for MAC, DAC and RBAC can be found.

To enforce privacy requirements, simple obligations by a "house", which stores private facts, to keep these facts secret can be expressed in XACML. This XACML framework is, thus, one example of a Architecture/Mechanism which can be used to enforce both access control and privacy requirements.

A wide range of cryptographic protocols are available, which are extensively discussed in [5]. The encryption/decryption is performed to validate the access, so the architecture is still a reference monitor, but the mechanism is encryption. We see from [5] that cryptographic mechanisms can be used to protect hierarchical arrangements of rights, so it is a valid mechanism to enforce an RBAC model, for example.

For Privacy of facts, the mechanisms seem to be the same as those used for standard access control. One exception is when an access purpose must be matched with an intended purpose, for which mechanisms need to be enhanced. For statistical privacy, special techniques are required. Statistical databases typically contain large amounts of data against which only summary queries are allowed. The objectives are expressed in complex definitions of k-anonymity or differential privacy, and the architecture for implementing these objectives is usually a specialized database system.

4. The User

To add another dimension to this discussion, we need to look at users. In the database textbooks, it is common to distinguish between "casual" or even "naive" users, and the database administrator who is supposed to understand all the details and implications of a complex data model. It clearly requires an expert to express access control requirements directly in XACML. Slightly less expertise is required for someone to design roles for an RBAC system, and have them automatically translated into XACML. Neither of these users would be a naive user, as they would require technical knowledge to write raw XACML, and at least thorough knowledge of the application environment to design an RBAC system. Statistical databases are not used by what could be called "casual" users; they are used by statisticians or researchers. Many repositories of facts concerning individuals are used and managed by these casual users. In a social network, it is a naive user who is specifying "privacy" settings (one could argue that what the user is doing is specifying access control to the data which they initially own). Any model provided to such a user needs to be one which allows them to understand the implications of what they are doing.

5. Discussion

As far as the Objective and Mechanism are concerned, we have highlighted some differences between maintaining privacy and enforcing access control. In both, there are scenarios where controlling reading of facts is the issue, and the same mechanisms can be used to achieve both objectives. Some models are only suitable for expert users. Others can be used by casual users. One could say that we have many adequate mechanisms, and if systems are not achieving their objectives, it might be that the objectives and models are not well understood by the user, or that the mechanisms are not being employed.

References

[1] A.L. Al-Harbi and S.L. Osborn. Mixing privacy with role-based access control. In Fourth International C* Conference on Computer Science & Software Engineering, C3S2E 2011, Montreal, Quebec, Canada, pp. 1-7, 2011.

[2] R. Agrawal, J. Kiernan, R. Srikant, and Y. Xu. Hippocratic databases. In VLDB 2002, Proceedings of 28th International Conference on Very Large Data Bases, pp. 143-154. Morgan Kaufmann, 2002.

[3] K. Barker, M. Askari, M. Banerjee, K. Ghazinour, B. Mackas, M. Majedi, S. Pun, and A. Williams. A data privacy taxonomy. In BNCOD, vol. 5588 of Lecture Notes in Computer Science, pp. 42-54. Springer, 2009.

[4] J.-W. Byun and N. Li. Purpose based access control for privacy protection in relational database systems. VLDB J., 17(4):603-619, 2008.

[5] A.V.D.M. Kayem, S.G. Akl, and P. Martin. Adaptive Cryptographic Access Control, vol.48 of Advances in Information Security. Springer, 2010.

[6] F. McSherry. Privacy integrated queries: an extensible platform for privacy-preserving data analysis. Commun. ACM, 53(9):89-97, 2010.

[7] OASIS. eXtensible Access Control Markup Language (XACML) ver. 2.0. http://docs.oasis-open.org/xacml/2.0/access_control-xacml-2.0-core-spec-os.pdf, 2005.

[8] OECD. Guidelines on the protection of privacy and transborder flows of personal data. http://www.oecd.org/document/18/0,3746,en_2649_34223_1815186_1_1_1_1,00.html, 1980.

[9] S.L. Osborn, R. Sandhu, and Q. Munawer. Conquering role-based access control to enforce mandatory and discretionary access control policies. ACM Trans. Information and System Security, 3(2):1-23, 2000.

[10] K. Renaud and D. Gàlvez-Cruz. Privacy: Aspects, definitions and a multi-faceted privacy preservation approach. In Hein S. Venter, Marijke Coetzee, and Marianne Loock, editors, Information Security South Africa Conference 2010. ISSA, Pretoria, South Africa, 2010.

[11] R.S. Sandhu. Engineering authority and trust in cyberspace: the om-am and rbac way. In ACM Workshop on Role-Based Access Control, pages 111-119, 2000.

[12] L. Sweeney. k-anonymity: A model for protecting privacy. International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems, 10(5):557-570, 2002.

[13] Merriam Webster. On-line dictionary. www.merriam-webster.com.

Acknowledgement 1. *The financial support of the Natural Sciences and Engineering Research Council of Canada is gratefully acknowledged.*

(S.L. Osborn)
Dept. of Computer Science,
The University of Western Ontario,
London, ON, Canada, N6A 5B7,

E-mail address, S.L. Osborn: sylvia@csd.uwo.ca

ON COMPLEX DIGITAL FORENSICS

MARTIN S. OLIVIER

Abstract. Science provides the basis for truth claims in forensics. Very little research has been done to explore the scientific basis of digital forensics. The work that has been done vary widely in what they propose; in most cases it is unclear how the philosophical remarks about such forensic science apply to digital forensics practice, or that the practical suggestions are a sufficient basis to claim that practice based on them is scientific.

This paper provides an initial exploration of the potential of decision problems from the field of algorithmics to form this scientific basis. There is no doubt that decision problems operate in the scientific domain and decision problems look similar to hypotheses to be of immediate practical use.

The paper suggests that, if decision problems are used in this manner, it is clear that current digital forensics have only scratched the surface of what is possible. Probabilistic complexity classes, for example, offer interesting possibilities for performing complex tests in relatively short times, with known error rates.

Using decision problems as a demarcation criterion makes it possible to distinguish between digital forensic science (or simply digital forensics) and digital forensic craft, that should rather be called digital investigative technique or some other suitable term that does not imply that its use leads to scientific truths.

1. Introduction

Forensics entails the use of science to determine matters of fact where such facts are required to settle disputes (for example, in courts of law) or to determine the root cause of an event of interest. Forensics employs the notion that scientific knowledge is true and hence a good basis to settle such disputes and/or determine causes. Digital forensics is that branch of forensics that studies evidence that exists is digital form.

In order to make such truth claims forensics has to be ‘scientific’. In some cases this is emphasised by using the term forensic science, which in this paper will be deemed to be synonymous with the term forensics. The notion of science (as well as the notion of truth) has been the subject of deep philosophical reflection over centuries; so much has been said that a paper that ultimately intends to deal with a small fraction of forensic science cannot hope to do justice to.

The obvious question then is what is the nature of digital forensic science or, with the same meaning, the science that underlies digital forensics? Cohen [5] is the only author who has provided a coherent answer to this question by describing an information physics ‘natural laws’ that apply to information and can be used as the basis for more complex truth claims. However, it is not yet clear that it is possible to always relate the behaviour of a complex system to truths about bits and related matters see, for example, Hofstadter’s argument [13] that a complex system may be more than the mere sum of its parts and may exhibit characteristics that are not present in the parts.

Key words and phrases. Digital forensics, forensic science, decision problems.

A recent newspaper story [15] provides some insight on what may go wrong if we rely on digital forensics that cannot be trusted it may negatively affect innocent people. However, simply discarding digital forensics because of a lack of trust turns the cyberworld into a safe haven for criminals who can exploit others without fear of being caught. Clearly a digital forensics is required that maximises the chances that the guilt of the guilty can be proven, and that will ideally never implicate an innocent party. If these requirements are met the inhabitants of cyberspace can proceed with trust even in those cases where the proactive security mechanisms fail. Note that this problem is not only present in digital forensics; other branches of forensics have also failed because they used junk science or pseudoscience [9, 18]. Regarding digital forensics, Caloyannides [3] boldly declares that “It is important for judges and juries to be highly sceptical of any claims by prosecution that digital ‘evidence’ proves anything at all.”

This paper will examine the suitability of algorithmics or algorithmic complexity theory to form the basis of digital forensics. The justification of positing algorithmics as this basis is deferred to later in the paper when required underlying issues have been discussed. From the outset it is important to note that the paper distinguishes between expert testimony and forensics. In many jurisdictions forensic evidence can only be introduced in a court case by means of expert testimony. However, not all expert testimony is based on forensics. Consider, for example, the medical doctor who testifies as an expert about the current standard of care for some ailment. This testimony will be partly based on medical training (including continuing education), partly on professional observation of what colleagues do, partly by standards that may have been published by national and international bodies and partly by local conditions (such as affordability of various treatment options). Clearly such testimony from an expert may be invaluable in a case where it is required. However, such evidence will not be classified as scientific evidence. In particular is this witness not basing evidence on forensic science.

The remainder of the paper is structured as follows. The next section reviews some characteristics of science, forensic science and expert testimony to provide context for the exploration of digital forensic science that follows. Section 3 initiates this exploration by discussing two simple (and common) scenarios at length. Section 4 uses these scenarios, the notion of decision problems and expectations about digital forensic science from the literature to begin to develop a theory of digital forensics that can claim to be scientific. Section 5 briefly mentions some competing theories. Section 6 concludes the paper.

2. ON SCIENCE, FORENSIC SCIENCE AND EXPERT TESTIMONY

As noted earlier the intention of the current paper is not to explore the notion of science in depth. In the philosophy of science the following three landmarks are most important for the purposes of this paper. Firstly, in the period before the Second World War a group known as the Vienna Circle developed the notion of logical positivism. According to them the only meaningful judgements were the tautologies from mathematics and logic, and verifiable empirical claims from science. Everything else was nonsense. The second landmark is Popper’s demarcation criterion for science: falsifiability. Only theories that can be falsified should be regarded as science. To be more specific, Popper foresees series of theories, where, when one theory is falsified, it is replaced by another theory that has greater explanatory power. Finally, Kuhn [16] describes (rather than defines) science as an endeavour where during periods of normal science, scientists solve puzzles using the paradigm then prevalent. Once an existing theory becomes unsustainable, it is replaced by a new theory (again with greater explanatory power) during what he calls a scientific revolution.

Clearly much has to be added to this admittedly superficial descriptions of science to make them useful for forensic purposes. A theory that can be falsified but has not been tested at all may qualify as science, but not as grounds for the conviction of an alleged criminal. Similarly, the mere fact that a scientist has followed the appropriate paradigm may not ensure the reliability of the results. Rather than looking at the philosophy of science for deeper understanding, we turn our attention to the law. Expert testimony in courts have a long history.

In 1782 a civil engineer and scientist testified in the Wells Harbour case in the UK. Rather than just surmising from current observations what caused the silting up of the harbour he claimed that it was "necessary to shew the natural causes by which the port of Wells has been formed" [19, p.150]. This was extraordinary since Mr Smeaton was testifying about something he did not observe, but derived from laws of nature [10]. He also did not derive those laws or even tested them. Normally such testimony would have been classified as hearsay, or even irrelevant to the specific case being heard. The opposing side did indeed attempt to get his evidence excluded. However, Lord Mansfield who was presiding over the trial wrote I cannot believe that when the question is, whether a defect arises from a natural or an artificial cause, the opinions of men of science are not to be received. The cause of the decay of the harbour is also a matter of science, and still more so whether the removal of the bank can be beneficial. On this such men as Mr. Smeaton alone can judge. Therefore we are of the opinion that his judgement, formed on facts, was very proper evidence. This is often cited as the first use of science (or forensic science) in a court of law.

Of course the use of science enabled more informed judgements to be made, but over time much pseudoscience developed where claims were made based on some set of theories that was not scientific at all. A relatively modern example of a challenge that faces courts is the use of a polygraph to obtain evidence. The validity of such evidence is the topic of much debate; many reject polygraphy as pseudoscience, while others consider it to be very reliable. Many government agencies, for example, consider polygraphy as useful [6]. Even where such evidence is not accepted, a suspect who volunteers for such a test scores some credibility points.

It is therefore important that the court acts as a gatekeeper to only allow 'valid' or 'true' science to be accepted as scientific evidence. Of course it should not be necessary to qualify science using words such as valid or true because science itself implies those characteristics.

The best known 'modern' test for admissibility of expert testimony is the Daubert standard used in the USA. In this case the court decided, amongst others, that [20]: *Faced with a proffer of expert scientific testimony under Rule 702, the trial judge, pursuant to Rule 104(a), must make a preliminary assessment of whether the testimony's underlying reasoning or methodology is scientifically valid and properly can be applied to the facts at issue. Many considerations will bear on the inquiry, including whether the theory or technique in question can be (and has been) tested, whether it has been subjected to peer review and publication, its known or potential error rate and the existence and maintenance of standards controlling its operation, and whether it has attracted widespread acceptance within a relevant scientific community. The inquiry is a flexible one, and its focus must be solely on principles and methodology, not on the conclusions that they generate.*

While this standard has been slightly revised by other courts many of the phrases of this judgement reverberates in the minds of those who are trying to establish a scientific foundation for some forensic discipline. The key phrases include tested, peer review and publication, widespread acceptance and, perhaps the most challenging of all (and not present in some later formulations of the standard, but still frequently highlighted) the known or potential error rate of the theory. Note that it is easy to critique Daubert the minority judgement that forms part of the decision cited above [20] is a good starting point for such critique. However, some mechanism is required to keep pseudoscience out of courts, and Daubert is arguably the most prominent current standard used for this purpose. Clearly then, for digital forensics to become (or remain) trusted, self reflection is necessary in the light of standards such as Daubert.

3. FORENSIC CRAFT AND SCIENCE

Let us consider just two typical scenarios encountered in ‘digital cases’ and distinguish between the craft and science involved[1]. The first and apparently most prevalent example is one where it is necessary to show that some data is present on (or absent from) some medium. The nature of the content to find may vary. In the simplest case it may be some byte sequence, such as some credit card number (say 1234-5678-9012-3456) or a specific MP3 file. In a more complex case it may be an (any) email sent between two specific parties or a (any) JPEG image depicting certain content. For ease of reference we will refer to the criteria used for searching as the search pattern even though the criteria may not be a typical pattern such as when the criteria specify a certain file type. To find the search pattern a number of subtasks need to be completed. Firstly the disc (or other media) content needs to be acquired in a forensically sound manner. Secondly, the search pattern needs to be located. Thirdly, it has to be demonstrated that the search pattern does occur on the media (and it may be necessary to indicate what the full details are, for example the content of an email that has been found based on sender and recipient). Or it may be useful to indicate that the search pattern does not exist on the media (or, a weaker claim, indicate that the content was not found on the media).

The first of these three steps is typically not a scientific activity. On physical scenes crime scene investigators (CSIs) or first responders or some other group rather than forensics scientists collect (or bag and tag) the evidence. Contamination of evidence is one of the main concerns and therefore the collectors use specific collection (or acquisition) protocols to the letter. Legal issues (such as authority to collect evidence and questions about jurisdiction) may also play a role, but once again set protocols is followed or a legal expert (rather than a forensic scientist) is consulted.

Some scientific questions may arise. If, for example, it is (or becomes) clear that the container used to collect, say, some chemical or biological material reacts to its contents, it implies that such a container may contaminate such evidence. It then becomes a question of science to find (or develop) a container for which it can be scientifically shown that it will not contaminate the evidence. Similarly, if physical evidence may degrade over time it may be necessary to develop a container that restricts such degradation by, for example maintaining the proper temperature or by preserving the evidence in some appropriate preservative compound. These are clearly questions for science. However, in some cases the scientific knowledge, such as the temperature at which evidence will degrade, is already known and it becomes purely a question of engineering to construct a container that (for our example) maintains the appropriate temperature. There may be a question about the type

[1] Note that many branches of computing combine craft and science and that a need to distinguish between craft and science or even between art, craft, techniques, engineering and science becomes necessary. See the paper by Gruner [11] about the nature of this discourse in the software engineering discipline as an example.

of science involved in this step. The fact that some biological material does not degrade below a certain temperature may be a question of 'pure' biological science, rather than forensic science. However, we do not explore this possible distinction between 'pure' and forensic science further in the current paper.

We claim that, for the current scenario where some search pattern is to be located, acquisition is in principle very similar to physical acquisition. We have known for many years how to image disc drives. We know that we ought to use write blockers to prevent contamination or, if write blockers are not available, to use an operating system that allows the disc to be mounted as a read-only device. In the latter case we may know that it is best to boot that operating system from a read-only medium and to then bag and tag the medium as evidence in case any questions are later raised about the reliability of the operating system regarding not writing to media that are mounted in a read-only mode. We know that we have to calculate some hash (such as MD5 or SHA1) for the content of the original media as well as our evidentiary copies to demonstrate the integrity of our copy. Note that very little of this process is based on science; most of it is a matter of common sense. Where science does play a role (for example in the integrity claims supported by message digests), that science is also widely used outside the realm of forensic science. Note again that this corresponds with physical forensics. The CSI who collects DNA from a suspect by brushing a swab inside the suspect's mouth to collect some saliva is typically not a scientist with a university degree in science[2]. Similarly, the officers who collects fingerprints from a crime scene or even the officers who spray Luminol to detect spilled blood are not usually scientists. Note that this does not mean that they may be unqualified or inexperienced their requirements and experience are just not as scientists and they are not expected to derive scientific truths.

A couple of remarks are in order about the digital forensic acquisition process described in the previous paragraph. Rather than imaging the device the CSI may simply seize the media and send it to the forensic laboratory to be imaged (and then analysed). However, the fact that imaging may occur in the laboratory does not make it a scientific process per se. The process of imaging described above has become known as "dead analysis", with many known shortcomings (such as its impact on business continuity for the entity being investigated). An alternative is so-called "live analysis" [1], which will not be explored further in the current paper. However, live analysis (also) desperately needs answers for the questions raised in this paper.

This concludes our discussion of the first step in the scenario described above. In summary, collection or acquisition in this scenario is primarily a technical activity (or craft), rather than a scientific activity.

Although we have distinguished between steps 2 (searching) and 3 (demonstrating presence or absence) above, the distinction does become blurred in many cases. We will, however, for the time being, use this distinction for the sake of exposition. What we do know at this point is that if digital forensics is a science, the science has to be part of step 2 or 3, since it was not present in step 1.

As noted, the second step of the given scenario entails searching for and finding (or not finding) the search pattern. Again we may use physical forensics as a point of

[2] As an example, to be formally recognised as a professional, candidate or certified natural scientist in South Africa a person has to meet the requirements specified by Act 27 of 2003 (Natural Scientific Professions Act,2003); this act establishes the South African Council for Natural Scientific Professions which is responsible for registration of scientists who meet the prescribed requirements. In general a four year degree followed by three years of professional experience or a higher degree followed by a shorter period of professional experience is required to register as a professional natural scientist.

departure. Consider an apparent murder case where the (possible) murder weapon needs to be found. It is possible that a knife is still stuck in the victim's body, in which case finding it is trivial and the process of finding it will not be considered forensic science. In a somewhat harder case the investigating officers may find a knife that they think may be the murder weapon in the suspect's home. Much now depends on the characteristics of the knife: Is it bloody? Is it similar to a set of knives from the victim's home and one such knife is missing from the victim's set? Does it have fragments of cloth stuck to it that correspond to the clothes the victim was wearing? Remember that we are assuming that there is some reason why the investigators think that this may be the weapon. If the knife is bloody, matches the set in the victim's house and contains 'obvious' fragments of cloth, the search may be over before the forensics have begun. Forensics will only come into play at step 3, where it needs to be proven that the found knife matches the victim's wound (and/or whatever other matches that may add scientific weight to the claim that the knife was indeed the murder weapon). However, if there are not such a multitude of indicators that the identified knife is the correct one forensics may begin to play a role much earlier in the search. For example, it may be determined from the wound that a serrated or smooth knife was used; it may be possible to determine the length (and perhaps other measurements of the blade); paint or other traces from the knife may determine its colour, and so on. The investigators can now proceed with a (non-forensic) search based on what they have learnt from the forensic scientists about the weapon they are looking for. Finally, if the murder weapon was some poison, searching for (traces of) it may be a pure forensic exercise.

Our digital forensic scenario requires us to find some specified data on the disc image. Let us now make a sacrilegious claim: In general any tool may be used to search for the data. Of course the tool needs to be suitable for performing the search: if we are looking for a type of file (rather than exact text) we need a tool that is able to search for such files. To illustrate, suppose the investigator copies the image to a hard disc of a computer and then boots the computer from this disc. Suppose the investigator opens the email application and uses its search fields to find the email messages between the two parties that are of interest. Or suppose we are indeed looking for some pattern; suppose the investigator uses grep or some other pattern matching program to find the required files. And, in any of these cases, suppose the investigators find what they were looking for. Is there any reason to object because 'non-forensic' or 'untrusted' tools have been used? I claim that there ought to be no objection. The claim is based on the assumption that we will during step 3 prove that the search pattern does indeed exist on the medium. Whatever methods we used to locate it are irrelevant.

Objections to these claims may come from multiple sources. Firstly, the notion of using untrusted tools in a regular forensic laboratory is unthinkable. Who knows what such tools may do to the evidence and in what way they may contaminate the evidence. But in the digital world we have the luxury of working with copies of evidence. Even if we destroy a copy we can just make a new copy from our master copy, check the message digests and no harm has been done (besides our time that may have been wasted).

Another objection may be that the non-forensic tool we are using may be 'biased' in some way; for some peculiar reason it may find incriminating evidence, but miss the exculpatory evidence. Say, for example, A emails a ransom note to B and five minutes later emails a note that it was an April fool joke. This behaviour may still be illegal, but these messages may be interpreted very differently depending on whether both or only the first message are discovered. This objection clearly has merit in some cases. However, the sad reality is that in many (possibly most) current digital forensic investigations this makes little difference: In so many current investigations investigators are looking for contraband; if the suspect is guilty hundreds (or more) of examples are typically found. Exculpatory evidence (if it can exist) will have a very different form from what is being searched for. Say the disc contains many illegal (or unlicensed) MP3 files. Then it does not matter whether we find all of them; yet another MP3 will not serve as exculpatory evidence. Exculpatory evidence may exist in the form of a letter from a copyright holder granting permission to the suspect to copy their MP3 files without licences for, say, research purposes. This letter will typically be produced by the other party to explain the presence of the files. However, our claim that any tool may be used is dangerous when it is necessary to find all occurrences of the data of interest, or if the want to conclude that the data does not occur on the media at all. For such cases we need a tool we can trust; however, even for such cases there is no reason to use non-forensic tools if using them holds some benefit such as the ability to find at least some occurrences faster than the trusted tool.

The final objection against the use of any tool to be considered may come from those who infer that our untested tool may go outside the boundaries of what we are legally allowed to access. This certainly is not the intention. The proper analogy to use when using these non-forensic tools is not the physical forensic laboratory, but the police officer who searches a room for evidence. This can only be done once an appropriate warrant has been issued and then the search has to be confined to the limits set out in the warrant. If this officer wants to use a flashlight to look into a dark corner of the room, it is ridiculous to require that it has to be a forensically sound flashlight. If the officer wants to read a label on a box that may be accessed and needs reading glasses, there is no need to ensure that they are forensically sound glasses. But when the officer looks into a cupboard that is beyond the limits of the warrant, evidence obtained will be inadmissible (in addition to punitive measures against the officer that should result). So, when using arbitrary tools to search data it is necessary to ensure that the limits of the warrant are respected. In many cases tools (such as grep) are simple enough to restrict to search within limits. Alternatively, the ‘forbidden’ areas of the disc may be redacted or the allowed areas may be copied to a clean disk. Either option, if executed correctly, will avoid any possible problems.

One practical consequence is this: If the investigator gives a copy of the (redacted) evidence to his or her sysadmin who is a Unix toolset, bash and scripting guru with the request to use his or her ingenuity to find the search pattern, whatever is found ought to be admissible. (This of course assumes that the sysadmin is authorised to access the evidence.) Note again that what the sysadmin does is not science irrespective of how brilliant the search strategy may be.

We have spent an inordinate amount of space to the simple issue of searching for specific data in some data set. However, my sense is that most current forensic investigations occur in this space and that many who are looking for the science in digital forensic science are looking for it in this space. To illustrate the first point just consider the types of investigations that fit in this category. It includes searching for contraband, deleted logs, entries in the registry that indicate (former) presence of a specific program or device, credit card numbers, IP addresses, events in logs, events or modified files within some time period, fragments of known files and many more. Science may play a role in optimising the search strategies. However, the forensic investigation does not pose any specific requirements. Therefore it seems inappropriate to consider ‘forensic searching’ as a relevant problem area for this scenario. Some search algorithms may hold certain benefits for forensics (and quite possibly other fields); for example strategies that yield initial results early in a specific search domain may be beneficial.

However, the requirements change when it is necessary to know that the search pattern does not occur on the disc at all or to find all instances of the search pattern. Similarly, issues arise when there is only sufficient time available to search a fraction of the available data. These cases are revisited after step 3 of the scenario has been considered. Step 3 entails proving that the search pattern exists or (equivalently) revealing the details of the found search pattern (by, for example, revealing the credit card number found if a pattern conforming to a credit card number was used as search pattern). In its first form the requirement is clearly that a decision problem has to be answered: Does the given search pattern occur on the disc? Decision problems are well known from the field of algorithmics [12] (or computational complexity). And, from that same field we know the second formulation above is computationally equivalent to the decision problem. And thus we find ourselves with a problem for which a solid theoretical framework exists and can be answered in a scientific manner. In the scenario under discussion the question about the presence of the search pattern may be answered positively in an incontrovertible manner by simply pointing to where the data occurs on the image. Formulated in its current form the problem is tractable and answerable in absolute terms. The error rate is 0scientist can answer this question with absolute scientific certainty in the witness box. If scenario 1 deals with the possession of contraband, finding contraband on the disc allows the prosecution to introduce the disc image as evidence. If contraband has not been found it is possible to simply not introduce that disc as evidence. However, the defence is potentially faced with a bigger challenge: they want to prove that no contraband occurs on the disc (or any other disc either). Suppose that the message digests of files containing contraband are known. Then they simply have to compute the message digests of all the files on the disc and show that none of those digests corresponds with any of the contraband digests. This is again clearly a decision problem. However, this pushes the 'burden of proof' to step 2 of the scenario and there is no step 3 where one can simply point to the fact that nothing has been found. Ideally the defence needs to know that their search algorithm is correct and that the search problem itself is tractable. The issue of correctness may again be addressed from the perspective of algorithmics where the algorithm is formally proven correct (and where the accuracy of the algorithm is therefore 100%)[3]. A less desirable alternative is where trust develops in a certain search tool where opposing parties use it (and other tools) over many years and nobody finds any contraband missed by the other party. However, this only becomes scientific at the point where one can move from mere induction ('it has worked thus far and will therefore probably work in the next case as well') to where one may express one's confidence in the tool in scientific terms. Formal testing of the tool seems useful in this regard.

In addition to correctness the defence in our example ideally wants the search problem to be tractable (or even if it is tractable in general, they want the answer to be available before it is needed for testimony in court, for example). As indicated by Cohen [5] the field of computational complexity may provide us with the answer to the dilemma of whether it is even worth starting the computation. However, there may be another alternative available: a probabilistic algorithm may provide an

[3] Note that correctness of the algorithm does not ensure correctness of the program; a simple option is to use multiple independent tools in parallel with the (probably valid) assumption that these independent programs will not contain coding errors that let them all fail in the same manner. However, this brings us back to an assumption, rather than a scientific fact. A deeper review of the field of software correctness is required than what can justifiably be provided here to be certain that the tool is correct.

answer that is correct with a given certainty. Executing the probabilistic algorithm repeatedly increases the certainty (or finds a counterexample). If the problem is intractable, probabilistic algorithms may provide us with a scientific answer with a quantifiable error rate. Even if the problem is tractable but requires more time than is available, it may be possible to use a much simpler probabilistic algorithm and run it repeatedly. This will again yield a scientifically valid answer with a quantified probability of being incorrect.

This, at long last, brings us to the end of scenario 1 that set out to locate or prove the absence of some data on a disc image in a scientific manner. We now turn our attention to just one other scenario that illustrates a different case where the craft may be turned into science.

Scenario 2 deals with file carving. File systems organise files in blocks, sectors, clusters or some other units (henceforth just referred to as blocks). Files typically consist of multiple blocks that are linked together using metadata. If these links are destroyed the file is effectively lost even though the file contents may physically still be present on the disc. The links may be lost because of an attack or some accident. It is, for example, possible that a user deletes a file because it is no longer deemed necessary. Deleting the file typically deletes the links, but not the block contents. If it turns out that the information in such a deleted or lost file is important the question arises whether the blocks can be reassembled into the initial file. Such reassembly is known as file carving. Note that while the blocks are unlinked some of them may be reused for other files; therefore it is sometimes at best possible to carve a partial file.

Obviously file carving requires deep knowledge of the details of file systems. The carver needs to know how the metadata links blocks together in the specific file system as well as the other minutiae of the file system. In addition the carver needs to know the details of the file formats of the files being carved to recognise neighbouring blocks. In essence the carver is solving a jigsaw puzzle that has many extraneous pieces and where a few required pieces may be missing.

Now suppose that the carved file is used as evidence in a court case. Does the carver have scientific grounds to claim that the file has been reconstructed correctly? An intuitive answer may be that the mere fact that, say, a JPEG file that has been reconstructed from blocks scattered over a disc now successfully opens in an image viewer is sufficient evidence that reconstruction was done correctly. It seems just too improbable that a file with an incorrect block somewhere will still 'work'. But suppose the disc contained several versions of a given file with only minor differences between the versions. Is it not then possible that the reconstructed file may contain blocks from different versions forming a carved file that never existed in that exact form? And can it be guaranteed that there are no other situations where a combination of inappropriate blocks may seem like a valid file?

A somewhat different approach is to ask what can be said about the reconstructed file that is scientifically true (and hence truly forensic science). One example is the question whether the reconstructed file conforms to the expected format. File formats are often specified using some formal notation, such as a grammar. If not, it is in many cases possible to create a grammar-based specification from whatever specification exists possibly even from reverse engineering an authoritative piece of software that creates such files. The notion of syntax checking is well understood from the field of compiler construction. The question whether the reconstructed file is syntactically correct is therefore one example of a question that may be phrased as a decision problem and answered in a scientific manner. Many other properties may be checked in this manner. If certain values in a file are expected to have some relationship to one another this may be verified. The time stamps in a log file, for example, are supposed to be ordered according to time. In some (rare) cases it may be possible to show that no other blocks on the disc can possibly be part of a file of the given type. It may be possible to show that the blocks in the carved file are arranged on disc in a manner consistent with the block allocation strategy used by the operating system. It may be possible to allocate all blocks on the disc to files that are all syntactically, semantically and positionally correct. It may be possible to test all permutations of blocks (possibly after filtering those out that cannot possibly form part of the given file type) and show that the reconstructed file is the only permutation starting from some block that yields a syntactically valid file. Based on these scientific facts the expert may then offer a professional opinion about the correctness of the reconstructed file. The opinion may take into account the complexity of the format, the consistency of the reconstructed file with other available information and other attributes of the file the professional may deem relevant. It is important to distinguish between science and opinion though. Different forensic scientists should arrive at the same scientific answers to questions that can be answered by forensic science. If their opinions differ, so be it. They are opinions after all, and should have less evidentiary weight than scientific facts. However, note that if it can be shown that an opinion is inconsistent with facts, that opinion is refuted.

Note that this second scenario conveniently ignored the fact that many real programs do not faithfully implement file format standards. It is therefore possible that an original file may not pass the syntax check and if such a file is reconstructed correctly it should fail the syntax check. However, this may possibly be addressed by not only using de jure specifications, but also de facto specifications. We ignore this issue in the remainder of the paper.

To conclude note that this distinction between fact and opinion is also present in traditional (physical) forensic science. The DNA scientist cannot ‘place’ a person at a crime scene. The scientist can state as a scientific fact that, say, a hair and some saliva come from the same donor. The additional (non-scientific) information that the hair was found at the crime scene and the saliva sample was obtained from the suspect (as well as some convincing argument that there is no other logical explanation for the suspect’s hair to be at the crime scene) is required to be certain that the suspect was indeed at the crime scene.

4. DIGITAL FORENSIC SCIENCE

The two scenarios discussed earlier in this paper show that decision problems may indeed provide the scientific basis for digital forensics for some cases; in such cases decision problems may be used to distinguish between forensic science and expert opinion. Those two scenarios are insufficient to claim that decision problems can be used as the underlying theoretical base of all of digital forensic science. However, it is a strategy that seems worth exploring. As Garfinkel [8] points out, locating incriminating information (such as contraband) in large datasets was the original challenge for digital forensics and the field needs to urgently cast its net wider to remain relevant. A digital forensic science based on decision problems (and the accompanying algorithmics or complexity theory) provides much scope for forensics to develop beyond its current state. Garfinkels identification of the original challenge of digital forensics coincides with scenario 1 provided earlier in this paper. Much of digital forensics was originated by finding ways of solving crimes (or finding digital evidence) that may be useful to address such crimes. If we decide that decision problems underly digital forensics it also becomes possible to develop digital forensic science from the top down by determining what can and what cannot be proven by viewing the extensive body of knowledge about tractable and intractable problems from this new perspective [4].

An initial argument that decision problems should form the basis of digital forensic science may read as follows. Many (or most) digital forensic investigators will be comfortable with characterising the examination process as a set of hypotheses that are tested and then rejected or not rejected[5]. The work by Carrier [4] is a seminal text that frames digital forensics using hypothesis testing. The idea of using decision problems and determining the answers they yield (or concluding that they cannot be answered) appear rather similar to hypothesis testing. Hypothesis testing, however, typically assumes natural variation and testing a hypothesis is about determining whether minute differences between an observation and an ideal value may be ascribed to this natural variation. Digital data, on the other hand, being discrete, does not display such natural variation. The millions of statements produced by a bank on a monthly basis are not a little wrong each month because of natural variation. If the statement is not exactly correct it is because something is amiss. Natural variation may be introduced in a digital system because of external physical influences. The time that data needs to traverse a network is one such example; this may result in a natural variation between times recorded in a log at the transmitter and times recorded for the same messages at the receiver. However, it is not clear that these differences are indeed natural. The digital realm is one that is inherently artificial. Users can influence congestion on the network and hence the differences in times. In fact, such times are often affected by multiple natural and artificial causes that make it impractical to measure a given characteristic and associate it with scientific accuracy with some specific cause or condition. Decision problems may therefore fit digital data better than hypothesis testing would for forensic purposes.

Note that decision problems, just like hypotheses, do not prescribe how an examination should be conducted, but clearly delineates what may be offered as evidence. An 'accepted' hypothesis makes a truth claim as does a decision problem that has been decided.

The remainder of this section reviews the (well known) classes into which decision problems fall [12]. The intention is twofold. Firstly, it shows how much of the field remains unexplored from a forensic perspective and therefore indicates a direction into which future forensic research may grow. It also shows how error rates naturally become an issue when decision problems become more complex. This potentially lends some credibility for a forensic investigator who claims 100% accuracy for a result based on a simple decision problem (relative to a whole field of varying complexity where error rates are no longer zero).

In general decision problems fall in one of four categories: they are decidable in polynomial time, probabilistically decidable in polynomial time, intractable or undecidable. The second category in this list gives us our first glimpse of what error rates may mean in the context of decidable digital forensics. We return to this topic below. When the question of interest is polynomially decidable there is no inherent need to quantify error rates. However, even polynomial time algorithms may sometimes be too 'expensive': to search a petabyte of information in $O(n)$ at 1 megabyte per second will take just over 30 years. A probabilistic algorithm that does not sample every byte of the petabyte and that yields a result that is reliable enough but terminates within some reasonable time will be preferable over

[4] Note that Garfinkels plea for an extension of digital forensics refers primarily to the extension of technology used for digital forensics that is, to digital forensic craft rather than digital forensic science.

[5] Note that the mere use of hypothesis testing (outside a body of theory) would not be sufficient to make an activity scientific. For more details see [2, Chapter 5].

the absolutely correct $O(n)$ algorithm. In general, given the large data sets that digital forensics often has to deal with, it may be necessary to approximate the algorithm with an even faster one (one that, for example, only uses a fraction of the n inputs) if the results of the probabilistic algorithm are correct enough that is, if the error rate can be quantified and it is deemed small enough to sufficiently substantiate the claim that it supports.

Probabilistic algorithms (also known as randomised algorithms) are algorithms that use a random number to determine their behaviour. The type of probabilistic algorithm alluded to above is a Monte Carlo algorithm one that always terminates in polynomial time and produces an answer with a known error rate. Monte Carlo algorithms may be true biased, false-biased, or unbiased. When a true-biased algorithm returns true the answer to the problem is indeed true, which is often written as yes. When it returns false (or no), however, it may be wrong with some known (small) probability. The converse is true for false-biased algorithms. Unbiased algorithms may yield incorrect results (with some small probability) when they return either true or false. Monte Carlo algorithms are deigned such that the random number determines the execution of the algorithms, such that one execution of the algorithm is independent from the next and the probability of error from two executions of the algorithm are then the product of the probability of error during a single execution. To reach a particular level of certainty it is necessary to repeat the execution of the algorithm a sufficient number of times so that the combined error is small enough. Note that, if a true-biased algorithm returns yes during any execution, the final answer is true. It is only when it repeatedly returns no that the answer is no with a probability of error en, where e is the probability of error for a single execution and n is the number of executions. The same applies to false-biased algorithms, except that a no result is certain and a yes result is reached with a margin of error. We do not consider two-sided errors further in the current paper.

The class of problems that are solvable by probabilistic algorithms are known as the bounded-error probabilistic polynomial (**BPP**) class of problems. Let P (as usual) denote the class of problems that are solvable in polynomial time. Then we contend that probabilistic algorithms are indicated for any problems in **BPP - P**. (Note that is is possible and many indeed conjecture that **BPP = P**.) As noted, Probabilistic algorithms may also be useful for problem in **P**, where available time simply does not allow execution of an exhaustive algorithm, even though it may be tractable.

5. ALTERNATIVE PERSPECTIVES ON ERROR RATES AND DIGITAL FORENSIC SCIENCE

As noted earlier, others have proposed strategies to deal with accuracy (or known error rates) of digital evidence. Cohen [5], for example, notes that the error rates of CPUs are known and suggests that this may be used to quantify the accuracy of digital evidence. However, such random CPU errors do not necessarily translate to specific error rates in digital evidence. In many cases data is, for example, subjected to error checks (such as integrity checks in databases, digitally signed messages, ordinary parity checks for memory, and so on). Some errors may cause a program to crash, rather than produce incorrect results.

Yet other errors may be inconsequential such as when the colour of a single pixel on a screen is somewhat wrong. The fact is that such errors are extremely rare and of the few that occur, many will have no impact on evidence that is collected. If it does affect the evidence it is possible that it may affect it in such a way that it is obvious that something is wrong. Once all of this is taken into account it is easy to see that errors may occur, but that this will occur so rarely that it is safe to ignore the possibility. However, with all these factor impacting on the error rates it becomes impossible to quantify the known error rates of our forensic techniques.

An earlier approach to describe (rather than quantify) error rates is Caseys certainty scale. It, for example, postulates that an event that has been logged in two independent logs may be accepted as fact with more certainty than an event only logged in one log (but this certainty will still be very low if the two logs are not properly secured). While this makes sense, it is not a scientific truth. An event logged in a number of highly secure, independent logs may lead to a high level of certainty that it really occurred. But it is possible that the administrators of all those systems colluded and entered a fake entry in all the logs. In contrast, an event logged in a single, unreliable log may indeed have occurred. The higher degree of certainty is basedat least in parton the assumption that a group of trusted individuals associated with independent systems will very rarely collude. While this is probably true, the average digital forensic scientist is not qualified to testify about human natureand questions of human nature should arguably not be part of the domain of digital forensics. In any case, it seems unlikely that even a social scientist will be able to accurately estimate this probability. This does not mean that Caseys certainty scale is useless; it does mean that the certainty scale may be unsuitable to derive scientific facts. It may be very useful for an expert to express an opinion once the scientific facts have been determined.

Finally, Garfinkel et al [7] emphasise the ability to independently verify test results as the hallmark of science and encourage the development of standardised corpora that may be used for independent testing (and provide some such corpora).

6. Conclusion

This paper identified a possible basis to ensure that digital forensics is indeed scientific, namely decision problems from the field of algorithmics. It illustrated that decision problems may indeed be useful for some investigative problems. Decision problems also help to talk about facets of science such as truth and error rates. It provides a possible explanation for why it is currently hard to talk about such issues, because current research has only scratched the surface of this domain (once such research is rephrased in terms of decision problems).

Decision problems may be helpful to guide the construction of digital forensic tools that can be certified as reliable. Much remains to be done. Many other investigative scenarios need to be considered to determine whether decision problems form an appropriate solution, or whether there are better options to obtain scientifically valid evidence for such scenarios. Decision problems also potentially delimit the scope of digital forensics and delineation is often a source of contention. Do authorship attribution [14] and source camera identification [17], for example, still form part of digital forensics or are they really about human and physical attributes that just happen to be represented in a digital format, but may just as well have been presented in a non-digital manner? If the proposal contained in this paper is accepted as a viable option by the digital forensics community only time will provide definitive answers to these latter questions.

References

[1] F. Adelstein. Live forensics: Diagnosing your system without killing it first. Communications of the ACM, 49(2):6366, February 2006.

[2] M. Bunge. Philosophy of Science: From Problem to Theory, volume 1. Transaction Publishers, 1998.

[3] M.A. Caloyannides. Digital “Evidence” is Often Evidence of Nothing, pages 334339. IGI, 2006.

[4] B.D. Carrier. A Hypothesis-based Approach to Digital Forensic Investigations. PhD thesis, Purdue University, 2006.

[5] F. Cohen. Digital Forensic Evidence Examination. Fred Cohen & Associates, 3rd edition, 2012.

[6] J.J. Furedy. The North American polygraph and psychophysiology: Disinterested, uninterested, and

[7] S. Garfinkel, P. Farrell, V. Roussev, and G. Dinolt. Bringing science to digital forensics with standardized forensic corpora. Digital Investigation, 6:S2S11, 2009.

[8] S.L. Garfinkel. Digital forensics research: The next 10 years. Digital Investigation, 7(Supplement):S64S73, August 2010.

[9] P.C. Giannelli. Wrongful convictions and forensic science: The need to regulate crime labs. Working Paper.

[10] T. Golan. Laws of Men and Laws of Nature: The History of Scientific Expert Testimony in England and America. Harvard University Press, 2007.

[11] S. Gruner. Software engineering between technics and science recent discussions about the foundations and the scientificness of a rising discipline. Journal for General Philosophy of Science, 41:237260, 2010.

[12] D. Harel. Algorithmics: The Spirit of Computing. Pearson Education, 2nd edition, 1992.

[13] D.R. Hofstadter. Gódel, Escher, Bach: An Eternal Golden Braid. Harvester Press, 1979.

[14] P. Juola. Authorship attribution. In Foundations and Trends in Information Retrieval, Vol. 1, pp. 233334, 2006.

[15] R. Koppl and M.M. Ferraro. Digital devices and miscarriages of justice. The Dayly Caller, 2012. Online: http://dailycaller.com/2012/06/15/digital-devices-and-miscarriages-of-justice/.

[16] T.S. Kuhn. The Structure of Scientific Revolutions. University of Chicago Press, 3rd edition, 1996.

[17]] M.S. Olivier. Using sensor dirt for toolmark analysis of digital photographs. In Indrajit Ray and Sujeet Shenoi, editors, Advances in Digital Forensics IV, pages 193206. Springer, 2008.

[18] M.J. Saks and D.L. Faigman. Failed forensics: How forensic science lost its way and how it may yet find it. Annual Review of Law and Social Science, 4:149171, 2008.

[19] J. Smeaton. Reports of the late John Smeaton, F.R.S., made on various occasions, in the course of his employment as a civil engineer, vol. II. M. Taylor, 2nd edition, 1837.

[20] U.S. Supreme Court . Daubert v. Merrell Dow Pharmaceuticals, inc., 509 U.S. 579 (1993). Technical Report pp.92102, Certiorari to the United Sstates Court of Appeals for the Ninth Ccircuit, 1993.

Acknowledgement 1. *The author would like to thank Stefan Gruner for providing some valuable insights after reading an earlier draft of this paper and Candice le Sueur for proofreading an almost final version of the paper. The author obviously remains responsible for any errors or inaccuracies that remain.*

(Martin S. Olivier)
ICSA RESEARCH GROUP, COMPUTER SCIENCE,
UNIVERSITY OF PRETORIA, SOUTH AFRICA
E-mail address, M.S. Olivier: ms.olivier@olivier.ms
URL: http://mo.co.za

Aktuelle Technische Berichte des Hasso-Plattner-Instituts

Band	ISBN	Titel	Autoren / Redaktion
62	978-3-86956-212-4	**Covering or Complete? Discovering Conditional Inclusion Dependencies**	Jana Bauckmann, Ziawasch Abedjan, Ulf Leser, Heiko Müller, Felix Naumann
61	978-3-86956-194-3	**Vierter Deutscher IPv6 Gipfel 2011**	Christoph Meinel, Harald Sack (Hrsg.)
60	978-3-86956-201-8	**Understanding Cryptic Schemata in Large Extract-Transform-Load Systems**	Alexander Albrecht, Felix Naumann
59	978-3-86956-193-6	**The JCop Language Specification**	Malte Appeltauer, Robert Hirschfeld
58	978-3-86956-192-9	**MDE Settings in SAP: A Descriptive Field Study**	Regina Hebig, Holger Giese
57	978-3-86956-191-2	**Industrial Case Study on the Integration of SysML and AUTOSAR with Triple Graph Grammars**	Holger Giese, Stephan Hildebrandt, Stefan Neumann, Sebastian Wätzoldt
56	978-3-86956-171-4	**Quantitative Modeling and Analysis of Service-Oriented Real-Time Systems using Interval Probabilistic Timed Automata**	Christian Krause, Holger Giese
55	978-3-86956-169-1	**Proceedings of the 4th Many-core Applications Research Community (MARC) Symposium**	Peter Tröger, Andreas Polze (Eds.)
54	978-3-86956-158-5	**An Abstraction for Version Control Systems**	Matthias Kleine, Robert Hirschfeld, Gilad Bracha
53	978-3-86956-160-8	**Web-based Development in the Lively Kernel**	Jens Lincke, Robert Hirschfeld (Eds.)
52	978-3-86956-156-1	**Einführung von IPv6 in Unternehmensnetzen: Ein Leitfaden**	Wilhelm Boeddinghaus, Christoph Meinel, Harald Sack
51	978-3-86956-148-6	**Advancing the Discovery of Unique Column Combinations**	Ziawasch Abedjan, Felix Naumann
50	978-3-86956-144-8	**Data in Business Processes**	Andreas Meyer, Sergey Smirnov, Mathias Weske
49	978-3-86956-143-1	**Adaptive Windows for Duplicate Detection**	Uwe Draisbach, Felix Naumann, Sascha Szott, Oliver Wonneberg
48	978-3-86956-134-9	**CSOM/PL: A Virtual Machine Product Line**	Michael Haupt, Stefan Marr, Robert Hirschfeld
47	978-3-86956-130-1	**State Propagation in Abstracted Business Processes**	Sergey Smirnov, Armin Zamani Farahani, Mathias Weske
46	978-3-86956-129-5	**Proceedings of the 5th Ph.D. Retreat of the HPI Research School on Service-oriented Systems Engineering**	Hrsg. von den Professoren des HPI

ISBN 978-3-86956-204-9
ISSN 1613-5652