Joaquin Garcia-Alfaro · Georgios Lioudakis
Nora Cuppens-Boulahia · Simon Foley
William M. Fitzgerald (Eds.)

LNCS 8247

Data Privacy Management and Autonomous Spontaneous Security

8th International Workshop, DPM 2013, and 6th International Workshop, SETOP 2013 Egham, UK, September 12–13, 2013 Revised Selected Papers

Springer

Lecture Notes in Computer Science 8247

Commenced Publication in 1973
Founding and Former Series Editors:
Gerhard Goos, Juris Hartmanis, and Jan van Leeuwen

Editorial Board

For further volumes:
http://www.springer.com/series/7410

Joaquin Garcia-Alfaro · Georgios Lioudakis
Nora Cuppens-Boulahia · Simon Foley
William M. Fitzgerald (Eds.)

Data Privacy Management and Autonomous Spontaneous Security

8th International Workshop, DPM 2013, and
6th International Workshop, SETOP 2013
Egham, UK, September 12–13, 2013
Revised Selected Papers

Springer

Editors
Joaquin Garcia-Alfaro
Telecom SudParis
Evry
France

Georgios Lioudakis
National Technical University of Athens
Athens
Greece

Nora Cuppens-Boulahia
Telecom Bretagne
Cesson Sévigné
France

Simon Foley
University College Cork
Cork
Ireland

William M. Fitzgerald
IDA Ovens
EMC Information Systems International
Cork
Ireland

ISSN 0302-9743 ISSN 1611-3349 (electronic)
ISBN 978-3-642-54567-2 ISBN 978-3-642-54568-9 (eBook)
DOI 10.1007/978-3-642-54568-9
Springer Heidelberg New York Dordrecht London

Library of Congress Control Number: 2014934122

LNCS Sublibrary: SL4 – Security and Cryptology

Printed on acid-free paper

Springer is part of Springer Science+Business Media (www.springer.com)

Foreword from the DPM 2013 Program Chairs

This volume contains the proceedings of the 8th Data Privacy Management International Workshop (DPM 2013), held in Egham, UK, at Royal Holloway, University of London, during September 12–13, 2013, in conjunction with the 18th annual European research event in Computer Security (ESORICS 2013) symposium. It includes a revised version of the papers selected for presentation at the workshop. Previous issues of the DPM workshop were held in 2012 in Pisa (Italy), 2011 in Leuven (Belgium), 2010 in Athens (Greece), 2009 in Saint Malo (France), 2007 in Istanbul (Turkey), 2006 in Atlanta (USA), and 2005 in Tokyo (Japan).

The aim of DPM is to promote and stimulate the international collaboration and research exchange on areas related to the management of privacy-sensitive information. This is a very critical and important issue for organizations and end-users. It poses several challenging problems, such as translation of high-level business goals into system level privacy policies, administration of sensitive identifiers, data integration and privacy engineering, among others.

In response to the call for participation, 46 submissions were received. Each submission was evaluated on the basis of significance, novelty, and technical quality. All submissions went through a careful anonymous review process (three or more reviews per submission) aided by 49 Technical Program Committee members and 31 additional referees. In the end, 13 full papers, accompanied by five short papers, were presented at the event. The final program also included three invited talks by Steven J. Murdoch (University of Cambridge), Emil Lupu (Imperial College London), and John Borking (former Privacy Commissioner and Board Member of the Dutch Data Protection Authority in The Hague). Our special thanks to Steven, Emil, and John for accepting our invitation and for their presence during the event and talks.

We would like to thank everyone who helped at organizing the event, including all the members of the Organizing Committee of both ESORICS and DPM 2013. In particular, we would like to highlight and acknowledge the tremendous efforts of the ESORICS 2013 General Chair Keith Mayes and his team. Thank you Keith for all your help and support with DPM. Our gratitude goes also to Pierangela Samarati, Steering Committee Chair of the ESORICS Symposium, for all her arrangements to make possible the satellite events. Our special thanks to the General Chairs of DPM 2013, Josep Domingo-Ferrer and Maryline Laurent, as well as Steering Committee member Guillermo Navarro-Arribas, for their unconditional help since the beginning of this event. Last but by no means least, we thank all the DPM 2013 Program Committee members, additional reviewers, all the authors who submitted papers, and all the workshop attendees.

Finally, we want to acknowledge the support received from the sponsors of the workshop: Institute Mines-Telecom, CNRS Samovar UMR 5157, Telecom SudParis, UNESCO Chair in Data Privacy, and National Technical University of Athens.

January 2014

Joaquin Garcia-Alfaro
Georgios Lioudakis

8th International Workshop on Data Privacy Management—DPM 2013

Program Committee Chairs

Joaquin Garcia-Alfaro	Telecom SudParis, France
Georgios Lioudakis	National Technical University of Athens, Greece

Workshop General Chairs

Josep Domingo-Ferrer	Universitat Rovira i Virgili, Spain
Maryline Laurent	Telecom SudParis, France

Program Committee

Esma Aimeur	Université de Montreal, Canada
Michel Barbeau	Carleton University, Canada
John Borking	Borking Consultancy, The Netherlands
Jens-Matthias Bohli	NEC Laboratories Europe, Germany
Ana Cavalli	Telecom SudParis, France
Frederic Cuppens	Telecom Bretagne, France
Nora Cuppens-Boulahia	Telecom Bretagne, France
Roberto Di Pietro	Roma Tre University of Rome, Italy
Nicola Dragoni	Technical University of Denmark, Denmark
Christian Duncan	Quinnipiac University, USA
David Evans	University of Derby, UK
Sara Foresti	Università degli Studi di Milano, Italy
Sebastien Gambs	University of Rennes 1, France
Flavio D. Garcia	Radboud University Nijmegen, The Netherlands
Paolo Gasti	New York Institute of Technology, USA
Francesca Gaudino	Baker & McKenzie Law Firm, Italy
Stefanos Gritzalis	University of the Aegean, Greece
Marit Hansen	Unabhängiges Landeszentrum für Datenschutz, Germany
Artur Hecker	Telecom ParisTech, France
Jordi Herrera	Autonomous University of Barcelona, Spain
Iakovos Venieris	National Technical University of Athens, Greece
Dimitra Kaklamani	National Technical University of Athens, Greece
Panos Kampanakis	Cisco Systems, USA
Georgia Kapitsaki	University of Cyprus, Cyprus

Sokratis Katsikas	University of Piraeus, Greece
Evangelos Kranakis	Carleton University, Canada
Jean Leneutre	Telecom ParisTech, France
Giovanni Livraga	Università degli Studi di Milano, Italy
Javier Lopez	University of Malaga, Spain
Brad Malin	Vanderbilt University, USA
Sotirios Maniatis	Hellenic Authority for Communications Privacy, Greece
Chris Mitchell	Royal Holloway, UK
Refik Molva	Eurecom, France
Krish Muralidhar	University of Kentucky, USA
Guillermo Navarro-Arribas	Autonomous University of Barcelona, Spain
Silvio Ranise	Fondazione Bruno Kessler, Italy
Kai Rannenberg	Goethe University Frankfurt, Germany
Indrajit Ray	Colorado State University, USA
Yves Roudier	Eurecom, France
Mark Ryan	University of Birmingham, UK
Claudio Soriente	ETH Zürich, Switzerland
Alessandro Sorniotti	IBM Research, Switzerland
Traian M. Truta	Northern Kentucky University, USA
Yasuyuki Tsukada	NTT Communication Science Laboratories, Japan
Jens Weber	University of Victoria, Canada
Lena Wiese	University of Göttingen, Germany
Yanjiang Yang	Institute for Infocomm Research, Singapore
Nicola Zannone	Eindhoven University of Technology, The Netherlands
Melek Önen	Eurecom, France

Steering Committee

Josep Domingo-Ferrer	Universitat Rovira i Virgili, Spain
Joaquin Garcia-Alfaro	Telecom SudParis, France
Guillermo Navarro-Arribas	Autonomous University of Barcelona, Spain
Vicenç Torra	Artificial Intelligence Research Institute, Spain

Additional Reviewers

Achilleas Achilleos	Christian Kahl
Ahmad Sabouri	David Galindo
Alessio Di Mauro	David Nuñez
Ana Nieto	Elisa Costante
Anderson Morais	Eugenia Papagiannakopoulou
Anis Bkakria	Fatbardh Veseli
Aouadi Mohamed	Flavio Lombardi

Foreword from the SETOP 2013 Program Chairs

These are the proceedings of the 6th International Workshop on Autonomous and Spontaneous Security (SETOP 2013).

The purpose of this workshop is to bring together researchers to explore challenges in the automated configuration of security. In this volume you will find papers on a range of topics related to authentication and authorization, mobile security and vulnerabilities.

The workshop program also included invited talks by Steven Murdoch (University of Cambridge, UK) on "Quantifying and Measuring Anonymity" and by Emil Lupu (Imperial College London) on "Pervasive Autonomous Systems: Challenges in Policy based Adaptation and Security."

As with previous years, SETOP was a satellite workshop of the European Symposium on Research in Computer Security (ESORICS). We are grateful to the ESORICS 2013 Organizing Committee for agreeing to host SETOP-2013 and especially to ESORICS General Chair Keith Mayes for his assistance and support.

We are grateful to the many people who contributed to the success of the workshop. The members of the Program Committee and external reviewers. The Publications Chair, William Fitzgerald assembled the workshops proceedings and ensured its timely publication.

Finally, the workshops would not be possible without the authors who submitted papers, the presenters, and attendees.

We hope you enjoy reading the proceedings.

Nora Cuppens-Boulahia
Simon Foley

6th International Workshop on Autonomous and Spontaneous Security—SETOP 2013

Program Committee Chairs

Research Track

Simon Foley	University College Cork, Ireland
Nora Cuppens-Boulahia	Telecom Bretagne, France

Industrial Track

Edgardo Montes de Oca	Montimage, France

Workshop General Chairs

Ana Cavalli	Telecom SudParis, France
Frédéric Cuppens	Telecom Bretagne, France

Publicity and Publication Chair

William Fitzgerald	University College Cork, Ireland

Webmaster

Said Oulmakhzoune	Telecom Bretagne, France

Program Committee

Fabien Autrel	Telecom Bretagne, France
Gildas Avoine	Catholic University of Louvain, Belgium
Michele Bezzi	SAP Research, France
Christophe Bidan	Supelec, France
Carlo Blundo	University of Salerno, Italy
Joan Borrell-Viader	UAB, Spain
Jordi Castella-Roca	Rovira i Virgili University, Spain
Iliano Cervesato	Carnegie Mellon University, Qatar
Stelvio Cimato	Università degli Studi di Milano, Italy
Mauro Conti	Università di Padova, Italy
Ernesto Damiani	Università degli Studi di Milan, Italy
Sabrina De Capitani di Vimercati	Università degli Studi di Milano, Italy

Josep Domingo-Ferrer	Rovira i Virgili University, Spain
William Fitzgerald	University College Cork, Ireland
Sara Foresti	Università degli Studi di Milano, Italy
Jerome Francois	University of Luxembourg, Luxembourg
Joaquin Garcia-Alfaro	Telecom SudParis, France
Stefanos Gritzalis	University of the Aegean, Greece
Olivier Heen	Technicolor, France
Wei Jiang	Missouri University of S&T, USA
Sokratis Katsikas	University of Piraeus, Greece
Florian Kerschbaum	SAP Research, France
Evangelos Kranakis	Carleton University, Canada
Marie Noelle Lepareux	Thales, France
Javier Lopez	University of Malaga, Spain
Giovanni Livraga	Università degli Studi di Milano, Italy
Wissam Mallouli	Montimage, France
Guillermo Navarro-Arribas	Autonomous University of Barcelona, Spain
Marie Nuadi	EADS-Cassidian, France
Andreas Pashalidis	K.U. Leuven, Belgium
Nicolas Prigent	Supelec, France
Yves Roudier	Eurecom, France
Thierry Sans	Carnegie Mellon University, Qatar
George Spanoudakis	City University London, UK
Radu State	University of Luxembourg, Luxembourg
Ari Takanen	Codenomicon, Finland
Bachar Wahbi	Percevio, France

Steering Committee

Ana-Rosa Cavalli	Telecom SudParis, France
Frédéric Cuppens	Telecom Bretagne, France
Nora Cuppens-Boulahia	Telecom Bretagne, France
Jean Leneutre	Telecom ParisTech, France
Yves Roudier	Eurecom, France

Contents

Keynote Address

Quantifying and Measuring Anonymity

Steven J. Murdoch(✉)

University of Cambridge Computer Laboratory, Cambridge, UK
Steven.Murdoch@cl.cam.ac.uk
http://www.cl.cam.ac.uk/~sjm217/

Abstract. The design of anonymous communication systems is a relatively new field, but the desire to quantify the security these systems offer has been an important topic of research since its beginning. In recent years, anonymous communication systems have evolved from obscure tools used by specialists to mass-market software used by millions of people. In many cases the users of these tools are depending on the anonymity offered to protect their liberty, or more. As such, it is of critical importance that not only can we quantify the anonymity these tools offer, but that the metrics used represent realistic expectations, can be communicated clearly, and the implementations actually offer the anonymity they promise. This paper will discuss how metrics, and the techniques used to measure them, have been developed for anonymous communication tools including low-latency networks and high-latency email systems.

1 Introduction

Anonymous communication systems seek to hide patterns visible in communications to obscure relationships between people and the activities they carry out, typically over the Internet. Such systems have become increasingly popular as a result of the Internet developing into an important tool in the support and promotion of human rights. Examples of uses include the publication of videos showing human rights abuses, journalists soliciting information on government corruption, and law enforcement agencies monitoring websites operated by organized crime.

In all these examples there are motivated individuals who would want to discover the identity of the users of the anonymous communication system. Therefore it is of critical importance that the level of protection that the anonymous communication system provides is well understood. Overestimating the level might result in users putting themselves at unacceptable amounts of risk; underestimating the level might result in users avoiding using a system unnecessarily.

The task of measuring the level of anonymity offered by anonymous communication tools is challenging particularly because of the narrow safety margins which they necessarily offer. A system operating perfectly can only hide the real sender or receiver of a message within the ranks of the users of that system. An attacker who wants to de-anonymise a user can often also take into account

J. Garcia-Alfaro et al. (Eds.): DPM 2013 and SETOP 2013, LNCS 8247, pp. 3–13, 2014.
DOI: 10.1007/978-3-642-54568-9_1,

auxiliary information collected through means other than monitoring the anonymous communication system.

For example, suppose a company discovers that a whistleblower has leaked documents, sent through an anonymous communication system, proving that management have authorised the bribing of government officials. If that anonymous communication system only had a million users that day, then there are at most a million candidates for who leaked the document. Intersecting the set of users of the system with the set of people who had access to the documents in question might leave only a handful of possibilities. Even a small amount of information disclosed by the anonymous communication system could leave the whistleblower singled out.

In contrast, encryption systems draw their strength from the large number of possible keys that could have been used to encrypt the information – far more than the number of users of the system. Adding to the key length imposes a linear cost to users of the system but increases the time needed to attack the system exponentially. As a result, modern encryption systems have a very large safety margin and so even serious weaknesses in encryption algorithms rarely have a practical effect on their security.

Therefore research on anonymous communication systems has focussed on improving security through increasing their number of users and decreasing the information disclosed to an observer. However, achieving either of these goals typically comes at a significant cost to users by reducing network capacity. As a result, it is not feasible to achieve the same safety margins that encryption systems offer and so it is important to develop ways to accurately measure the level of protection offered by anonymous communication systems. Then appropriate design choices can be made to provide the right trade-off between performance and security.

2 Email Mixes

One of the early applications of anonymous communication technology was to email. In a scheme proposed by Chaum [2] a user selects one or more "mixes" as a path through which his message should be sent. Messages are encrypted by a sender under multiple layers of public-key encryption (Fig. 1). Outside each layer of encryption is the address of the next mix, which allows messages to be routed. This mix can remove the next layer of encryption, and will find the address of the next mix in the path to which the message should be sent. Once the message reaches the last mix in the path, the plaintext of the message will be available along with the address of the ultimate destination of the message.

Each mix will see the immediate source of the message and the immediate destination. Therefore the first mix will know the sender's address but not the recipient's, and the last mix will know the recipient's address, but not the sender's. Similarly, someone observing messages flowing through the network will not be able to match incoming messages to outgoing messages based on the content because a decryption operation is carried out at each step which only

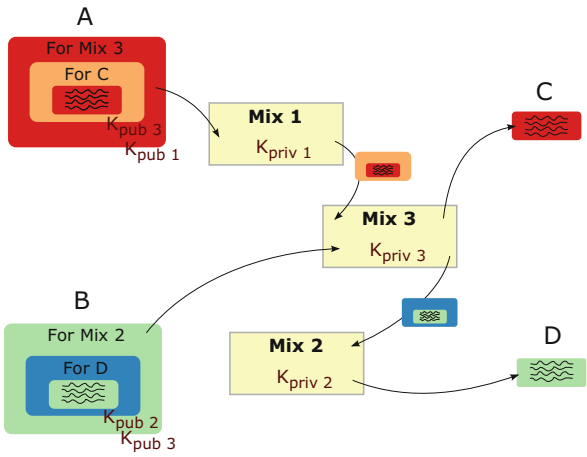

Fig. 1. A two-hop mix network. A is sending a message to C, via Mix 1 then Mix 3. B is sending a message to D via Mix 3 then Mix 2

a specific mix has the private key necessary to perform, and message lengths are fixed. Messages are also delayed at each mix, for a random period of time or until a particular number of messages have been received by a mix (or some combination of these) so as to complicate matching based on the time messages are sent and received.

In this way, the email mix network provides "unlinkability" [6] to messages because the attacker should not be able to link which messages entering the mix network correspond to which messages leaving the mix network. The mix network can also be seen to offer anonymity to its users – for each message leaving the network it should not be possible to establish its sender and for each message entering the network it should not be possible to establish its recipient. An attacker does however know a list of possible candidate senders for each message which leaves the network – the "sender anonymity set". Similarly there is a "recipient anonymity set" for each message sent.

2.1 Measuring Anonymity

Much of the research on email mixes has focussed on how to quantify the anonymity provided. Berthold et al. [1] proposed to simply count the size ("cardinality") of the anonymity set: a larger list of candidates for the true sender or receiver corresponds to better anonymity. By taking the logarithm of the set size, base 2, this quantity can be expressed in bits. An ideal anonymous communication system will have an anonymity set size of the number of users and the probability of each user being the sender or recipient of a particular message will be equal. Looking at the anonymity set as a probability distribution over possible senders/receivers of a message, the ideal anonymous communication system produces the uniform distribution.

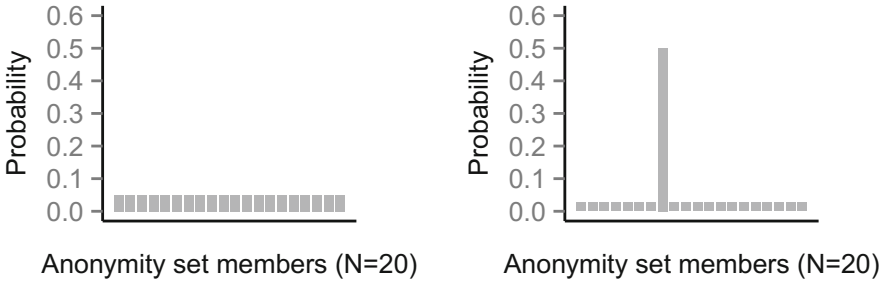

Fig. 2. Two possible distributions over a 20-element anonymity set. The left distribution is uniform (all elements at $\frac{1}{20}$); the right has one element at probability $\frac{1}{2}$ and the others at $\frac{1}{38}$

However real anonymous communication systems will not achieve this ideal. It is typically possible to distinguish senders from recipients by observing the direction of flow of data. Also by taking into account that it will be unlikely (for usability reasons) that mixes will delay messages for a long period of time, not every possible sender/recipient will be equally likely the true sender/recipient. In an extreme case an attacker may know that a single user may almost certainly be the sender of a message yet based on cardinality this system is indistinguishable from an ideal one of the same size, as shown in Fig. 2.

For this reason, other proposed metrics take into account the unevenness of the probability distribution. One such metric is the "degree of anonymity" proposed by Reiter et al. [7]. Although originally developed for analysing a system for anonymising web traffic it can equally be applied to email mixes. The 6 point scale is described in Table 1.

The degree of anonymity metric differentiates between the two anonymity set distributions of Fig. 2. The left graph shows that users are beyond suspicion whereas the right is barely probable innocence. For all reasonable purposes, the left graph corresponds to a better system so taking into account the unevenness of the distribution has produced a better metric, but ignoring the cardinality of the set has a weakness too.

For example, an anonymity set probability distribution over 101 senders, with the most likely sender having probability 0.01 and others probability 0.0099 offers possible innocence. Whereas an uniform anonymity set probability distribution over 4 senders has each sender assigned a probability of 0.25. Although the latter system has a better degree of anonymity, the probability of an attacker successfully identifying a user is much higher than the former.

It therefore follows that both cardinality and unevenness of distribution should be taken into account, and so Shannon entropy was proposed as a metric by Serjantov and Danezis [8]. Here, if the probability that user i was the true sender is p_i, and there are N members of the anonymity set, then the entropy

Table 1. The 6-point degree of anonymity scale

	Degree	Attacker observation
Best anonymity	Absolute privacy	No evidence whether or not a sender sent any message
	Beyond suspicion	A sender sent a message, but all senders are equally likely to have sent any message
	Probable innocence	A sender is no more likely to have been the originator of a message than to not have been
	Possible innocence	A sender has a nontrivial probability of not being the originator of a message
	Exposed	The originator has been identified
Worst anonymity	Provably exposed	The originator has been identified and the identity can be proven to others

of the anonymity set S is:

$$\mathrm{H}(S) = -\sum_{i=1}^{N} p_i \log_2 (p_i)$$

For the probability distributions in Fig. 2, the left distribution has entropy ≈ 4.32 bits (the same as the cardinality, in bits – $\log_2(20)$), but the right distribution only has entropy ≈ 3.12. The anonymity set discussed above, of 101 senders with one at probability 0.01 and others at 0.0099, gives entropy 6.66 bits (only 10^{-5} % less than the entropy of the uniform distribution over 101 senders). Whereas the uniform distribution over 4 senders is 2. We can see that entropy takes into account both cardinality and unevenness, and also gives similar values to similar distributions, but it is still possible to find examples which raise the question of whether entropy is the best metric.

For example, in Fig. 3 the two very different distributions have the same entropy. However, from the perspective of an attacker the anonymity might be very different. The de-anonymisation of communications is seldom used as an end in itself, but rather to guide further investigation. An attacker analysing the left distribution would need to investigate 10 senders before getting a 50 % probability of having found the right sender. In contrast the attacker could achieve the same goal with the right distribution after trying only one user.

One way of differentiating between the two distributions is to note that the number of users is rarely under direct control of the system designer so a reasonable metric could examine the ratio between the security of the ideal system for a given user base to the actual security achieved for the same user base. This metric was proposed as the "degree of anonymity" by Diaz et al. [3], but to differentiate from the Crowds degree in Table 1 we will use the term "nor-