Nikita Borisov
Philippe Golle (Eds.)

LNCS 4776

Privacy Enhancing Technologies

7th International Symposium, PET 2007
Ottawa, Canada, June 2007
Revised Selected Papers

Springer

Nikita Borisov Philippe Golle (Eds.)

Privacy Enhancing Technologies

7th International Symposium, PET 2007
Ottawa, Canada, June 20-22, 2007
Revised Selected Papers

Springer

Volume Editors

Nikita Borisov
University of Illinois at Urbana-Champaign
Department of Electrical and Computer Engineering
1308 West Main Street, Urbana, IL 61801-2307, USA
E-mail: nikita@uiuc.edu

Philippe Golle
Palo Alto Research Center
3333 Coyote Hill Road, Palo Alto, CA 94304, USA
E-mail: Philippe.Golle@parc.com

Library of Congress Control Number: 2007938055

CR Subject Classification (1998): E.3, C.2, D.4.6, K.6.5, K.4, H.3, H.4

LNCS Sublibrary: SL 4 – Security and Cryptology

ISSN 0302-9743
ISBN-10 3-540-75550-0 Springer Berlin Heidelberg New York
ISBN-13 978-3-540-75550-0 Springer Berlin Heidelberg New York

Springer-Verlag Berlin Heidelberg holds the exclusive right of distribution and reproduction of this work, for a period of three years starting from the date of publication.

Springer is a part of Springer Science+Business Media

springer.com

Printed in Germany

Typesetting: Camera-ready by author, data conversion by Scientific Publishing Services, Chennai, India
Printed on acid-free paper SPIN: 12171590 06/3180 5 4 3 2 1 0

Lecture Notes in Computer Science 4776

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

Editorial Board

David Hutchison
Lancaster University, UK

Takeo Kanade
Carnegie Mellon University, Pittsburgh, PA, USA

Josef Kittler
University of Surrey, Guildford, UK

Jon M. Kleinberg
Cornell University, Ithaca, NY, USA

Friedemann Mattern
ETH Zurich, Switzerland

John C. Mitchell
Stanford University, CA, USA

Moni Naor
Weizmann Institute of Science, Rehovot, Israel

Oscar Nierstrasz
University of Bern, Switzerland

C. Pandu Rangan
Indian Institute of Technology, Madras, India

Bernhard Steffen
University of Dortmund, Germany

Madhu Sudan
Massachusetts Institute of Technology, MA, USA

Demetri Terzopoulos
University of California, Los Angeles, CA, USA

Doug Tygar
University of California, Berkeley, CA, USA

Moshe Y. Vardi
Rice University, Houston, TX, USA

Gerhard Weikum
Max-Planck Institute of Computer Science, Saarbruecken, Germany

Lecture Notes in Computer Science

Sublibrary 4: Security and Cryptology

Vol. 4784: W. Susilo, J.K. Liu, Y. Mu (Eds.), Provable Security. X, 237 pages. 2007.

Vol. 4779: J.A. Garay, A.K. Lenstra, M. Mambo, R. Peralta (Eds.), Information Security. XIII, 437 pages. 2007.

Vol. 4776: N. Borisov, P. Golle (Eds.), Privacy Enhancing Technologies. X, 273 pages. 2007.

Vol. 4752: A. Miyaji, H. Kikuchi, K. Rannenberg (Eds.), Advances in Information and Computer Security. XIII, 460 pages. 2007.

Vol. 4734: J. Biskup, J. López (Eds.), Computer Security – ESORICS 2007. XIV, 628 pages. 2007.

Vol. 4727: P. Paillier, I. Verbauwhede (Eds.), Cryptographic Hardware and Embedded Systems - CHES 2007. XIV, 468 pages. 2007.

Vol. 4691: T. Dimitrakos, F. Martinelli, P.Y.A. Ryan, S. Schneider (Eds.), Formal Aspects in Security and Trust. VIII, 285 pages. 2007.

Vol. 4677: A. Aldini, R. Gorrieri (Eds.), Foundations of Security Analysis and Design IV. VII, 325 pages. 2007.

Vol. 4657: C. Lambrinoudakis, G. Pernul, A M. Tjoa (Eds.), Trust and Privacy in Digital Business. XIII, 291 pages. 2007.

Vol. 4637: C. Kruegel, R. Lippmann, A. Clark (Eds.), Recent Advances in Intrusion Detection. XII, 337 pages. 2007.

Vol. 4622: A. Menezes (Ed.), Advances in Cryptology - CRYPTO 2007. XIV, 631 pages. 2007.

Vol. 4593: A. Biryukov (Ed.), Fast Software Encryption. XI, 467 pages. 2007.

Vol. 4586: J. Pieprzyk, H. Ghodosi, E. Dawson (Eds.), Information Security and Privacy. XIV, 476 pages. 2007.

Vol. 4582: J. López, P. Samarati, J.L. Ferrer (Eds.), Public Key Infrastructure. XI, 375 pages. 2007.

Vol. 4579: B.M. Hämmerli, R. Sommer (Eds.), Detection of Intrusions and Malware, and Vulnerability Assessment. X, 251 pages. 2007.

Vol. 4575: T. Takagi, T. Okamoto, E. Okamoto, T. Okamoto (Eds.), Pairing-Based Cryptography – Pairing 2007. XI, 408 pages. 2007.

Vol. 4521: J. Katz, M. Yung (Eds.), Applied Cryptography and Network Security. XIII, 498 pages. 2007.

Vol. 4515: M. Naor (Ed.), Advances in Cryptology - EUROCRYPT 2007. XIII, 591 pages. 2007.

Vol. 4499: Y.Q. Shi (Ed.), Transactions on Data Hiding and Multimedia Security II. IX, 117 pages. 2007.

Vol. 4464: E. Dawson, D.S. Wong (Eds.), Information Security Practice and Experience. XIII, 361 pages. 2007.

Vol. 4462: D. Sauveron, K. Markantonakis, A. Bilas, J.-J. Quisquater (Eds.), Information Security Theory and Practices. XII, 255 pages. 2007.

Vol. 4450: T. Okamoto, X. Wang (Eds.), Public Key Cryptography – PKC 2007. XIII, 491 pages. 2007.

Vol. 4437: J.L. Camenisch, C.S. Collberg, N.F. Johnson, P. Sallee (Eds.), Information Hiding. VIII, 389 pages. 2007.

Vol. 4392: S.P. Vadhan (Ed.), Theory of Cryptography. XI, 595 pages. 2007.

Vol. 4377: M. Abe (Ed.), Topics in Cryptology – CT-RSA 2007. XI, 403 pages. 2006.

Vol. 4356: E. Biham, A.M. Youssef (Eds.), Selected Areas in Cryptography. XI, 395 pages. 2007.

Vol. 4341: P.Q. Nguyên (Ed.), Progress in Cryptology - VIETCRYPT 2006. XI, 385 pages. 2006.

Vol. 4332: A. Bagchi, V. Atluri (Eds.), Information Systems Security. XV, 382 pages. 2006.

Vol. 4329: R. Barua, T. Lange (Eds.), Progress in Cryptology - INDOCRYPT 2006. X, 454 pages. 2006.

Vol. 4318: H. Lipmaa, M. Yung, D. Lin (Eds.), Information Security and Cryptology. XI, 305 pages. 2006.

Vol. 4307: P. Ning, S. Qing, N. Li (Eds.), Information and Communications Security. XIV, 558 pages. 2006.

Vol. 4301: D. Pointcheval, Y. Mu, K. Chen (Eds.), Cryptology and Network Security. XIII, 381 pages. 2006.

Vol. 4300: Y.Q. Shi (Ed.), Transactions on Data Hiding and Multimedia Security I. IX, 139 pages. 2006.

Vol. 4298: J.K. Lee, O. Yi, M. Yung (Eds.), Information Security Applications. XIV, 406 pages. 2007.

Vol. 4296: M.S. Rhee, B. Lee (Eds.), Information Security and Cryptology – ICISC 2006. XIII, 358 pages. 2006.

Vol. 4284: X. Lai, K. Chen (Eds.), Advances in Cryptology – ASIACRYPT 2006. XIV, 468 pages. 2006.

Vol. 4283: Y.Q. Shi, B. Jeon (Eds.), Digital Watermarking. XII, 474 pages. 2006.

Vol. 4266: H. Yoshiura, K. Sakurai, K. Rannenberg, Y. Murayama, S.-i. Kawamura (Eds.), Advances in Information and Computer Security. XIII, 438 pages. 2006.

Vol. 4258: G. Danezis, P. Golle (Eds.), Privacy Enhancing Technologies. VIII, 431 pages. 2006.

Vol. 4249: L. Goubin, M. Matsui (Eds.), Cryptographic Hardware and Embedded Systems - CHES 2006. XII, 462 pages. 2006.

Vol. 4237: H. Leitold, E.P. Markatos (Eds.), Communications and Multimedia Security. XII, 253 pages. 2006.

此为试读,需要完整PDF请访问: www.ertongbook.com

Vol. 4236: L. Breveglieri, I. Koren, D. Naccache, J.-P. Seifert (Eds.), Fault Diagnosis and Tolerance in Cryptography. XIII, 253 pages. 2006.

Vol. 4219: D. Zamboni, C. Krügel (Eds.), Recent Advances in Intrusion Detection. XII, 331 pages. 2006.

Vol. 4189: D. Gollmann, J. Meier, A. Sabelfeld (Eds.), Computer Security – ESORICS 2006. XI, 548 pages. 2006.

Vol. 4176: S.K. Katsikas, J. López, M. Backes, S. Gritzalis, B. Preneel (Eds.), Information Security. XIV, 548 pages. 2006.

Vol. 4117: C. Dwork (Ed.), Advances in Cryptology - CRYPTO 2006. XIII, 621 pages. 2006.

Vol. 4116: R. De Prisco, M. Yung (Eds.), Security and Cryptography for Networks. XI, 366 pages. 2006.

Vol. 4107: G. Di Crescenzo, A. Rubin (Eds.), Financial Cryptography and Data Security. XI, 327 pages. 2006.

Vol. 4083: S. Fischer-Hübner, S. Furnell, C. Lambrinoudakis (Eds.), Trust and Privacy in Digital Business. XIII, 243 pages. 2006.

Vol. 4064: R. Büschkes, P. Laskov (Eds.), Detection of Intrusions and Malware & Vulnerability Assessment. X, 195 pages. 2006.

Vol. 4058: L.M. Batten, R. Safavi-Naini (Eds.), Information Security and Privacy. XII, 446 pages. 2006.

Vol. 4047: M.J.B. Robshaw (Ed.), Fast Software Encryption. XI, 434 pages. 2006.

Vol. 4043: A.S. Atzeni, A. Lioy (Eds.), Public Key Infrastructure. XI, 261 pages. 2006.

Vol. 4004: S. Vaudenay (Ed.), Advances in Cryptology - EUROCRYPT 2006. XIV, 613 pages. 2006.

Vol. 3995: G. Müller (Ed.), Emerging Trends in Information and Communication Security. XX, 524 pages. 2006.

Vol. 3989: J. Zhou, M. Yung, F. Bao (Eds.), Applied Cryptography and Network Security. XIV, 488 pages. 2006.

Vol. 3969: Ø. Ytrehus (Ed.), Coding and Cryptography. XI, 443 pages. 2006.

Vol. 3958: M. Yung, Y. Dodis, A. Kiayias, T.G. Malkin (Eds.), Public Key Cryptography - PKC 2006. XIV, 543 pages. 2006.

Vol. 3957: B. Christianson, B. Crispo, J.A. Malcolm, M. Roe (Eds.), Security Protocols. IX, 325 pages. 2006.

Vol. 3956: G. Barthe, B. Grégoire, M. Huisman, J.-L. Lanet (Eds.), Construction and Analysis of Safe, Secure, and Interoperable Smart Devices. IX, 175 pages. 2006.

Vol. 3935: D.H. Won, S. Kim (Eds.), Information Security and Cryptology - ICISC 2005. XIV, 458 pages. 2006.

Vol. 3934: J.A. Clark, R.F. Paige, F.A.C. Polack, P.J. Brooke (Eds.), Security in Pervasive Computing. X, 243 pages. 2006.

Vol. 3928: J. Domingo-Ferrer, J. Posegga, D. Schreckling (Eds.), Smart Card Research and Advanced Applications. XI, 359 pages. 2006.

Vol. 3919: R. Safavi-Naini, M. Yung (Eds.), Digital Rights Management. XI, 357 pages. 2006.

Vol. 3903: K. Chen, R. Deng, X. Lai, J. Zhou (Eds.), Information Security Practice and Experience. XIV, 392 pages. 2006.

Vol. 3897: B. Preneel, S. Tavares (Eds.), Selected Areas in Cryptography. XI, 371 pages. 2006.

Vol. 3876: S. Halevi, T. Rabin (Eds.), Theory of Cryptography. XI, 617 pages. 2006.

Vol. 3866: T. Dimitrakos, F. Martinelli, P.Y.A. Ryan, S. Schneider (Eds.), Formal Aspects in Security and Trust. X, 259 pages. 2006.

Vol. 3860: D. Pointcheval (Ed.), Topics in Cryptology – CT-RSA 2006. XI, 365 pages. 2006.

Vol. 3858: A. Valdes, D. Zamboni (Eds.), Recent Advances in Intrusion Detection. X, 351 pages. 2006.

Vol. 3856: G. Danezis, D. Martin (Eds.), Privacy Enhancing Technologies. VIII, 273 pages. 2006.

Vol. 3786: J.-S. Song, T. Kwon, M. Yung (Eds.), Information Security Applications. XI, 378 pages. 2006.

Vol. 3108: H. Wang, J. Pieprzyk, V. Varadharajan (Eds.), Information Security and Privacy. XII, 494 pages. 2004.

Vol. 2951: M. Naor (Ed.), Theory of Cryptography. XI, 523 pages. 2004.

Vol. 2742: R.N. Wright (Ed.), Financial Cryptography. VIII, 321 pages. 2003.

Foreword

The PET community has grown in size considerably since the first PET workshop was held in 2000. With this growth came an increase in the number and quality of submissions. PET has become a premier venue for publishing original research on privacy enhancing technologies, and the current acceptance ratio puts PET in the same league as other highly selective security and privacy venues. To appropriately reflect this evolution, the PET workshop is changing its name to the *Privacy Enhancing Technologies Symposium.*

PET 2007 was held at the University of Ottawa, Canada, on June 20–22, 2007. We received 84 full-paper submissions, of which 16 were selected for presentation at the symposium. PET also included a keynote address and two panel discussions. PET was once again collocated with the IAVoSS Workshop on Trustworthy Elections (WOTE 2007), with a full day of plenary sessions. All participants were free to attend sessions from both events.

The program chairs would like to thank, first of all, the authors, speakers, and panelists for their contribution to the content of the workshop. We would also like to thank the program committee for their hard work of a month of reviews and two more weeks of intense discussions, helping to ensure a program of high scientific quality. As well, we would like to acknowledge the contribution of the external reviewers, who assisted the program committee with the reviews. A special thanks is due to the designers of the Websubmission and Webreview software at K.U. Leuven for allowing us to use their software to help with the selection process, and to Thomas Herlea for his help in getting the software up and running.

Our general chair, Carlisle Adams, did an outstanding job taking care of the local arrangements and making sure the symposium ran smoothly. We also would like to thank Jeremy Clark for designing and maintaining the PET 2007 Website. We are very grateful to Josh Benaloh, the chair of WOTE 2007, for his help in coordinating the two events. Finally, PET 2007 was made possible, and more affordable, thanks to our sponsors: Microsoft, ORNEC, Bell Privacy Centre of Excellence, PGP Corporation, and Google. We are extremely grateful for their generous support.

The Award for Outstanding Research in Privacy Enhancing Technologies was given this year to Stephen C. Bono, Matthew Green, Adam Stubblefield, Ari Juels, Aviel D. Rubin, and Michael Szydlo for their paper "Security Analysis of a Cryptographically-Enabled RFID Device." The award is sponsored by Microsoft and by the Office of the Information and Privacy Commissioner of Ontario and the winners were selected through an independent prize committee headed by George Danezis to whom we are thankful.

July 2007

Nikita Borisov
Philippe Golle

Organization

Organizers

General Chair	Carlisle Adams (University of Ottawa, Canada)
Program Chairs	Nikita Borisov (University of Illinois at Urbana-Champaign, USA)
	Philippe Golle (Palo Alto Research Center, USA)
PET Prize	George Danezis (K.U. Leuven, Belgium)
Stipends	Roger Dingledine (The Tor Project, USA)

Program Committee

Alessandro Acquisti (Carnegie Mellon University, USA)
Mikhail Atallah (Purdue University, USA)
Michael Backes (Saarland University, Germany)
Alastair Beresford (University of Cambridge, UK)
Jean Camp (Indiana University, USA)
George Danezis (K.U. Leuven, Belgium)
Claudia Díaz (K.U. Leuven, Belgium)
Roger Dingledine (The Tor Project, USA)
Cynthia Dwork (Microsoft Research, USA)
Simson Garfinkel (Harvard University, USA)
Ian Goldberg (University of Waterloo, Canada)
Susan Hohenberger (Johns Hopkins University, USA)
Dennis Kügler (Federal Office for Information Security, Germany)
Bradley Malin (Vanderbilt University, USA)
David Martin (University of Massachusetts at Lowell, USA)
Nick Mathewson (The Tor Project, USA)
David Molnar (University of California at Berkeley, USA)
Steven Murdoch (University of Cambridge, UK)
Andreas Pfitzmann (Dresden University of Technology, Germany)
Mike Reiter (University of North Carolina at Chapel Hill, USA)
Andrei Serjantov (The Free Haven Project, UK)
Vitaly Shmatikov (University of Texas at Austin, USA)
Paul Syverson (Naval Research Laboratory, USA)
Matthew Wright (University of Texas at Arlington, USA)

External Reviewers

Mike Bergmann
Alexander Böttcher
Katrin Borcea-Pfitzmann
Sebastian Clauß
Richard Clayton
Markus Duermuth
David Evans
Anna Lisa Ferrara
Elke Franz
Bikas Gurung
Thomas Heydt-Benjamin
Yong Ho Hwang
Ponnurangam Kumaraguru
Haim Levkowitz
Benyuan Liu
Matteo Maffei
Sasha Romanosky
Sandra Steinbrecher
Carmela Troncoso
Lasse Øverlier

Sponsors

Microsoft
ORNEC
Bell Privacy Centre of Excellence
PGP Corporation
Google

Table of Contents

此为试读,需要完整PDF请访问: www.ertongbook.com

Attacking Unlinkability: The Importance of Context

Matthias Franz[1], Bernd Meyer[1], and Andreas Pashalidis[2]

[1] Siemens AG, Corporate Technology,
Otto-Hahn-Ring 6, 81739 München, Germany
{matthias.franz,bernd.meyer}@siemens.com
[2] NEC Europe Ltd, Network Laboratories
Kurfürsten-Anlage 36, 69115 Heidelberg, Germany
andreas.pashalidis@netlab.nec.de

Abstract. A system that protects the unlinkability of certain data items (e. g. identifiers of communication partners, messages, pseudonyms, transactions, votes) does not leak information that would enable an adversary to link these items. The adversary could, however, take advantage of hints from the context in which the system operates. In this paper, we introduce a new metric that enables one to quantify the (un)linkability of the data items and, based on this, we consider the effect of some simple contextual hints.

1 Introduction

A number of privacy-preserving systems, such as mix networks, anonymous credential systems, and secret voting schemes, protect the unlinkability of certain data items of interest. Mix networks, in particular, protect the unlinkability of the messages that enter the network with respect to their recipients, the messages that leave the network with respect to their senders, and, hence, the identifiers of communicating parties with respect to communication sessions. Since their introduction [9], a number of different mix network variants have been proposed (see, for example, [4,19,26,33,34]), some of which have also been implemented and deployed. Anonymous credentials, on the other hand, protect the unlinkability of the pseudonyms and the transactions with respect to the users they correspond to. Since their introduction into the digital world [10], a number of anonymous credential systems have been proposed (see, for example, [7,8,11,12,13,14,29,32,38]). Secret voting schemes protect the unlinkability of votes with respect to the users who cast them. Such schemes have evolved from ostracism [24] to sophisticated cryptosystems; for an overview of the current state of the art the reader is referred to [1].

The problem of analysing how well the above types of system protect unlinkability has received some attention during recent years. The focus of most works is, however, on mix networks (see, for example, [2,15,16,25,27,30]). This is not surprising since mix networks provide the basis for anonymous communication

N. Borisov and P. Golle (Eds.): PET 2007, LNCS 4776, pp. 1–16, 2007.

and are, as such, necessary for preserving privacy in a number of settings, including the setting of anonymous credentials [17] and, sometimes, the setting of voting systems (see, for example, [6]).

An adversary that wishes to link the protected items may use information that is leaked by the system during its operation, or hints from the environment in which the system operates. In contrast to existing literature, the focus of this paper is on the latter. That is, we study a number of simple contextual hints and their effect on unlinkability. Our results apply to *all* types of unlinkability-protecting system, including mix networks, anonymous credentials, and secret voting schemes. The rest of the paper is organised as follows. Section 2 introduces the metric for unlinkability that is used throughout the paper. Section 3 examines seven different types of hint and their effect on unlinkability. Finally, Section 4 concludes.

2 Measuring Unlinkability

Consider a set of elements A and a partition $\pi \vdash A$ of that set. Note that we do not distinguish between π and the equivalence relation it defines. In the sequel, we write $a_1 \equiv_\pi a_2$ if the elements $a_1, a_2 \in A$ lie in the same equivalence class of π, and $a_1 \not\equiv_\pi a_2$ otherwise. Let $\tau \vdash A$ denote a 'target' partition, chosen uniformly at random. We use entropy as a metric for unlinkability. That is, the unlinkability of the elements in a set A against an adversary $\mathcal{A}$ is defined as

$$\mathcal{U}_A(\mathcal{A}) = -\sum_{\pi \in \Pi} \Pr(\pi = \tau) \log_2(\Pr(\pi = \tau)),$$

where $\Pi = \{P : P \vdash A\}$ denotes the set of partitions of A and $\Pr(\pi = \tau)$ denotes, in $\mathcal{A}$'s view, the probability that π is the target partition τ. We further define the *degree* of unlinkability of the elements in A against an adversary $\mathcal{A}_H$ with access to a hint H about τ as

$$\mathcal{D}_A(\mathcal{A}_H) = \frac{\mathcal{U}_A(\mathcal{A}_H)}{\mathcal{U}_A(\mathcal{A}_\emptyset)},$$

where $\mathcal{A}_\emptyset$ denotes the adversary without any hints. That is, $\mathcal{A}_\emptyset$ knows A but has no information about τ. The set of candidate partitions for $\mathcal{A}_\emptyset$ is therefore $\Pi_A(\mathcal{A}_\emptyset) = \Pi$, i. e. the set of all partitions of A. The number $|\Pi_A(\mathcal{A}_\emptyset)| = B_{|A|}$ of such partitions, a Bell number [3,35], is given by the recursive formula

$$B_{n+1} = \sum_{k=0}^{n} \binom{n}{k} B_k \tag{1}$$

where $B_0 = 1$.[1] Since τ is chosen uniformly at random, the unlinkability of the elements in A is therefore at its maximum, i. e. $\mathcal{U}_A(\mathcal{A}_\emptyset) = \log_2(B_{|A|})$ bits. This

[1] The first few Bell numbers are 1, 1, 2, 5, 15, 52, 203, 877, 4140, 21147.

is the best case from a privacy point of view: all partitions of A are equally likely to be the target partition τ.

Remark 1: In the setting of unlinkability-protecting systems, the goal of the adversary is to identify a target partition from an 'anonymity set' of candidate partitions. The fact that the information-theoretic metric we use for unlinkability is identical to the metric introduced for anonymity in [18,36], is therefore natural.

Remark 2: $\mathcal{U}_A$ is a measure of the information that is contained in the probability distribution that the adversary assigns to the set of all partitions of A. Since we assume that τ is selected uniformly at random, this distribution is, a priori, uniform. However, a hint may enable the adversary to change his view such that, a posteriori, some partitions are more likely than others. The hints we consider in this paper enable the adversary to exclude a number of candidate partitions (i.e. to reduce the size of the 'anonymity set') while the remaining partitions remain equally likely.

Example: Consider an anonymous help line where a clerk offers advice over the telephone. Suppose that, one day, the clerk receives four calls, denoted $A = \{\lambda_1, \lambda_2, \lambda_3, \lambda_4\}$. Without any additional information, all $B_4 = 15$ partitions of A constitute valid ways to link these calls. Since without any additional information all these partitions are equally likely, the unlinkability of the calls is, in this case, $\log_2(15) \simeq 3.9$ bits, and the degree of unlinkability is $\log_2(15)/\log_2(15) = 1$.

The clerk, however, has some additional information: he realised that the calls λ_1 and λ_2 were made by men, and that the calls λ_3 and λ_4 by women (however, the clerk does not know whether or not the same person called twice). This hint effectively rules out all partitions where λ_1 or λ_2 appears in the same equivalence class as λ_2 or λ_4. In particular, only four partitions remain valid, namely $\{(\lambda_1, \lambda_2), (\lambda_3, \lambda_4)\}$, $\{(\lambda_1, \lambda_2), (\lambda_3), (\lambda_4)\}$, $\{(\lambda_1), (\lambda_2), (\lambda_3, \lambda_4)\}$, and $\{(\lambda_1), (\lambda_2), (\lambda_3), (\lambda_4)\}$. Since these four partitions are equally likely, the unlinkability of the calls is, in this case, $\log_2(4) = 2$ bits, and the degree of unlinkability is $\log_2(2)/\log_2(15) \simeq 0.52$.

3 The Importance of Context

In this section, we examine seven types of hint that an adversary may obtain from the operational context of the system. In particular, we examine hints that reveal to the adversary (a) the number of equivalence classes in τ, (b) the cardinality of equivalence classes in τ, (c) the fact that all equivalence classes in τ have a given cardinality, (d) a 'reference partition' the equivalence classes of which have exactly one representative in each equivalence class in τ, (e) a set of element pairs that are equivalent in τ, (f) a set of element pairs that are not equivalent in τ, and (g) a combination of (e) and (f).

3.1 The Number of Equivalence Classes

Consider an adversary $\mathcal{A}_{H_1}$ with a hint $H_1 = (\alpha)$, where $\alpha \in \mathbb{N}$ and $1 \leq \alpha \leq |A|$, that reveals how many equivalence classes τ has. $\mathcal{A}_{H_1}$ can restrict its attention

to $\Pi_A(\mathcal{A}_{H_1}) = \{P : P \vdash A, |P| = \alpha\}$, i. e. the partitions that divide A into exactly α equivalence classes. The number of such partitions, which is a Stirling number of the second kind [22], is given by

$$|\Pi_A(\mathcal{A}_{H_1})| = \frac{1}{\alpha!}\sum_{k=0}^{\alpha}(-1)^k\binom{\alpha}{k}(\alpha-k)^{|A|}.$$

Since τ is chosen uniformly at random, the unlinkability of the elements in A is $\mathcal{U}_A(\mathcal{A}_{H_1}) = \log_2(|\Pi_A(\mathcal{A}_{H_1})|)$ bits. Figure 1 shows the degree of unlinkability $\mathcal{D}_A(\mathcal{A}_{H_1})$ as a function of $|A|$.

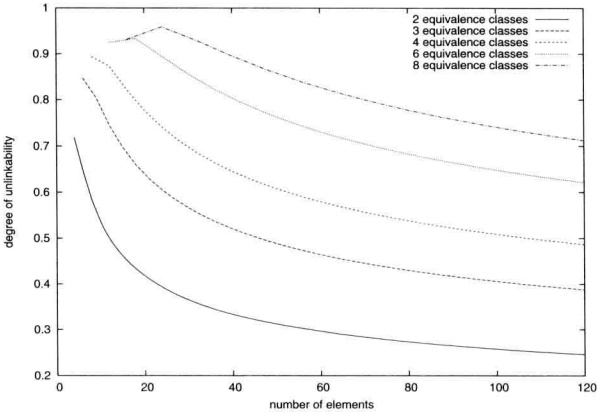

Fig. 1. Degree of unlinkability $\mathcal{D}_A(\mathcal{A}_{H_1})$ of elements in a set A as a function of $|A|$, if it is known that they must be divided into α equivalence classes

How to obtain this hint: The number α typically is the number of users in a system. In the setting of mix networks, this number may be known to the operator of the network if users are required to register themselves or pay a fee. Otherwise, obtaining such a hint may be tricky due to the possibility of sybil attacks [20]. Whether or not it is straightforward to obtain this hint in the setting of anonymous credentials depends on the application. In the case of cash, for example, the financial institution is very likely to know how many users participate in the system. Similarly, in the case of demographic or personal credentials (such as age certificates or driving licences), the issuing authority is also likely to know the number of users in the system. In the setting of secret voting, there exist multiple ways to obtain the number of voters. The number of cast ballots, for example, may be conclusive about the number of voters.

3.2 The Cardinality of Equivalence Classes

Consider an adversary $\mathcal{A}_{H_2}$ with a hint $H_2 = (\beta_1, \beta_2, \ldots, \beta_\alpha)$, where $\sum_{i=1}^{\alpha}\beta_i = |A|$ and $1 < \alpha < |A|$, that reveals the sizes of the equivalence classes in τ. That is,

if $\tau = \{T_1, T_2, \ldots, T_\alpha\} \vdash A$, H_2 reveals that $|T_1| = \beta_1$, $|T_2| = \beta_2$, and so on. $\mathcal{A}_{H_2}$ can restrict its attention to $\Pi_A(\mathcal{A}_{H_2}) = \{P : P = \{T_1, T_2, \ldots, T_\alpha\} \vdash A, \forall 1 \leq i \leq \alpha, |T_i| = \beta_i\}$, i. e. the partitions that divide A into exactly α equivalence classes with cardinalities $\beta_1, \beta_2, \ldots, \beta_\alpha$. The number of such partitions is given by

$$|\Pi_A(\mathcal{A}_{H_2})| = \frac{|A|!}{\prod_{i=1}^{\alpha}(\beta_i!)\prod_{i=1}^{|A|}(\gamma_i!)} \tag{2}$$

where, for all $1 \leq i \leq |A|$, $\gamma_i = |\{\beta \in H_2 : \beta = i\}|$ (for a proof see Appendix B). That is, γ_i is the number of equivalence classes in τ that have cardinality i. Since τ is chosen uniformly at random, the unlinkability of the elements in A is $\mathcal{U}_A(\mathcal{A}_{H_2}) = \log_2(|\Pi_A(\mathcal{A}_{H_2})|)$ bits. It is perhaps worth noting that there exist hints of type H_2 which do not reveal any information as to whether any two given elements are equivalent or not. This is in contrast to what is claimed in [37] (see Appendix A).

As a special case, consider an adversary $\mathcal{A}_{H_3}$ with a hint $H_3 = (\alpha)$, where $\alpha \in \mathbb{N}$ divides $|A|$, that reveals the fact that τ has α equivalence classes *of the same cardinality* $|A|/\alpha$. $\mathcal{A}_{H_3}$ can restrict its attention to $\Pi_A(\mathcal{A}_{H_3}) = \{P : P \vdash A, \forall p \in P, |p| = |A|/\alpha\}$, i. e. the partitions that divide A into exactly α equivalence classes of equal cardinality $|A|/\alpha$. The number of such partitions is given by

$$|\Pi_A(\mathcal{A}_{H_3})| = \frac{|A|!}{\alpha!((|A|/\alpha)!)^\alpha} \tag{3}$$

(for a proof see Appendix B). Since τ is chosen uniformly at random, the unlinkability of the elements in A is $\mathcal{U}_A(\mathcal{A}_{H_3}) = \log_2(|\Pi_A(\mathcal{A}_{H_3})|)$ bits. Figure 2 shows the degree of unlinkability $\mathcal{D}_A(\mathcal{A}_{H_3})$ as a function of $|A|$.

How to obtain this hint:[2] In the setting of mix networks, this hint may be obtained if it is known how many messages each user sends in each session. In the setting of anonymous credentials, it is possible to obtain this hint if it is known how many pseudonyms each user has. In the setting of secret voting, this hint may be obtained if it is known how many ballots each user has cast.

3.3 A Reference Partition

Consider an adversary $\mathcal{A}_{H_4}$ with a hint $H_4 = (\rho)$, consisting of a 'reference partition' $\rho = \{R_1, R_2, \ldots, R_{|A|/\alpha}\} \vdash A$ such that, for all $1 \leq i \leq |A|/\alpha$, $|R_i| = \alpha$ (note that α divides $|A|$), and that reveals the fact that each of the equivalence classes of τ contains exactly one element from R_i. $\mathcal{A}_{H_4}$ can restrict its attention to $\Pi_A(\mathcal{A}_{H_4}) = \{P : P \vdash A, P \text{ is a transversal of } \rho\}$, i. e. the partitions that divide A into α equivalence classes of equal cardinality $|A|/\alpha$, where each class contains exactly one element from each of $R_1, R_2, \ldots, R_{|A|/\alpha}$. The number of such partitions is given by

$$|\Pi_A(\mathcal{A}_{H_4})| = (\alpha!)^{(|A|/\alpha)-1} \tag{4}$$

[2] This paragraph refers to hints of both type H_2 and H_3.

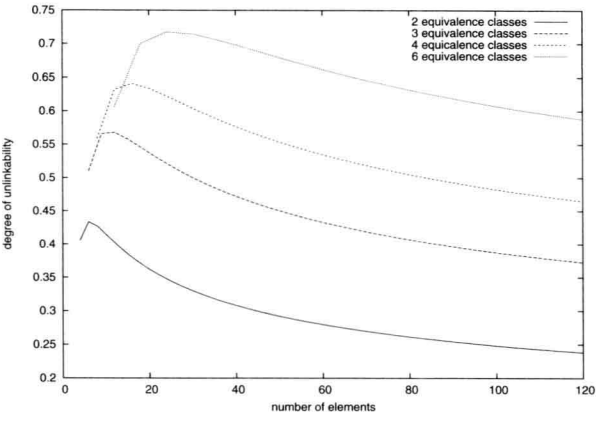

Fig. 2. Degree of unlinkability $\mathcal{D}_A(\mathcal{A}_{H_3})$ of elements in a set A as a function of $|A|$, if it is known that they must be divided into α equivalence classes of equal cardinality $|A|/\alpha$

(for a proof see Appendix C). Since τ is chosen uniformly at random, the unlinkability of the elements in A is $\mathcal{U}_A(\mathcal{A}_{H_4}) = \log_2(|\Pi_A(\mathcal{A}_{H_4})|)$ bits. Figure 3 shows the degree of unlinkability $\mathcal{D}_A(\mathcal{A}_{H_4})$ as a function of $|A|$.

How to obtain this hint: In the setting of mix networks this hint may be obtained if each of the α users sends exactly one message through the network in β communication sessions. An adversary that wishes to divide the set of messages that leave the network (there are $\alpha \cdot \beta$ of them) into α subsets of equal cardinality β, such that each subset contains the messages sent by a single user,

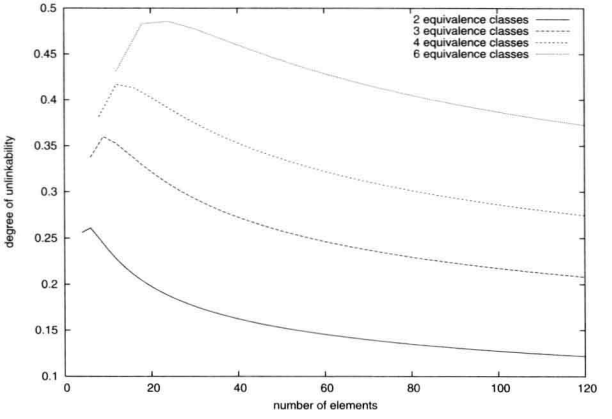

Fig. 3. Degree of unlinkability $\mathcal{D}_A(\mathcal{A}_{H_4})$ of elements in a set A as a function of $|A|$, if it is known that they must be divided into α equivalence classes of equal cardinality $|A|/\alpha$, such that each class contains exactly one element from each equivalence class of a given partition

此为试读,需要完整PDF请访问: www.ertongbook.com