

Bruce Christianson
Bruno Crispo
James A. Malcolm
Michael Roe (Eds.)

LNCS 3957

Security Protocols

12th International Workshop
Cambridge, UK, April 2004
Revised Selected Papers



Springer

Bruce Christianson Bruno Crispo
James A. Malcolm Michael Roe (Eds.)

Security Protocols

12th International Workshop
Cambridge, UK, April 26-28, 2004
Revised Selected Papers

Volume Editors

Bruce Christianson
University of Hertfordshire
Computer Science Department
Hatfield AL10 9AB, UK
E-mail: b.christianson@herts.ac.uk

Bruno Crispo
Vrije Universiteit
Department of Computer Science
De Boelelaan 1081, 1081 HV Amsterdam, The Netherlands
E-mail: crispo@cs.vu.nl

James A. Malcolm
University of Hertfordshire
Computer Science Department
Hatfield AL10 9AB, UK
E-mail: j.a.malcolm@herts.ac.uk

Michael Roe
Microsoft Research Ltd.
7 J.J. Thomson Avenue
Cambridge CB3 0FB, UK
E-mail: mroe@microsoft.com

Library of Congress Control Number: 2006932036

CR Subject Classification (1998): E.3, F.2.1-2, C.2, K.6.5, J.1, K.4.1, D.4.6

LNCS Sublibrary: SL 4 – Security and Cryptology

ISSN	0302-9743
ISBN-10	3-540-40925-4 Springer Berlin Heidelberg New York
ISBN-13	978-3-540-40925-0 Springer Berlin Heidelberg New York

This work is subject to copyright. All rights are reserved, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, re-use of illustrations, recitation, broadcasting, reproduction on microfilms or in any other way, and storage in data banks. Duplication of this publication or parts thereof is permitted only under the provisions of the German Copyright Law of September 9, 1965, in its current version, and permission for use must always be obtained from Springer. Violations are liable to prosecution under the German Copyright Law.

Springer is a part of Springer Science+Business Media
springer.com

© Springer-Verlag Berlin Heidelberg 2006
Printed in Germany

Typesetting: Camera-ready by author, data conversion by Scientific Publishing Services, Chennai, India
Printed on acid-free paper SPIN: 11861386 06/3142 5 4 3 2 1 0

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

Preface

Here are the proceedings of the 12th International Workshop on Security Protocols. We hope that you will enjoy them, and that they will cause you to think at least one heretical thought. Please write or e-mail and share it with us.

Our theme this workshop was “Authentic Privacy.” Traditionally we have based authentication upon a rather strong notion of identity, and have then built other security services on top of authentication. Perhaps if we want a more nuanced notion of privacy, then we need to re-examine some of our assumptions, particularly when attackers and defenders share the same resources and infrastructure.

The position papers published here have been revised by the participants in the workshop, and are followed by edited (heavily in some cases) transcripts of parts of the discussions which they led.

Our thanks to Sidney Sussex College Cambridge for the use of their facilities, to Johanna Hunt at the University of Hertfordshire for organizing the logistics of the workshop and orchestrating the production of these proceedings, to Lori Klimaszewska of the University of Cambridge Computing Service for transcribing the audio tapes (in which “viruses without halos” could have caused havoc but didn’t), and to Donald Hunt for impeccable copyediting.

Finally, it is both a sadness and a pleasure to pay our tribute to David Wheeler, one of the original forty-niners at the Cambridge Computer Laboratory and author of the initial orders for EDSAC. The second version of initial orders *is* the Platonic bootstrap.

These workshops grew out of a series of informal meetings, which migrated between David’s office in the old Computer Laboratory tower and the front room of the Eagle across the road, and where so many of us were touched forever by our encounters with his decades of fearless thought. Time was finally called while these proceedings were being prepared.

February 2006

Bruce Christianson
Bruno Crispo
James Malcolm
Michael Roe

Previous Proceedings in This Series

The proceedings of previous International Workshops on Security Protocols have also been published by Springer as *Lecture Notes in Computer Science*, and are occasionally referred to in the text:

11th Workshop (2003), LNCS 3364, ISBN 3-540-28389-7
10th Workshop (2002), LNCS 2845, ISBN 3-540-20830-5
9th Workshop (2001), LNCS 2467, ISBN 3-540-44263-4
8th Workshop (2000), LNCS 2133, ISBN 3-540-42566-7
7th Workshop (1999), LNCS 1796, ISBN 3-540-67381-4
6th Workshop (1998), LNCS 1550, ISBN 3-540-65663-4
5th Workshop (1997), LNCS 1361, ISBN 3-540-64040-1
4th Workshop (1996), LNCS 1189, ISBN 3-540-63494-5

Lecture Notes in Computer Science

For information about Vols. 1–4082

please contact your bookseller or Springer

- Vol. 4228: D.E. Lightfoot, C.A. Szyperski (Eds.), *Modular Programming Languages*. X, 415 pages. 2006.
- Vol. 4206: P. Dourish, A. Friday (Eds.), *UbiComp 2006: Ubiquitous Computing*. XIX, 526 pages. 2006.
- Vol. 4193: T.P. Runarsson, H.-G. Beyer, E. Burke, J.J. Merelo-Guervós, L. D. Whitley, X. Yao (Eds.), *Parallel Problem Solving from Nature - PPSN IX*. XIX, 1061 pages. 2006.
- Vol. 4192: B. Mohr, J.L. Träff, J. Worringer, J. Dongarra (Eds.), *Recent Advances in Parallel Virtual Machine and Message Passing Interface*. XVI, 414 pages. 2006.
- Vol. 4188: P. Sojka, I. Kopeček, K. Pala (Eds.), *Text, Speech and Dialogue*. XIV, 721 pages. 2006. (Sublibrary LNAI).
- Vol. 4187: J.J. Alferes, J. Bailey, W. May, U. Schwertel (Eds.), *Principles and Practice of Semantic Web Reasoning*. XI, 277 pages. 2006.
- Vol. 4186: C. Jesshope, C. Egan (Eds.), *Advances in Computer Systems Architecture*. XIV, 605 pages. 2006.
- Vol. 4185: R. Mizoguchi, Z. Shi, F. Giunchiglia (Eds.), *The Semantic Web – ASWC 2006*. XX, 778 pages. 2006.
- Vol. 4184: M. Bravetti, M. Núñez, G. Zavattaro (Eds.), *Web Services and Formal Methods*. X, 289 pages. 2006.
- Vol. 4183: J. Euzenat, J. Domingue (Eds.), *Artificial Intelligence: Methodology, Systems, and Applications*. XIII, 291 pages. 2006. (Sublibrary LNAI).
- Vol. 4180: M. Kohlhase, OMDoc – An Open Markup Format for Mathematical Documents [version 1.2]. XIX, 428 pages. 2006. (Sublibrary LNAI).
- Vol. 4178: A. Corradini, H. Ehrig, U. Montanari, L. Ribeiro, G. Rozenberg (Eds.), *Graph Transformations*. XII, 473 pages. 2006.
- Vol. 4176: S.K. Katsikas, J. Lopez, M. Backes, S. Gritzalis, B. Preneel (Eds.), *Information Security*. XIV, 548 pages. 2006.
- Vol. 4175: P. Bücher, B.M.E. Moret (Eds.), *Algorithms in Bioinformatics*. XII, 402 pages. 2006. (Sublibrary LNBI).
- Vol. 4169: H.L. Bodlaender, M.A. Langston (Eds.), *Parameterized and Exact Computation*. XI, 279 pages. 2006.
- Vol. 4168: Y. Azar, T. Erlebach (Eds.), *Algorithms – ESA 2006*. XVIII, 843 pages. 2006.
- Vol. 4165: W. Jonker, M. Petković (Eds.), *Secure, Data Management*. X, 185 pages. 2006.
- Vol. 4163: H. Bersini, J. Carneiro (Eds.), *Artificial Immune Systems*. XII, 460 pages. 2006.
- Vol. 4162: R. Kráľovič, P. Urzyczyn (Eds.), *Mathematical Foundations of Computer Science 2006*. XV, 814 pages. 2006.
- Vol. 4159: J. Ma, H. Jin, L.T. Yang, J.J.-P. Tsai (Eds.), *Ubiquitous Intelligence and Computing*. XXII, 1190 pages. 2006.
- Vol. 4158: L.T. Yang, H. Jin, J. Ma, T. Ungerer (Eds.), *Autonomic and Trusted Computing*. XIV, 613 pages. 2006.
- Vol. 4156: S. Amer-Yahia, Z. Bellahsene, E. Hunt, R. Unland, J.X. Yu (Eds.), *Database and XML Technologies*. IX, 123 pages. 2006.
- Vol. 4155: O. Stock, M. Schaerf (Eds.), *Reasoning, Action and Interaction in AI Theories and Systems*. XVIII, 343 pages. 2006. (Sublibrary LNAI).
- Vol. 4153: N. Zheng, X. Jiang, X. Lan (Eds.), *Advances in Machine Vision, Image Processing, and Pattern Analysis*. XIII, 506 pages. 2006.
- Vol. 4152: Y. Manolopoulos, J. Pokorný, T. Sellis (Eds.), *Advances in Databases and Information Systems*. XV, 448 pages. 2006.
- Vol. 4151: A. Iglesias, N. Takayama (Eds.), *Mathematical Software - ICMS 2006*. XVII, 452 pages. 2006.
- Vol. 4150: M. Dorigo, L.M. Gambardella, M. Birattari, A. Martinoli, R. Poli, T. Stützle (Eds.), *Ant Colony Optimization and Swarm Intelligence*. XVI, 526 pages. 2006.
- Vol. 4149: M. Klusch, M. Rovatos, T.R. Payne (Eds.), *Cooperative Information Agents X*. XII, 477 pages. 2006. (Sublibrary LNAI).
- Vol. 4148: J. Vounckx, N. Azemard, P. Maurine (Eds.), *Integrated Circuit and System Design*. XVI, 677 pages. 2006.
- Vol. 4146: J.C. Rajapakse, L. Wong, R. Acharya (Eds.), *Pattern Recognition in Bioinformatics*. XIV, 186 pages. 2006. (Sublibrary LNBI).
- Vol. 4144: T. Ball, R.B. Jones (Eds.), *Computer Aided Verification*. XV, 564 pages. 2006.
- Vol. 4139: T. Salakoski, F. Ginter, S. Pyysalo, T. Pahikkala, *Advances in Natural Language Processing*. XVI, 771 pages. 2006. (Sublibrary LNAI).
- Vol. 4138: X. Cheng, W. Li, T. Znati (Eds.), *Wireless Algorithms, Systems, and Applications*. XVI, 709 pages. 2006.
- Vol. 4137: C. Baier, H. Hermanns (Eds.), *CONCUR 2006 – Concurrency Theory*. XIII, 525 pages. 2006.
- Vol. 4136: R.A. Schmidt (Ed.), *Relations and Kleene Algebra in Computer Science*. XI, 433 pages. 2006.
- Vol. 4135: C.S. Calude, M.J. Dinneen, G. Păun, G. Rozenberg, S. Stepney (Eds.), *Unconventional Computation*. X, 267 pages. 2006.

- Vol. 4134: K. Yi (Ed.), *Static Analysis*. XIII, 443 pages. 2006.
- Vol. 4133: J. Gratch, M. Young, R. Aylett, D. Ballin, P. Olivier (Eds.), *Intelligent Virtual Agents*. XIV, 472 pages. 2006. (Sublibrary LNAI).
- Vol. 4132: S. Kollias, A. Stafylopatis, W. Duch, E. Oja (Eds.), *Artificial Neural Networks – ICANN 2006, Part II*. XXXIV, 1028 pages. 2006.
- Vol. 4131: S. Kollias, A. Stafylopatis, W. Duch, E. Oja (Eds.), *Artificial Neural Networks – ICANN 2006, Part I*. XXXIV, 1008 pages. 2006.
- Vol. 4130: U. Furbach, N. Shankar (Eds.), *Automated Reasoning*. XV, 680 pages. 2006. (Sublibrary LNAI).
- Vol. 4129: D. McGookin, S. Brewster (Eds.), *Haptic and Audio Interaction Design*. XII, 167 pages. 2006.
- Vol. 4128: W.E. Nagel, W.V. Walter, W. Lehner (Eds.), *Euro-Par 2006 Parallel Processing*. XXXIII, 1221 pages. 2006.
- Vol. 4127: E. Damiani, P. Liu (Eds.), *Data and Applications Security XX*. X, 319 pages. 2006.
- Vol. 4126: P. Barahona, F. Bry, E. Franconi, N. Henze, U. Sattler, *Reasoning Web*. X, 269 pages. 2006.
- Vol. 4124: H. de Meer, J.P. G. Sterbenz (Eds.), *Self-Organizing Systems*. XIV, 261 pages. 2006.
- Vol. 4121: A. Biere, C.P. Gomes (Eds.), *Theory and Applications of Satisfiability Testing - SAT 2006*. XII, 438 pages. 2006.
- Vol. 4119: C. Dony, J.L. Knudsen, A. Romanovsky, A. Tripathi (Eds.), *Advanced Topics in Exception Handling Components*. X, 302 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. 4115: D.-S. Huang, K. Li, G.W. Irwin (Eds.), *Computational Intelligence and Bioinformatics, Part III*. XXI, 803 pages. 2006. (Sublibrary LNBI).
- Vol. 4114: D.-S. Huang, K. Li, G.W. Irwin (Eds.), *Computational Intelligence, Part II*. XXVII, 1337 pages. 2006. (Sublibrary LNAI).
- Vol. 4113: D.-S. Huang, K. Li, G.W. Irwin (Eds.), *Intelligent Computing, Part I*. XXVII, 1331 pages. 2006.
- Vol. 4112: D.Z. Chen, D. T. Lee (Eds.), *Computing and Combinatorics*. XIV, 528 pages. 2006.
- Vol. 4111: F.S. de Boer, M.M. Bonsangue, S. Graf, W.-P. de Roeer (Eds.), *Formal Methods for Components and Objects*. VIII, 447 pages. 2006.
- Vol. 4110: J. Díaz, K. Jansen, J.D.P. Rolim, U. Zwick (Eds.), *Approximation, Randomization, and Combinatorial Optimization*. XII, 522 pages. 2006.
- Vol. 4109: D.-Y. Yeung, J.T. Kwok, A. Fred, F. Roli, D. de Ridder (Eds.), *Structural, Syntactic, and Statistical Pattern Recognition*. XXI, 939 pages. 2006.
- Vol. 4108: J.M. Borwein, W.M. Farmer (Eds.), *Mathematical Knowledge Management*. VIII, 295 pages. 2006. (Sublibrary LNAI).
- Vol. 4106: T.R. Roth-Berghofer, M.H. Göker, H. A. Güvenir (Eds.), *Advances in Case-Based Reasoning*. XIV, 566 pages. 2006. (Sublibrary LNAI).
- Vol. 4105: B. Gunsel, A.K. Jain, A. M. Tekalp, B. Sankur (Eds.), *Multimedia, Content Representation, Classification and Security*. XIX, 804 pages. 2006.
- Vol. 4104: T. Kunz, S.S. Ravi (Eds.), *Ad-Hoc, Mobile, and Wireless Networks*. XII, 474 pages. 2006.
- Vol. 4103: J. Eder, S. Dustdar (Eds.), *Business Process Management Workshops*. XI, 508 pages. 2006.
- Vol. 4102: S. Dustdar, J.L. Fiadeiro, A. Sheth (Eds.), *Business Process Management*. XV, 486 pages. 2006.
- Vol. 4099: Q. Yang, G. Webb (Eds.), *PRICAI 2006: Trends in Artificial Intelligence*. XXVIII, 1263 pages. 2006. (Sublibrary LNAI).
- Vol. 4098: F. Pfenning (Ed.), *Term Rewriting and Applications*. XIII, 415 pages. 2006.
- Vol. 4097: X. Zhou, O. Sokolsky, L. Yan, E.-S. Jung, Z. Shao, Y. Mu, D.C. Lee, D. Kim, Y.-S. Jeong, C.-Z. Xu (Eds.), *Emerging Directions in Embedded and Ubiquitous Computing*. XXVII, 1034 pages. 2006.
- Vol. 4096: E. Sha, S.-K. Han, C.-Z. Xu, M.H. Kim, L.T. Yang, B. Xiao (Eds.), *Embedded and Ubiquitous Computing*. XXIV, 1170 pages. 2006.
- Vol. 4095: S. Nolfi, G. Baldassarre, R. Calabretta, J.C. T. Hallam, D. Marocco, J.-A. Meyer, O. Miglino, D. Parisi (Eds.), *From Animals to Animats 9*. XV, 869 pages. 2006. (Sublibrary LNAI).
- Vol. 4094: O. H. Ibarra, H.-C. Yen (Eds.), *Implementation and Application of Automata*. XIII, 291 pages. 2006.
- Vol. 4093: X. Li, O.R. Zaïane, Z. Li (Eds.), *Advanced Data Mining and Applications*. XXI, 1110 pages. 2006. (Sublibrary LNAI).
- Vol. 4092: J. Lang, F. Lin, J. Wang (Eds.), *Knowledge Science, Engineering and Management*. XV, 664 pages. 2006. (Sublibrary LNAI).
- Vol. 4091: G.-Z. Yang, T. Jiang, D. Shen, L. Gu, J. Yang (Eds.), *Medical Imaging and Augmented Reality*. XIII, 399 pages. 2006.
- Vol. 4090: S. Spaccapietra, K. Aberer, P. Cudré-Mauroux (Eds.), *Journal on Data Semantics VI*. XI, 211 pages. 2006.
- Vol. 4089: W. Löwe, M. Südholt (Eds.), *Software Composition*. X, 339 pages. 2006.
- Vol. 4088: Z.-Z. Shi, R. Sadananda (Eds.), *Agent Computing and Multi-Agent Systems*. XVII, 827 pages. 2006. (Sublibrary LNAI).
- Vol. 4087: F. Schwenker, S. Marinai (Eds.), *Artificial Neural Networks in Pattern Recognition*. IX, 299 pages. 2006. (Sublibrary LNAI).
- Vol. 4085: J. Misra, T. Nipkow, E. Sekerinski (Eds.), *FM 2006: Formal Methods*. XV, 620 pages. 2006.
- Vol. 4084: M.A. Wimmer, H.J. Scholl, Å. Grönlund, K.V. Andersen (Eds.), *Electronic Government*. XV, 353 pages. 2006.
- Vol. 4083: S. Fischer-Hübner, S. Furnell, C. Lambri-noudakis (Eds.), *Trust and Privacy in Digital Business*. XIII, 243 pages. 2006.

Table of Contents

Introduction: Authentic Privacy	1
<i>Bruce Christianson</i>	
Limits to Anonymity When Using Credentials	4
<i>Andreas Pashalidis, Chris J. Mitchell *</i>	
Discussion	13
A Cryptographic Framework for the Controlled Release of Certified	
Data	20
<i>Endre Bangerter*, Jan Camenisch, Anna Lysyanskaya</i>	
Discussion	43
One User, Many Hats; and, Sometimes, No Hat: Towards a Secure Yet	
Usable PDA	51
<i>Frank Stajano*</i>	
Discussion	65
Authentication Components: Engineering Experiences and	
Guidelines	68
<i>Pasi Eronen*, Jari Arkko</i>	
Discussion	78
Accountable Privacy	83
<i>Mike Burmester, Yvo Desmedt, Rebecca N. Wright*,</i>	
<i>Alec Yasinsac</i>	
Discussion	96
Toward a Broader View of Security Protocols	106
<i>Matt Blaze*</i>	
Discussion	121
Privacy, Control and Internet Mobility	133
<i>Tuomas Aura*, Alf Zugenmaier</i>	
Discussion	146
Controlling Who Tracks Me	151
<i>Denis Bohm, Mik Lamming, Robert N. Mayo*, Jeff Morgan,</i>	
<i>Kan Zhang</i>	
Discussion	155
BLIND: A Complete Identity Protection Framework for End-Points	163
<i>Jukka Ylitalo*, Pekka Nikander</i>	
Discussion	177

Privacy Is Linking Permission to Purpose	179
<i>Fabio Massaccesi*, Nicola Zannone</i>	
Discussion	192
Establishing Trust with Privacy.....	199
<i>Laurent Bussard*, Refik Molva</i>	
Discussion	210
Safe and Private Data Sharing with Turtle: Friends Team-Up and Beat the System	213
<i>Bogdan C. Popescu*, Bruno Crispo,</i>	
<i>Andrew S. Tanenbaum</i>	
Discussion	221
The Dancing Bear: A New Way of Composing Ciphers.....	231
<i>Ross Anderson*</i>	
Discussion	239
Identity and Location	246
<i>Dieter Gollmann*</i>	
Discussion	251
Security of Emergent Properties in Ad-Hoc Networks (Transcript of Discussion)	256
<i>Virgil Gligor</i>	
Pseudonymity in the Light of Evidence-Based Trust	267
<i>Daniel Cvrček*, Václav Matyáš Jr.</i>	
Discussion	275
Secure Processor Consistent with Both Foreign Software Protection and User Privacy Protection.....	276
<i>Mikio Hashimoto*, Hiroyoshi Haruki, Takeshi Kawabata</i>	
Discussion	287
Why Are We Authenticating (Transcript of Discussion)	291
<i>Mark Lomas</i>	
Anonymous Authentication	299
<i>Partha Das Chowdhury, Bruce Christianson*,</i>	
<i>James Malcolm</i>	
Discussion	306
Towards a Mechanism for Discretionary Overriding of Access Control....	312
<i>Erik Rissanen*, Babak Sadighi Firozabadi,</i>	
<i>Marek Sergot</i>	
Discussion	320

Last Orders 324
 David Wheeler

Author Index 325

Introduction: Authentic Privacy

Bruce Christianson

University of Hertfordshire, UK

Well hello, and welcome to the twelfth Security Protocols Workshop. When this all started we had no idea what a juggernaut we were creating, and it's particularly nice to see so many young people here. [Laughter]

There's a tradition that we start by spending five minutes introducing the theme and then we don't mention it again for the next 48 hours. This year being no exception, I shall now explain this year's theme.

There was a time when on the Internet you would see something called privacy being talked about as if it were an absolute good like, say, free education or health care. I think those days are now pretty thoroughly gone. It's clear that beyond a certain point amplifying privacy actually gives the attackers of the system more of an advantage than it gives the defenders, and there's also a point beyond which increased privacy doesn't really seem to be particularly useful to legitimate users of the system. These remarks are intended to be controversial by the way. [Laughter]

But what is privacy? Everyone is using different definitions of the word. When they say privacy, some people mean anonymity, some mean uncorrelatability, *i.e.* that the different things that you do can't be correlated. Some mean that you can tell who someone is but not what they're doing, some mean that you can tell exactly what is being done but not who is doing it, some mean that you can tell who's doing it, and what they're doing, but you can't tell why they're doing it. Quite often people slide between different definitions in the same paper, and sometimes they are smuggling: the sliding definition is a false bottom in the privacy suitcase, and they're using this to smuggle some different concept in.

Matt Blaze: There's another definition that the privacy people who are not technologists are more concerned with. For example, the OECD¹ privacy guidelines. This is more concerned with personally identifiable information, the ability to control the propagation of this, and perhaps how the data is used.

Reply: Yes indeed, and this whole side of the business is about how you control information once the cat is out of the bag: how do you track where it goes, what it does, and what the result is used for. And what sanction do you have against people who are misusing it when they themselves may have some form of anonymity or privacy?

Well, having almost gotten away with saying that unlimited privacy is bad — whatever privacy is taken to mean — the contrary alternative is to say that there's no privacy at all. This extreme alternative point of view says that com-

¹ www.oecd.org

puter systems should be a panopticon², so that every single thing that everyone does should be visible to all.³ This approach has a certain twisted intellectual appeal.

Now we've got to be a little careful here. Some operations in the underware⁴ like using a cryptographic key, or entering a password, have to be in some sense private because that's the nature of them. But we're computer scientists, we're used to working at lots of different levels of abstraction, and the question is, is there a nice level of abstraction at which privacy is not an issue, or at least is not something anyone particularly desires⁵.

Some of you may remember the Cambridge active badge system that used to operate in the old Computer Laboratory. One of the nice features of that system was that, yes indeed, you could ping anybody and find out where they were and then speculate about what they were doing there, but the fact that you had pinged them was instantly visible to the person whom you pinged. They knew who was pinging them and where you were, the loss of privacy was symmetric.

It's quite nice to believe that something like this could be done on a larger scale, but it really only seems to be viable in a relatively closed community, which the Computer Lab at that time perhaps was. [Laughter] As soon as you start to have a more open environment, it becomes hard to see how to implement the kind of counter-mechanisms that would be needed to ensure that loss of privacy was symmetric.

But it's a nice idea that somebody who abuses the protocols that are supposed to protect the security of the system is punished by losing some element of their privacy (in whatever way we interpret privacy). We might perhaps see information of a personal nature about them having something done with it that the system would not be able to do had they not breached the protocol.

This is a bit like what happens if you double-spend some forms of digital cash, where breaking the protocol is what releases the information that allows the privacy loss to take place. But we don't, I think, really have a strong enough cross-domain infrastructure to be able to do very much along that line yet.

There's also the issue of by whom personal information, or private information, is (or should be) held. It's clear that traditional answers along the lines of the system, or the administrator, or the government, are no longer candidates for being the right answer.

I think I'd like to argue that getting any further with the development of protocols for authentication, or of other protocols which need to be audited by third parties who are not necessarily trusted by the first two parties, might require rather gentler notions of personal identity than we've been in the habit

² See Jeremy Bentham, "Panopticon, or The Inspection-House", 1787, Crecheff (published 1791, London + Dublin) see <http://cartome.org/panopticon2.htm>

³ Michael Roe points out that this analogy is not quite as straightforward as it seems. The inmates of Bentham's panopticon do retain some privacy, because they cannot see one another (and so are prevented from acting in combination). It is only the overlookers who have no privacy at all.

⁴ Or whatever we call it now.

⁵ After all, this is just about the case with other properties such as determinism.

of tying authentication to. Traditionally we've tied authentication to a rather strong notion of personal identity, and then we've pinned everything else on the back of that strong form of authentication.

Maybe if we want a more nuanced notion of privacy, and a more graceful degradation in the event of misuse of information, then we have to look at different ways of separating, or trying to orthogonalize, the issues of who's doing it, what are they doing, why are they doing it, who knows about it, and who else is affected.

At the moment we have a single authentication mechanism living in the basement that we try to make (or just hope will) do service for everything. We encapsulate this strong authentication as a service, and then encourage everyone else to build their secure service on top of it. It's not clear that this is the correct way forward in a more open environment. Perhaps strong authentication now belongs in the attic, along with the first Mrs Rochester.

This is a workshop not a conference, and the intention is that the presenters should be leading a discussion, or in some cases, trying to keep up with a discussion that has gone in a different direction to the direction in which they intended their talk to go. Please don't allow yourself to become constrained by the presentation which you prepared before you came.

Likewise, within the normal limits of academic debate (no personal attacks, no hitting with the closed fist) you're free to make whatever points from the floor you wish. The only proviso is that if you break somebody else's idea, then you are under an obligation to help them sort out the pieces at teatime. A correctly broken idea is often more interesting than a flawlessly polished one.

Limits to Anonymity When Using Credentials

Andreas Pashalidis* and Chris J. Mitchell

Royal Holloway, University of London,
Egham, Surrey, TW20 0EX, United Kingdom,
{A.Pashalidis, C.Mitchell}@rhul.ac.uk
<http://www.isg.rhul.ac.uk>

Abstract. This paper identifies certain privacy threats that apply to anonymous credential systems. The focus is on timing attacks that apply even if the system is cryptographically secure. The paper provides some simple heuristics that aim to mitigate the exposure to the threats and identifies directions for further research.

Keywords: anonymous credential systems, pseudonym systems, unlinkability, privacy, timing attacks.

1 Introduction

Credential systems allow subjects to prove possession of attributes to interested parties. In a sound credential system subjects first need to obtain a structure termed a *credential* from an entity termed the credential *issuer*. The issuer encodes some well-defined set of attributes together with their values into the credential which is then passed on, or ‘granted’, to the subject. Only after having gone through this process can the subject prove possession of those (and only those) attributes that are encoded in the credential. During this latter process, the interested party is said to ‘verify the credential’ and is therefore called a *verifier*. Subjects are typically human users, issuers are typically well-known organisations with authority over the attributes they encode into the credentials they issue, and verifiers typically are service providers that perform attribute-based access control.

An example of a credential system is a Public Key Infrastructure (PKI). In a PKI, credentials are public key certificates that bind together subject attributes such as subject name, public key, its issue and expiry dates, and so on. The credential issuer is the Certification Authority (CA); it grants public key certificates according to some subject registration procedure. Finally, credential verifiers are the entities within the PKI that accept the certificates issued by the CA.

In conventional credential systems (e.g. a PKI), issuers and verifiers identify any given subject by a system-wide identifier. This has a potentially severe impact on the subject’s privacy, as it enables issuers and verifiers to combine their knowledge about the subject. Indeed, they can construct individual transaction

* The author is sponsored by the State Scholarship Foundation of Greece.

histories for all the subjects in the system, simply by correlating credential-related events using these identifiers.

Over the last 20 years, a significant amount of research has been performed on credential systems that try to address the above privacy issue (see, for example, [2,3,4,6,7,8,10,11]). These systems are known as *anonymous* credential systems. In an anonymous credential system, subjects establish a different identifier with each issuer and verifier they wish to interact with, where we assume throughout that these pseudonyms cannot be connected to the subject's true identity. These identifiers, termed the subject's *pseudonyms*, are unlinkable, i.e. they do not possess any connection with one another. This means that it is infeasible, for colluding issuers and verifiers, to decide with certainty whether or not any given pair of pseudonyms belongs to the same subject¹. While a subject obtains a credential under the pseudonym that was established with the issuer, proof of its possession² takes place under the pseudonym established with the verifier. Of course, in order for the system to remain sound, subjects should only be able to successfully prove possession of credentials that they were indeed issued by some legitimate issuer.

In this paper, we consider practical limits to the level of pseudonym unlinkability (and, thus, subject privacy) offered by anonymous credential systems. In particular, assuming the soundness and security of such a system, we consider how timing attacks, launched by colluding issuers and verifiers, may affect pseudonym unlinkability. Finally, we outline possible pragmatic approaches to minimising exposure to such attacks.

The paper is structured as follows. The next section outlines the assumptions we make about anonymous credential systems, section 3 discusses the issue of encoding freshness into credentials and section 4 presents the timing attacks. Section 5 provides some simple heuristics to counter the attacks and section 6 concludes, giving directions for further research.

2 A General Model for Anonymous Credential Systems

A number of anonymous credential systems have been proposed in the literature, each with its own particular set of entities, underlying problems, assumptions and properties. This section presents the model of anonymous credential systems on which the rest of the paper is based. It is intended to be as general as possible, in order to be consistent with the majority of existing schemes.

We consider an anonymous credential system to involve three types of player: subjects, issuers and verifiers. We refer to issuers and verifiers, collectively, as 'organisations'. It is assumed that subjects establish at least one pseudonym with each organisation with which they wish to interact. These pseudonyms are assumed to be indistinguishable, meaning that they do not bear any connection

¹ Assuming that at least two subjects exist within the system.

² Proving possession of a credential amounts to proving possession of the attributes that are encoded within the credential. We refer to this process also as the *showing* of a credential.

to the identity of the subject they belong to. We further assume that pseudonyms are unlinkable, i.e. two pseudonyms for the same subject cannot be linked to each other. Subjects may obtain credentials, i.e. structures that encode a well-defined, finite set of attributes together with their values, from issuers. They may subsequently show those credentials to verifiers, i.e. convince them that they possess (possibly a subset of) the encoded attributes. A credential is issued under a pseudonym that the subject has established with its issuer, and it is shown under the pseudonym that the subject has established with the relevant verifier.

It is assumed that the anonymous credential system is sound. This means that it offers *pseudonym owner protection*, i.e. that only the subject that established a given pseudonym can show credentials under it. Soundness also implies *credential unforgeability*; the only way that subjects may prove possession of a credential is by having obtained it previously from a legitimate issuer. In some applications, it is required that the system offers the stronger property of *credential non-transferability*. This property guarantees that no subject can prove possession of a credential that it has not been issued, even if the subject colludes with other subject(s) that may have (legitimately) obtained such a credential. In other words, a system that offers non-transferability prohibits credential sharing, whereas a system that offers only unforgeability, does not. (Of course, the degree of protection against credential sharing is always limited, since if one subject gives all its secrets to another subject then the latter subject will always be able to impersonate the former and use its credentials.) We require that credentials are bound to the subject to which they have been issued. We therefore assume that either the system offers non-transferability or that in practice subjects do not share their credentials.

It is assumed further that the system properly protects privacy in that a subject's transactions with organisations do not compromise the unlinkability of its pseudonyms. We note, however, that this unlinkability can only be guaranteed up to a certain point, as credential *types* potentially reveal links between pseudonyms. The type of a credential is defined as the collection of attribute values that are encoded into the credential. An organisation, for example, that issues demographic credentials containing the fields **sex** and **age group**, with possible values of {male, female} and {18-, 18-30, 30-50, 50+} respectively, may actually issue up to 8 different types of credential (one for each combination of values). To see how credential types can be exploited to link subject's pseudonyms, consider the following trivial scenario. At time τ , a credential of type t is shown under the pseudonym p . However, suppose that up to time τ , only one credential of type t has been issued, and this was done under pseudonym p' . It follows, under the assumption that credentials are bound to subjects, that the two pseudonyms p, p' belong to the same subject; the colluding organisations can successfully link those two pseudonyms.

We note that, as part of credential showing, some anonymous credential systems allow subjects to reveal only a subset of the encoded attributes; in the above example it may be possible for the subject to reveal only the value of