

**John Ioannidis
Angelos Keromytis
Moti Yung (Eds.)**

LNCS 3531

Applied Cryptography and Network Security

**Third International Conference, ACNS 2005
New York, NY, USA, June 2005
Proceedings**



Springer

John Ioannidis Angelos Keromytis
Moti Yung (Eds.)

Applied Cryptography and Network Security

Third International Conference, ACNS 2005
New York, NY, USA, June 7-10, 2005
Proceedings



Volume Editors

John Ioannidis
Columbia University
Center for Computational Learning Systems
NY, USA
E-mail: ji@cs.columbia.edu

Angelos Keromytis
Moti Yung
Columbia University
Department of Computer Science
NY, USA
E-mail: {angelos,moti}@cs.columbia.edu

Library of Congress Control Number: 2005926829

CR Subject Classification (1998): E.3, C.2, D.4.6, H.3-4, K.4.4, K.6.5

ISSN	0302-9743
ISBN-10	3-540-26223-7 Springer Berlin Heidelberg New York
ISBN-13	978-3-540-26223-7 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
springeronline.com

© Springer-Verlag Berlin Heidelberg 2005
Printed in Germany

Typesetting: Camera-ready by author, data conversion by Olgun Computergrafik
Printed on acid-free paper SPIN: 11496137 06/3142 5 4 3 2 1 0

Preface

The 3rd International Conference on Applied Cryptography and Network Security (ACNS 2005) was sponsored and organized by ICISA (the International Communications and Information Security Association). It was held at Columbia University in New York, USA, June 7–10, 2005. This conference proceedings volume contains papers presented in the academic/research track.

ACNS covers a large number of research areas that have been gaining importance in recent years due to the development of the Internet, wireless communication and the increased global exposure of computing resources. The papers in this volume are representative of the state of the art in security and cryptography research, worldwide.

The Program Committee of the conference received a total of 158 submissions from all over the world, of which 35 submissions were selected for presentation at the academic track. In addition to this track, the conference also hosted a technical/ industrial/ short papers track whose presentations were also carefully selected from among the submissions. All submissions were reviewed by experts in the relevant areas.

Many people and organizations helped in making the conference a reality. We would like to take this opportunity to thank the Program Committee members and the external experts for their invaluable help in producing the conference's program. We also wish to thank Michael E. Locasto for his help in all technical and technological aspects of running the conference and Sophie Majewski for the administrative support in organizing the conference. We wish to thank the graduate students at Columbia University's Computer Science Department who helped us as well.

We wish to acknowledge the financial support of our sponsors, and their employees who were instrumental in the sponsorship process: Morgan Stanley (Ben Fried), Gemplus (David Naccache), and Google (Niels Provos).

Finally, we would like to thank all the authors who submitted papers to the conference; the continued support of the security and cryptography research community worldwide is what really enabled us to have this conference.

May 2005

John Ioannidis
Angelos Keromytis
Moti Yung

ACNS 2005

3rd International Conference on Applied Cryptography and Network Security

New York, USA

June 7–10, 2005

Sponsored and organized by the
International Communications and Information Security Association (ICISA)

In coöperation with
Columbia University, USA

General Chair

John Ioannidis Columbia University

Program Chairs

Moti Yung Columbia University and RSA Labs
Angelos Keromytis Columbia University

Program Committee

Scott Alexander Telcordia, USA
Tuomas Aura Microsoft Research, UK
David Brumley CMU, USA
Ran Canetti IBM Research, USA
Marc Dacier Eurecom, France
Ed Dawson Queensland University of Technology, Australia
Glenn Durfee PARC, USA
Virgil Gligor University of Maryland, USA
Peter Gutman University of Auckland, New Zealand
Goichiro Hanaoka National Institute of Advanced Industrial Science
and Technology (AIST), Japan
Amir Herzberg Bar Ilan University, Israel
Russ Housley Vigilsec, USA
John Ioannidis Columbia University, USA
Sotiris Ioannidis University of Pennsylvania, USA

Stas Jarecki	UC Irvine, USA
Ari Juels	RSA Laboratories, USA
Angelos Keromytis	Columbia University, USA
Aggelos Kiayias	University of Connecticut, USA
Tanja Lange	Ruhr-Universität Bochum, Germany
Dong Hoon Lee	Korea University, South Korea
Fabio Massacci	University Trento, Italy
Atsuko Miyaji	JAIST, Japan
Frederic Muller	DCSSI Crypto Lab, France
Kaisa Nyberg	Nokia, Finland
Bart Preneel	K.U.Leuven, Belgium
Vassilis Prevelakis	Drexel University, USA
Niels Provos	Google, USA
Pierangela Samarati	University of Milan, Italy
Tomas Sander	HP, USA
Dan Simon	Microsoft Research, USA
Tsuyoshi Takagi	T.U. Darmstadt, Germany
Wen-Guey Tzeng	NCTU, Taiwan
Dan Wallach	Rice University, USA
Susanne Wetzel	Stevens Institute, USA
Moti Yung	Columbia University, USA
Jianying Zhou	I2R, Singapore
Lidong Zhou	Microsoft Research, USA

External Reviewers

Kouichiro Akiyama, Kostas Anagnostakis, Farooq Anjum, N. Asokan, Nuttapong Attrapadung, Roberto Avanzi, Dirk Balfanz, Lejla Batina, Steven Bellovin, Josh Benaloh, Enrico Blanzieri, Christophe de Canniere, Alvaro Cardenas, Roberto Cascella, Jae-Gwi Choi, Stelvio Cimato, Jared Cordasco, Giovanni di Crescenzo, Eric Cronin, Stefano Crosta, Yang Cui, Ernesto Damiani, Seiji Doi, Wenliang Du, Detlef Duehnlein, Patrick Felke, Pierre-Alain Fouque, Eiichiro Fujisaki, Abhrajit Ghosh, Philippe Golle, Michael Greenwald, Shai Halevi, Nick Howgrave-Graham, Seokhie Hong, Omer Horvitz, Tetsu Iwata, Eliane Jaulmes, Markus Kaiser, Tetsutaro Kobayashi, Sébastien Kunz-Jacques, Kaoru Kurosawa, Klaus Kursawe, Joseph Liu, Dahlia Malkhi, Gwenaëlle Martinet, Mitsuru Matsui, Breno de Medeiros, Nele Mentens, Bernd Meyer, Ulrike Meyer, Ilya Mironov, Anderson Nascimento, Francesco De Natale, Svetla Nikova, Akihito Niwa, Takeshi Okamoto, Tatsuaki Okamoto, Renaud Pacalet, Young-Ho Park, Kirthika Parmeswaran, Guillaume Poupard, Vincent Rijmen, Michael Roe, Tomas Sander, Hisayoshi Sato, Nitesh Saxena, Katja Schmidt-Samoa, Micah Sherr, Diana Smetters, Masakazu Soshi, Jessica Staddon, Martijn Stam, Maria Striki, Gelareh Taban, Rajesh Talpade, Yuuko Tamura, Simon Tsang, Guillaume Urvoy-Keller, Frederik Vercauteren, Sabrina de Capitani di Vimercati, Camille Vuillaume, Zhiguo Wan, Guilin Wang, Kai Wirt, Hongjun Wu, Bennet Yee, Rui Zhang, Sheng Zhong, Huafei Zhu

Table of Contents

Two-Server Password-Only Authenticated Key Exchange	1
<i>Jonathan Katz, Philip MacKenzie, Gelareh Taban, and Virgil Gligor</i>	
Strengthening Password-Based Authentication Protocols Against Online Dictionary Attacks	17
<i>Peng Wang, Yongdae Kim, Vishal Kher, and Taekyoung Kwon</i>	
Cryptanalysis of an Improved Client-to-Client Password-Authenticated Key Exchange (C2C-PAKE) Scheme	33
<i>Raphael C.-W. Phan and Bok-Min Goi</i>	
Efficient Security Mechanisms for Overlay Multicast-Based Content Distribution	40
<i>Sencun Zhu, Chao Yao, Donggang Liu, Sanjeev Setia, and Sushil Jajodia</i>	
A Traitor Tracing Scheme Based on RSA for Fast Decryption	56
<i>John Patrick McGregor, Yiqun Lisa Yin, and Ruby B. Lee</i>	
N-Party Encrypted Diffie-Hellman Key Exchange Using Different Passwords	75
<i>Jin Wook Byun and Dong Hoon Lee</i>	
Messin' with Texas Deriving Mother's Maiden Names Using Public Records	91
<i>Virgil Griffith and Markus Jakobsson</i>	
Mitigating Network Denial-of-Service Through Diversity-Based Traffic Management	104
<i>Ashraf Matrawy, Paul C. van Oorschot, and Anil Somayaji</i>	
Searching for High-Value Rare Events with Uncheatable Grid Computing	122
<i>Wenliang Du and Michael T. Goodrich</i>	
Digital Signatures Do Not Guarantee Exclusive Ownership	138
<i>Thomas Pornin and Julien P. Stern</i>	
Thompson's Group and Public Key Cryptography	151
<i>Vladimir Shpilrain and Alexander Ushakov</i>	
Rainbow, a New Multivariable Polynomial Signature Scheme	164
<i>Jintai Ding and Dieter Schmidt</i>	
Badger – A Fast and Provably Secure MAC	176
<i>Martin Boesgaard, Thomas Christensen, and Erik Zenger</i>	
IDS False Alarm Reduction Using Continuous and Discontinuous Patterns	192
<i>Abdulrahman Alharby and Hideki Imai</i>	

Indexing Information for Data Forensics	206
<i>Michael T. Goodrich, Mikhail J. Atallah, and Roberto Tamassia</i>	
Model Generalization and Its Implications on Intrusion Detection	222
<i>Zhuowei Li, Amitabha Das, and Jianying Zhou</i>	
Intrusion-Resilient Secure Channels	238
<i>Gene Itkis, Robert McNeerney Jr., and Scott Russell</i>	
Optimal Asymmetric Encryption and Signature Paddings	254
<i>Benoît Chevallier-Mames, Duong Hieu Phan, and David Pointcheval</i>	
Efficient and Leakage-Resilient Authenticated Key Transport Protocol Based on RSA	269
<i>SeongHan Shin, Kazukuni Kobara, and Hideki Imai</i>	
Identity Based Encryption Without Redundancy	285
<i>Benoît Libert and Jean-Jacques Quisquater</i>	
OACerts: Oblivious Attribute Certificates	301
<i>Jiangtao Li and Ninghui Li</i>	
Dynamic k -Times Anonymous Authentication	318
<i>Lan Nguyen and Rei Safavi-Naini</i>	
Efficient Anonymous Roaming and Its Security Analysis	334
<i>Guomin Yang, Duncan S. Wong, and Xiaotie Deng</i>	
Quantifying Security in Hybrid Cellular Networks	350
<i>Markus Jakobsson and Liu Yang</i>	
Off-Line Karma: A Decentralized Currency for Peer-to-peer and Grid Applications	364
<i>Flavio D. Garcia and Jaap-Henk Hoepman</i>	
Building Reliable Mix Networks with Fair Exchange	378
<i>Michael K. Reiter, XiaoFeng Wang, and Matthew Wright</i>	
SCARE of the DES (Side Channel Analysis for Reverse Engineering of the Data Encryption Standard)	393
<i>Rémy Daudigny, Hervé Ledig, Frédéric Muller, and Frédéric Valette</i>	
Robust Key Extraction from Physical Unccloneable Functions	407
<i>B. Škorić, P. Tuyls, and W. Ophey</i>	
Efficient Constructions for One-Way Hash Chains	423
<i>Yih-Chun Hu, Markus Jakobsson, and Adrian Perrig</i>	
Privacy Preserving Keyword Searches on Remote Encrypted Data	442
<i>Yan-Cheng Chang and Michael Mitzenmacher</i>	

An Efficient Solution to the Millionaires' Problem Based on Homomorphic Encryption	456
<i>Hsiao-Ying Lin and Wen-Guey Tzeng</i>	
Non-interactive Zero-Knowledge Arguments for Voting	467
<i>Jens Groth</i>	
Short Signature and Universal Designated Verifier Signature Without Random Oracles	483
<i>Rui Zhang, Jun Furukawa, and Hideki Imai</i>	
Efficient Identity Based Ring Signature	499
<i>Sherman S.M. Chow, Siu-Ming Yiu, and Lucas C.K. Hui</i>	
New Signature Schemes with Coupons and Tight Reduction	513
<i>Benoît Chevallier-Mames</i>	
Author Index	529

Two-Server Password-Only Authenticated Key Exchange

Jonathan Katz^{1,*}, Philip MacKenzie², Gelareh Taban³, and Virgil Gligor³

¹ Dept. of Computer Science, University of Maryland
jkatz@cs.umd.edu

² DoCoMo USA Labs, USA
philmac@docomolabs-usa.com

³ Dept. of Electrical and Computer Engineering, University of Maryland
{gelareh,gligor}@eng.umd.edu

Abstract. Typical protocols for password-based authentication assume a single server which stores all the information (e.g., the password) necessary to authenticate a user. Unfortunately, an inherent limitation of this approach (assuming low-entropy passwords are used) is that the user's password is exposed if this server is ever compromised. To address this issue, a number of schemes have been proposed in which a user's password information is shared among multiple servers, and these servers cooperate in a threshold manner when the user wants to authenticate.

We show here a two-server protocol for this task assuming public parameters available to everyone in the system (as well as the adversary). Ours is the first provably-secure two-server protocol for the important *password-only* setting (in which the user need remember only a password, and not the servers' public keys), and is the first two-server protocol (in any setting) with a proof of security in the standard model.

1 Introduction

A well-known fact in the context of designing authentication systems is that human users/clients typically choose “weak”, low-entropy passwords from a relatively small space of possibilities. Unfortunately, protocols designed and proven secure for the case when clients use *cryptographic* (i.e., high-entropy) secrets are generally insecure when passwords (i.e., low-entropy secrets) are used instead; this is so because these protocols are typically not resistant to *off-line dictionary attacks* in which an eavesdropping adversary derives information about the password from observed transcripts of login sessions. In recent years, much attention has focused on designing password-based authenticated key-exchange protocols resistant to such attacks. (We remark that *on-line dictionary attacks* – in which an adversary simply attempts to log-in repeatedly, trying each possible password – cannot be prevented by cryptographic means but can be dealt with using other methods outside the scope of this work.)

* Supported by NSF CAREER award 0447075 and Trusted Computing grant 0310751.

Means of protecting against off-line dictionary attacks in a single-server setting were first suggested by Gong, et al. [23] in a “hybrid”, PKI-based model in which users are assumed to know the server’s public key in addition to a password. Bellare and Merritt [5] were the first to suggest protocols for what we term *password-only authenticated key exchange* (PAKE), where the clients are required to “store” (i.e., remember) *only* a short password and no additional information. These initial works (and others [6, 25, 30, 38]) were relatively informal and did not provide definitions or proofs of security. More recently, formal definitions and provably-secure protocols for the “hybrid” model have been given [7, 24], followed soon thereafter by formal models for the password-only setting [1, 8, 22] and associated protocols with proofs of security in the random oracle/ideal cipher models¹ [1, 8, 31] and in the standard model [20, 22, 27, 28]. (The protocols of [20, 27, 28] assume some public information which is available to all parties. Note, however, that since this information can be hard-coded into implementations of the protocol, clients do not need to memorize or store any high-entropy, cryptographic information as they are required to do in the PKI-based setting.)

Although the above protocols protect against off-line dictionary attacks, they do nothing to mitigate the concern that an adversary might obtain users’ passwords via *server compromise*. Such attacks represent a serious threat since they are potentially cost-effective (in that an adversary might be able to obtain thousands of users’ passwords by corrupting a single, poorly-protected server) and users frequently utilize the same password at multiple sites. Unfortunately, it is easy to show the *impossibility* of protecting against server compromise when a single server holds the necessary information to authenticate a user (assuming the only secret information held by the user is a low-entropy password). To protect against server compromise, Ford and Kaliski [19] thus proposed a *threshold* protocol in the PKI-based model, in which the authentication functionality is distributed across n servers who must all cooperate to authenticate a user. Their protocol remains secure (and, in particular, an adversary learns nothing about users’ passwords other than what it learns from its on-line password guesses) as long as $n - 1$ or fewer servers are compromised. Jablon [26] subsequently suggested a protocol with similar functionality in the password-only setting. Neither of these works, however, include rigorous definitions or proofs of security.

A number of provably-secure protocols for threshold password-based authentication have recently appeared. We summarize what is known:

- MacKenzie, et al. [35] showed a protocol in the “hybrid”, PKI-based setting which requires only t (out of n) servers to cooperate in order to authenticate a user, for any values of t, n (of course, security is only obtained as long as $t - 1$ or fewer servers are compromised). They prove security for their protocol in the random oracle model.

¹ In the random oracle model [2], parties are assumed to have “black-box” access to a random function. (The ideal cipher model assumes that parties have “black-box” access to a random keyed permutation, an even stronger assumption.) In practice, the random oracle is instantiated with a cryptographic hash function. It is known [10], however, that protocols secure in the random oracle model may not be secure for *any* choice of hash function.

- Di Raimondo and Gennaro [16] proposed a protocol in the password-only setting with a proof of security in the standard model. (A second protocol given in their paper, which we will not discuss further, achieves the weaker functionality in which the same session key is computed by all servers.) However, their protocol requires less than $1/3$ of the servers to be compromised (i.e., they require $n > 3t$) and thus does not give a solution for the two-server case². We remark that, in general, threshold cryptosystems for the two-party case do not follow immediately from threshold solutions for the case of general n ; see, e.g., the work of [17, 21, 32–34] in this regard.
- Brainard, et al. [9] have developed a two-server protocol (called “Nightingale” and being shipped by RSA Security), a variant of which has been proven secure in the random oracle model [37]. These protocols assume the PKI-based setting: as stated in the papers, the protocols assume a “secure channel” between the client and the server(s) which would in practice be implemented using public-key techniques such as SSL.

1.1 Our Contributions

We show here a two-server protocol for password-only authenticated key exchange, with proof of security in the standard model. Ours is the first *provably-secure* two-server protocol in the password-only setting, and is the first two-server protocol (in any setting) with a proof of security in the standard model.

Our protocol extends and builds upon the (non-threshold) protocol of Katz-Ostrovsky-Yung [28] and a more efficient variant of this protocol – called, for brevity, KOY* – described by Canetti, *et al.* [11] (we also introduce an additional modification that makes the protocol even more efficient). In Section 3 we describe a “basic” two-server protocol which is secure against a passive (i.e., “honest-but-curious”) adversary who has access to the entire state of one of the servers throughout its execution of the protocol, but cannot cause this server to deviate from its prescribed behavior. We believe this protocol is interesting in its own right (when the assumption on adversarial behavior is warranted), and anyway the basic protocol and its proof of security serve as a useful prelude to our second result. In Section 4 we show how to modify the basic protocol so as to achieve security against an *active* adversary who may arbitrarily deviate from the protocol. The changes we make consist of having the servers (efficiently) prove to each other that their computations were performed according to the protocol. Here, we make use of techniques developed by MacKenzie [32] for performing these proofs efficiently even in a concurrent setting.

The protocols we construct are relatively efficient. Each party in the basic two-server protocol performs roughly twice the amount of work as in the KOY* protocol. For the protocol secure against active adversaries, the work of the client stays the same but the work of the servers increases by a factor of roughly 2–4.

² The approach in their paper does not extend to the case $t \geq n/3$. The authors mention (without details) that “[i]t is possible to improve the fault-tolerance to $t < n/2 \dots$ ”, but even this would not imply a two-server solution.

2 Definitions and Preliminaries

We assume the reader is familiar with the model of [1] (building on [3, 4]) for password-based key exchange in the single-server case. Here, we generalize their model and present formal definitions for the case of two-server protocols. While the model presented here is largely equivalent to the model proposed by MacKenzie, et al. [35] (with the main difference that we do not assume a PKI), we can simplify matters a bit since we focus on the two-server setting exclusively. For convenience we describe the model for the case of a “passive” adversary first and then discuss briefly the modifications needed in the case of “active” adversaries. (As discussed below, in both the “passive” and “active” cases the adversary is free to interfere with all communication between the client and the servers. These cases only differ in the power of the adversary to control the actions of the corrupted servers: specifically, a “passive” adversary is unable to control the actions of corrupted servers, whereas a “active” adversary can.)

We first present a general overview of the system as we imagine it. For simplicity only, we assume that every client C in the system shares its password pw with exactly two servers A and B . In this case we say that servers A and B are *associated with* C . (Note that a single server may be associated with multiple clients.) In addition to holding password shares, these servers may also be provisioned with arbitrary other information which is *not* stored by C . Any such information is assumed to be provisioned by some incorruptible, central mechanism (a system administrator, say) at the outset of the protocol. Note that this does *not* represent an additional assumption or restriction in practice: the servers *must* minimally be provisioned with correct password shares anyway, and there is no reason why additional information cannot be provided to the servers at that time (in particular, the servers’ password shares are already high-entropy values, and the servers have no restriction – as the client does – on the amount of information they can store). An (honest) execution of a password-based key-exchange protocol between client C and associated servers A and B should result in the client holding independent session keys $sk_{C,A}$, $sk_{C,B}$, and servers A and B holding $sk_{A,C}$ and $sk_{B,C}$, respectively, with $sk_{C,A} = sk_{A,C}$ and $sk_{C,B} = sk_{B,C}$.

2.1 Passive Adversaries

We first describe the adversarial model under consideration. We assume an adversary who corrupts some servers at the outset of the protocol, such that for any client C at most one of the servers associated with C is corrupted. In the case of a passive adversary, a corrupted server continues to operate according to the protocol, but the adversary may monitor its internal state. Following [16, 35], we make the (largely conceptual) assumption that there exists a single *gateway* which is the only entity that *directly* communicates with the clients. This gateway essentially acts as a “router”, splitting messages from the clients to each of the two associated servers and aggregating messages from these servers to the client; we also allow the gateway to perform some simple, publicly-computable operations. Introduction of this gateway is not strictly necessary; however, it

enables a simplification of the formal model and also provides a straightforward way to quantify the number of on-line attacks made by an adversary.

During the course of the protocol, then, there may potentially be three types of communication: between the clients and the gateway, between the servers and the gateway, and between the servers themselves. We assume the adversary has the ability to eavesdrop on all of these. We further assume that the client-gateway communication is under the full control of the adversary, and thus the adversary can send messages of its choice to either of these entities, or may tamper with, delay, refuse to deliver, etc. any messages sent between clients and the gateway. On the other hand, in the case of a passive adversary we assume that the server-gateway and server-server communication is determined entirely by the protocol itself (i.e., a corrupted server follows the protocol exactly as specified)³. In addition, the adversary is assumed to see the entire internal state of any corrupted servers throughout their execution of the protocol. With the above in mind, we proceed to the formal definitions.

Participants, passwords, and initialization. We assume a fixed set of protocol participants (also called principals) each of which is either a client $C \in \text{Client}$ or a server $S \in \text{Server}$, where Client and Server are disjoint. Each $C \in \text{Client}$ is assumed to have a password pw_C chosen uniformly and independently from a “dictionary” of size N ⁴. (In fact, we further simplify matters and assume that passwords are chosen from the set $\{1, \dots, N\}$.) As noted earlier, we make the simplifying assumption that each client shares his password with exactly two other servers (and no more). If client C shares his password with the distinct servers A, B , then A (resp., B) holds a *password share* $pw_{C,A}$ (resp., $pw_{C,B}$); the mechanism for generating these shares depends on the protocol itself. We also allow each server to hold information in addition to these password shares. For example, as described in footnote 3, two servers associated with a particular client may be provisioned with a shared, symmetric key. As we have already discussed, the initialization phase during which this information is provisioned is assumed to be carried out by some trusted authority. Any information stored by a corrupted server is available to the adversary.

In general, additional information can be generated during the initialization phase. For example, in the “hybrid” password/PKI model [7, 24] public/secret key pairs are generated for each server and the secret key is given as input to the appropriate server, while the public key is provided to the appropriate client(s). For the protocol presented here, we require only the weaker requirement of a single set of public parameters which is provided to all parties.

³ For the case of a passive adversary, the assumption that the adversary cannot tamper with the server-server communication is easy to realize via standard use of message authentication codes or signatures (as the servers can store long-term keys for this purpose). The assumption that the adversary cannot tamper with the server-gateway communication is essentially for convenience/definitional purposes only, as the adversary can anyway tamper with client-gateway communication (and the gateway processes messages in a well-defined and predictable way).

⁴ As in other work, though, our proof of security may be adapted to handle arbitrary dictionaries and arbitrary distributions over these dictionaries.

Execution of the protocol. In the real world, a protocol determines how principals behave in response to input from their environment. In the formal model, these inputs are provided by the adversary. Each principal is assumed to be able to execute the protocol multiple times (possibly concurrently) with different partners; this is modeled by allowing each principal to have an unlimited number of *instances* [1, 4] with which to execute the protocol. We denote instance i of principal U as Π_U^i . A given instance may be used only once. The adversary is given oracle access to these different instances; furthermore, each instance maintains (local) state which is updated during the course of the experiment. In particular, each instance Π_U^i has associated with it the following variables, initialized as NULL or FALSE (as appropriate) during the initialization phase:

- sid_U^i , pid_U^i , and sk_U^i are variables containing the *session id*, *partner id*, and *session key(s)* for an instance, respectively. A client’s partner id will be a set of two servers; a server’s partner id will be a single client (viewed as a set for notational convenience). For C a client, sk_C^i consists of a pair $\text{sk}_{C,A}^i, \text{sk}_{C,B}^i$, where these are the keys shared with servers A and B , respectively. A server instance Π_S^i with partner C has only a single session key $\text{sk}_{S,C}^i$.
- term_U^i and acc_U^i are boolean variables denoting whether a given instance has terminated or accepted, respectively. state_U^i records any state necessary for execution of the protocol by Π_U^i .

As highlighted earlier, the adversary is assumed to have complete control over all communication between the client and the gateway. This is modeled via access to *oracles* which are essentially as in [1] and are described formally in the full version of this paper. Briefly, these include various Send oracles modeling “active” attacks in which the adversary tampers with communication between the client and the servers; an Execute oracle modeling passive eavesdropping attacks; a Reveal oracle which models possible leakage of session keys; and a Test oracle used to define security.

Sessions ids, partnering, correctness, and freshness. Session ids in our protocol are defined in a natural way, which then allows us to define notions of partnering, correctness, and freshness. Due to lack of space the details appear in the full version.

Advantage of the adversary. Informally, the adversary succeeds if it can guess the bit b used by the Test oracle on a “fresh” instance associated with a non-corrupted participant. Formally, we say an adversary A *succeeds* if it makes a single query $\text{Test}(U, U', i)$ regarding a fresh key $\text{sk}_{U,U'}^i$, and outputs a single bit b' with $b' = b$ (recall that b is the bit chosen by the Test oracle). We denote this event by Succ. Note that restricting the adversary to making its Test query regarding a fresh key is necessary for a meaningful definition of security. The advantage of adversary A in attacking protocol P is then given by:

$$\text{Adv}_{A,P}(k) \stackrel{\text{def}}{=} 2 \cdot \Pr[\text{Succ}] - 1,$$

where the probability is taken over the random coins used by the adversary as well as the random coins used during the course of the experiment.

It remains to define a secure protocol, as a PPT adversary can always succeed by trying all passwords one-by-one in an on-line impersonation attack. Informally, a protocol is secure if this is the best an adversary can do. Formally, we define in the full version what it means for an instance to represent an *on-line attack* (which, in particular, will not include instances used in *Execute* queries). The number of on-line attacks bounds the number of passwords the adversary could have tried in an on-line fashion, motivating the following definition:

Definition 1. *Protocol P is a secure two-server protocol for password-only authenticated key-exchange if there exists a constant c such that, for all dictionary sizes N and for all PPT adversaries A making at most $Q(k)$ on-line attacks and corrupting at most one server associated with each client, there exists a negligible function $\varepsilon(\cdot)$ such that $\text{Adv}_{A,P}(k) \leq c \cdot Q(k)/N + \varepsilon(k)$.*

Of course, we would optimally like to achieve $c = 1$; however, as in previous definitions and protocols [1, 22, 38] we will allow $c \neq 1$ as well. The proof of security for our protocol shows that we achieve $c = 2$, indicating that the adversary can (essentially) do no better than guess *two* passwords during each on-line attack.

Explicit mutual authentication. The above definition captures the requirement of *implicit* authentication only (and the protocol we present here achieves only implicit authentication). Using standard techniques, however, it is easy to add explicit authentication to any protocol achieving implicit authentication.

2.2 Active Adversaries

The key difference in the active case is that the adversary may now cause any corrupted servers to deviate in an arbitrary way from the actions prescribed by the protocol. Thus, if a server is corrupted the adversary controls all messages sent from this server to the gateway as well as messages sent from this server to any other server. As in the passive case, however, we continue to assume that communication between the gateway and any non-corrupted servers (as well as communication between two non-corrupted servers) is *not* under adversarial control. See footnote 3 for the rationale behind these conventions.

3 A Protocol Secure Against Passive Adversaries

3.1 Description of the Protocol

We assume the reader is familiar with the decisional Diffie-Hellman (DDH) assumption [15], strong one-time signature schemes, and the Cramer-Shoup encryption scheme [14] with labels. A high-level depiction of the protocol is given in Figures 1–3, and a more detailed description, as well as some informal discussion about the protocol, follows.

Initialization. During the initialization phase, we assume the generation of public parameters (i.e., a common reference string) which are then made available to

all parties. For a given security parameter k , the public parameters will contain a group \mathbb{G} (written multiplicatively) having prime order q with $|q| = k$; we assume the hardness of the DDH problem in \mathbb{G} . Additionally, the parameters include random generators $g_1, g_2, g_3, h, c, d \in \mathbb{G}^\times$ and a hash function $H : \{0, 1\}^* \rightarrow \mathbb{Z}_q$ chosen at random from a collision-resistant hash family.

As part of the initialization, each server S is provisioned with an El Gamal [18] public-/secret-key pair (pk_S, sk_S) , where $pk_S = g_1^{sk_S}$. If A and B are associated with the same client C , then A (resp., B) is given pk_B (resp., pk_A). We stress that, in contrast to the PKI-based model, the client is not assumed or required to know the public keys of any of the servers.

Passwords and password shares are provisioned in the following way: a password pw_C is chosen randomly for each client C and we assume that this password can be mapped in a one-to-one fashion to \mathbb{Z}_q . If A and B are the servers associated with a client C , then password shares $pw_{A,C}, pw_{B,C} \in \mathbb{Z}_q$ are chosen uniformly at random subject to $pw_{A,C} + pw_{B,C} = pw_C \bmod q$, with $pw_{A,C}$ given to server A and $pw_{B,C}$ given to server B . In addition, both A and B are given $\text{Com}_{A,C}, \text{Com}'_{A,C}, \text{Com}_{B,C}$, and $\text{Com}'_{B,C}$, where:

$$\begin{aligned} \text{Com}_{A,C} &\stackrel{\text{def}}{=} \text{ElG}_{g_3}(g_1^{pw_{A,C}}) = (g_1^{r_a}, g_3^{r_a} g_1^{pw_{A,C}}) \\ \text{Com}'_{A,C} &\stackrel{\text{def}}{=} \text{ElG}_{pk_A}(g_1^{pw_{A,C}}) \\ \text{Com}_{B,C} &\stackrel{\text{def}}{=} \text{ElG}_{g_3}(g_1^{pw_{B,C}}) = (g_1^{r_b}, g_3^{r_b} g_1^{pw_{B,C}}) \\ \text{Com}'_{B,C} &\stackrel{\text{def}}{=} \text{ElG}_{pk_B}(g_1^{pw_{B,C}}). \end{aligned}$$

Note that different public keys are used. Server A (resp., server B) is additionally given the randomness r_a (resp., r_b) used to construct $\text{Com}_{A,C}$ (resp., $\text{Com}_{B,C}$).

Protocol execution. At a high level one can view our protocol as two executions of the KOY* protocol, one between the client and server A (using server B to assist with the authentication), and one between the client and server B (using server A to assist with the authentication). Note that the assistance of the other server is necessary since the password information is split between the two servers. For efficiency, the signature and verification for the two executions are combined, and shares of the El Gamal encryption of the password sent by the servers (i.e., (F, G) in Figure 1) are fixed and stored by the servers.

When a client with password pw_C wants to initiate an execution of the protocol, this client computes Cramer-Shoup “encryptions” of pw_C for each of the two servers. In more detail (cf. Figure 1), the client begins by running a key-generation algorithm for a one-time signature scheme, yielding VK and SK. The client next chooses random $r_1 \in \mathbb{Z}_q$ and computes $A_a = g_1^{r_1}$, $B_a = g_2^{r_1}$, and $C_a = h^{r_1} \cdot pw_C$. The client then computes $\alpha_a = H(\text{Client}|\text{VK}|A_a|B_a|C_a)$ and sets $D = (cd^{\alpha_a})^{r_1}$. This exact procedure is then carried out again, using an independent random value $r_2 \in \mathbb{Z}_q$. The client sends

$$\text{msg}_1 \stackrel{\text{def}}{=} \langle \text{Client}, \text{VK}, A_a, B_a, C_a, D_a, A_b, B_b, C_b, D_b \rangle$$