05411 Abstracts Collection
Anonymous Communication and its Applications
— Dagstuhl Seminar —

Shlomi Dolev[1], Rafail Ostrovsky[2] and Andreas Pfitzmann[3]

[1] Ben Gurion Univ., IL
dolev@CS.bgu.ac.il
[2] Univ. California - Los Angeles, US
rafail@cs.ucla.edu
[3] TU Dresden, DE
pfitza@inf.tu-dresden.de

Abstract. From 09.10.05 to 14.10.05, the Dagstuhl Seminar 05411 "Anonymous Communication and its Applications" was held in the International Conference and Research Center (IBFI), Schloss Dagstuhl. During the seminar, several participants presented their current research, and ongoing work and open problems were discussed. Abstracts of the presentations given during the seminar as well as abstracts of seminar results and ideas are put together in this paper. The first section describes the seminar topics and goals in general. Links to extended abstracts or full papers are provided, if available.

Keywords. Anonymous Communication, Cryptography, Privacy, Security, Anonymity

Traffic analysis without direct observations

George Danezis (University of Leuven, B)

Anonymity system designers have traditionally tried to prevent global passive observers from linking communicating parties. Yet recently some proposed anonymous channels only protect from partial observers and assume that no real-world adversary can be powerful enough to observe all communications. We consider the challenge facing an adversary performing traffic analysis without having a direct view of the totality of the network. We show that the anonymity system can provide enough 'services' to this adversary to still be susceptible to some tracing, and present two attacks: one that degrades the anonymity of The Onion Router (Tor), and one that breaks the anonymity of a scheme based on universal re-encryption.

Keywords: Anonymity, Traffic analysis

Dagstuhl Seminar Proceedings 05411
Anonymous Communication and its Applications
http://drops.dagstuhl.de/opus/volltexte/2006/795

Improving the Decoding Efficiency of Private Search

George Danezis (University of Leuven, B)

We show two ways of recovering all matching documents, in the Ostrovsky et al. Private Search, while requiring considerably shorter buffers. Both schemes rely on the fact that documents colliding in a buffer position provide the sum of their plaintexts. Effcient decoding algorithms can make use of this property to recover documents never present alone in a buffer position.

Keywords: Private search, private information retrieval, cryptography

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2006/482

Anonymity Metrics Revisited

Claudia Diaz (University of Leuven, B)

In 2001, two information theoretic anonymity metrics were proposed: the "effective anonymity set size" and the "degree of anonymity". In this talk, we propose an abstract model for a general anonymity system which is consistent with the definition of anonymity on which the metrics are based. We revisit entropy-based anonymity metrics, and we apply them to Crowds, a practical anonymity system. We discuss the differences between the two metrics and the results obtained in the example.

Keywords: Anonymity, metrics, entropy

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2006/483

Tor: An Internet Anonymous Communication System

Roger Dingledine (The Free Haven Project)

What do the United States Department of Defense and the Electronic Frontier Foundation have in common? They are both funding the development of Tor (tor.eff.org), a free-software onion routing network that helps people around the world use the Internet in safety.

The public Tor network has over 250 servers on six continents, and averages over 100Mbit/s of traffic. Our users include ordinary citizens who want protection from identity theft and prying corporations, corporations who want to look at a competitor's website in private, and soldiers and aid workers in the Middle East who need to contact their home servers without fear of physical harm.

I'll give an overview of the Tor architecture, and talk about why you'd want to use it, what security it provides, and how user applications interface to it. Then we can start talking about whether Tor's free-route topology provides a good trade-off in terms of security and flexibility.

Research questions for Tor, or, Anonymity Problems I want solved

Roger Dingledine (The Free Haven Project)

Over the past few years we have begun to investigate such hard questions as how to quantify the level of security an anonymity system provides in various contexts and configurations. We have also begun to develop and deploy real anonymity networks, such as Tor and JAP. What are the next research steps in understanding anonymity, and what answers do we need before we can move on to building better deployed systems?

This talk has two parts. First we present the near-future questions we're asking and investigating right now, and then we present the longer-term questions that will likely require substantial changes or entirely new designs.

On Extractors, Error-Correction and Hiding All Partial Information

Yevgeniy Dodis (Courant Institute - New York, USA)

Randomness extractors allow one to obtain nearly perfect randomness from highly imperfect sources randomness, which are only known to contain "scattered" entropy. Not surprisingly, such extractors have found numerous applications in many areas of computer science including cryptography. Aside from extracting randomness, a less known usage of extractors comes from the fact that they hide all deterministic functions of their (high-entropy) input: in other words, extractors provide certain level of privacy for the imperfect source that they use. In the latter kind of applications, one typically needs extra properties of extractors, such as invertibility, collision-resistance or error-correction.

In this abstract we survey some of such usages of extractors, concentrating on several recent results by the speaker. The primitives we will survey include several flavors of randomness extractors, entropically secure encryption and perfect one-way hash functions. The main technical tools will include several variants of the leftover hash lemma, error correcting codes, and the connection between randomness extraction and hiding all partial information.

The copy of the survey can be found at http://theory.lcs.mit.edu/ yevgen/ps/ent-survey.ps

Keywords: Randomness extractors, biometrics, entropic security, fuzzy extractors, secure sketches, hiding all functions, error-correction

Full Paper:
http://theory.lcs.mit.edu/~yevgen/ps/ent-survey.ps

Anonymity and Beyond

Shlomi Dolev (Ben Gurion University, IL)

The talk summarizes three recent works:

- Buses for Anonymous Message Delivery (with Amos Beimel, J. Cryptology 2003).
- Polygonal Broadcast, Secret Maturity and the Firing Sensors, (With Ted Herman and Limor Lahiani, Add-Hoc Networks Journal 2005).
- Magnifying Computing Gaps (with Ephraim Korach and Galit Uzan, submitted).

Anonymous communication is achieved by using fixed routing and scheduling of "buses" (messages) that carry encrypted messages (seats) from every source to every destination. One such scheme uses a single bus with $O(n^2)$ seats that traverse the network using an Euler tour. Here there is only one message but the delay is $O(n)$.

The other extreme is when we have 2|E| buses, two buses for each link, traversing the link in opposite directions. A bus of an edge have seats for every source destination pair that choose the edge as part of their shortest path. Here there $O(|E|)$ messages but the transmission time is optimal. We present a cluster based schemes that trade time with communication. We also use random walk to cope with dynamic changes.

The second work is in the domain of sensor networks. The problem of simultaneous activation is considered where a captured sensor should not reveal the up-coming activity even during the time the encrypted command is distributed. Thus, the sensor should not know the nature of the command before it has to execute it. Unidirectional communication of satellite is used. The satellite repeatedly broadcasts public key and after a period of time larger than the command flooding time the satellite broadcasts the coupled private key. A sensor that would like to send an activation command, encrypts the command using a public key received from the satellite and flood the system with the encrypted command. When the sensors receive the satellite broadcast with the coupled private key they simultaneous reveal the command and act accordingly.

The third work scope is unidirectional encryption. Consider two, not necessarily identical, computationally powerful machines connected by a unidirectional communication link that should transfer a long stream of information in the presence of listening adversary that is slightly weaker. We present schemes that enhance the computation strength gap of the combinatorial exhaustive testing machines and the adversary. In other words, the gap between the amount of information decrypted by the adversary and the information decrypted by the receiver grows with time.

Abstracts and Slides

Stefan Köpsell (TU Dresden, D)

Abstract "AN.ON – a real world anonymity service"

The talk presents a real implementation of a public available anonymity service called AN.ON. The goals, fundamental architecture and the protocol details will be explained. At the end current research tasks will be introduced.

Abstract "Social Impacts of Anonymity Systems – more than Crime"

If one thinks about the impacts of anonymity systems on society one often thinks of child pornography, terrorism or at least crime at large. Of course this kind of impact exists and so the talk will focus on that problem. But the talk will also show that the problem is not that big as one may think it could be. Additional the talk will give at leas one example for a completely different influence on society.

Abstract "Challenges in real world anonymity systems"

There exist many theoretical approaches for anonymous communication in computer networks since more than two decades – like Chaum's Mixes, the DC-Net, Broadcast, Private Message Service, ISDN Mixes etc. But building an anonymous communication system for the real world is much more than simply implementing these theoretical approaches. One problem is that these approaches are based on their very own models of the world, i.e. models for the underlying communication infrastructure, models of the behaviour and incentives of senders and receives etc. As these models do not adequately describe the real world one has two possibilities:

- try to change the real world so that it fits to the model – to the best of my knowledge this seems to be unrealistic at the moment
- change the model so that it better reflects the real world

The later one has many implications on the proposed solutions for anonymous communication. This talk tries to spotlight some of the arising problems hopefully leading to a discussion about solutions.

Keywords: Anonymous communication

A Formal Treatment of Onion Routing

Anna Lysyanskaya (Brown Univ. - Providence, USA)

Anonymous channels are necessary for a multitude of privacy-protecting protocols. Onion routing is the best known way to achieve anonymity in practice. The idea of onion routing is to wrap a message in several layers of encryption and then route it through a path of randomly selected intermediate nodes.

Each node peels off a layer of encryption and forwards the resulting ciphertext to the next node. This process is repeated until all layers are removed. In the process, each intermediate router on the path must find out its two adjacent routers, but the protocol must conceal all other information. The informal security requirement is that, by the time the message arrives at destination, its origins are impossible to trace, even if some of the routers on the path are maliciously trying to identify its originator. Despite substantial implementation efforts, the cryptographic aspects of onion routing had not been sufficiently explored prior to our work. No satisfactory definitions of security had been given, and prior constructions have only had ad-hoc security analysis for the most part.

We provide a formal definition of onion-routing and then give an efficient and easy to implement onion routing scheme satisfying this definition in the PKI model.

Joint work of: Lysyanskaya, Anna; Camenisch, Jan

Full Paper:
http://dx.doi.org/10.1007/11535218_11

See also: Appeared in Crypto 2005

Privacy of Statistical Databases and Sensitivity of Functions

Kobbi Nissim (Ben Gurion University, IL)

Consider a trusted server that holds a database of sensitive information. The server would like to reveal global statistics about the population in the database, and yet hide information specific to individuals. These are conflicting goals. Exactly how these goals are defined, and what tradeoffs exist between them, is the focus of this paper. These questions are complementary to the design of secure function evaluation protocols; the goal here is to decide which functionalities are actually safe to learn. We describe a new framework for releasing noisy functions of the database while provably satisfying a strong definition of privacy. We define the *sensitivity* of a function, roughly, the amount that any single argument to the function can change its output, and show how adding noise proportional to the sensitivity is sufficient to preserve privacy. This framework improves on recent work of Blum et al. [**?**], in two ways: (i) a much richer set of queries is allowed, and (ii) for several functions, a significantly lower noise level is needed for proving privacy.

Our second main contribution is a separation between the power of interactive and non-interactive privacy mechanisms. We show that for any non-interactive mechanism S satisfying our definition of privacy, there exist low-sensitivity functions f which cannot be approximated at all based on $S(x)$, unless the database is very large–exponential in the number of attributes. In other words, when the data is complex, a non-interactive mechanism must be tailored to suit certain functions to the exclusion of others.

Finally, we present privacy definitions modeled after the notions of indistinguishability, semantic security, and simulation, and explain their equivalence. These definitions simplify those of previous work.

Cryptography from Anonymity

Rafail Ostrovsky (Univ. California - Los Angeles, USA)

Hiding one's identity is often easier than hiding one's secrets. In particular, the goal of anonymity becomes easier to achieve when the number of users in the system grows, whereas satisfactory solutions to the problem of privacy (as in the context of secure multi-party computation) become expensive when the number of users is large.

We suggest a novel approach for obtaining information-theoretic privacy and security by relying on anonymous (but non-private) communication as a *building block*. Our results go in two directions:

- **Feasibility.** We show that the non-private anonymity functionality can be used to implement unconditionally secure point-to-point channels, as well as secure multi-party computation as long as less than *half* of the players are corrupted.
- **Efficiency.** We show that anonymous channels can yield dramatic efficiency improvement for several natural secure computation tasks. In particular, we present the first solution to the problem of private information retrieval (PIR) which can handle multiple users while being (essentially) optimal with respect to both communication and computation.

Joint work of: Ishai, Yuval; Kushilevitz, Eyal; Sahai, Amit, Ostrovsky, Rafail

Cryptography from Anonymity

Rafail Ostrovsky (Univ. California - Los Angeles, USA)

There is a vast body of work on implementing anonymous communication. In this paper, we study the possibility of using anonymous communication as a building block, and show that one can leverage on anonymity in a variety of cryptographic contexts. Our results go in two directions.

Feasibility. We show that anonymous communication over insecure channels can be used to implement unconditionally secure point-to-point channels, broadcast, and generalmulti-party protocols that remain unconditionally secure as long as less than half of the players are maliciously corrupted.

Efficiency. We show that anonymous channels can yield substantial efficiency improvements for several natural secure computation tasks. In particular, we present the first solution to the problem of private information retrieval (PIR) which can handle multiple users while being close to optimal with respect to both communication and computation.

Joint work of: Ishai, Yuval; Kushilevitz, Eyal; Sahai, Amit; Ostrovsky, Rafail

Full Paper:
http://www.cs.ucla.edu/∼rafail/PUBLIC/75.html

Private Searching On Streaming Data

Rafail Ostrovsky (Univ. California - Los Angeles, USA)

In this talk, I will describe the problem of private searching on streaming data, where we can efficiently implement searching for documents under a secret criteria (such as presence or absence of a hidden combination of hidden keywords) under various cryptographic assumptions. The result can be viewed in a variety of ways: as a generalization of the notion of a Private Information Retrieval (to the more general queries and to a streaming environment as well as to public-key program obfuscation); as positive results on privacy-preserving data mining; and as a delegation of hidden program computation to other machines.

Joint work of: Ostrovsky, Rafail; Skeith, William

History of anonymity

Andreas Pfitzmann (TU Dresden, D)

Paul Baran stated in 1964, if you have a centrally trusted party, you can achieve nearly everything (e.g. confidentiality, anonymity, unobservability, ...) using link-encryption and dummy traffic.

In 1981, David Chaum published the first approach how to achieve anonymity and unobservability as well as pseudonyms and voting without central trust, i.e. it is enough to trust any 1 out of n parties, called MIXes. The base technology is asymmetric cryptography, so security is computational at best.

In 1985, David Chaum published solutions for anonymous communication (DC net), payment, and credentials, when the privacy properties are unconditional, i.e. provable in Claude Shannon's work of information theory.

Around 1990, MIXes and DC net have been pushed to their limits, i.e. as secure as can be, by Michael Waidner, Birgit Pfitzmann, and myself.

But in reality, systems work with quite weaker security, since usability (including delay and throughput) are essential to attract users - and without users, there is no anonymity.

So for the future, I hope to see theory and practice to meet.

Upper and Lower Bounds on Black-Box Steganography

Leonid Reyzin (Boston University, USA)

We study the limitations of steganography when the sender is not using any properties of the underlying channel beyond its entropy and the ability to sample from it. On the negative side, we show that the number of samples the sender must obtain from the channel is exponential in the rate of the stegosystem. On the positive side, we present the first secret-key stegosystem that essentially matches this lower bound regardless of the entropy of the underlying channel. Furthermore, for high-entropy channels, we present the first secret-key stegosystem that matches this lower bound statelessly (i.e., without requiring synchronized state between sender and receiver).

Keywords: Steganography, information hiding

Joint work of: Dedić, Nenad; Itkis, Gene; Reyzin, Leonid; Russell, Scott

Full Paper:
http://www.cs.bu.edu/∼reyzin/bbstego.html

See also: TCC 2005

Private Communication in a Dark Network

Oskar Sandberg (Chalmers - Göteborg, S)

This talk explores a different way to protect the privacy of participants and publishers on the Internet. Rather than attempting to build a mixing network, we build a "Dark network" (or "Friend to Friend" network) where users only reveal there identity to trusted friends, and yet can communicate with anybody in the network. We achieve this by exploring methods of using models of the "small world phenomenon" to route in social networks.

Keywords: Darknet, freenet, small world

Enhancing the security od perfect blind DL-signatures

Claus Peter Schnorr (Universität Frankfurt, D)

We enhance the security of Schnorr blind signatures against the novelone-more-forgery of SCHNORR and WAGNER which is possible even if the discrete logarithm is hard to compute. We show two limitations of this attack. Firstly, replacing the group G by the s-fold direct product $G^{\times s}$ increases the work of the attack, for a given number of signer interactions, to the s-power while increasing the work of the blind signature protocol merely by a factor s. Secondly, we bound the number of additional signatures per signer interaction that can be efficiently forged by known methods. That fraction of the additional forged signatures can be made arbitrarily small. Our security proofs assume both the random oracle and the generic group model.

Message Splitting and Intersection Attacks

Andrei Serjantov (The Free Haven Project)

First, we introduced the intersection attack which is the loss of anonymity resulting from observing a pattern of repeated communication of a sender to a set of users which is a subset of all possible recipients of messages. We discuss exact and probabilistic versions of this attack briefly and show that it is effective against any anonymous channel, powerful and difficult to protect against.

In the second part of the talk, we look at how the situation changes if we consider a weaker adversary model in which the attacker is not able to observe all of the sender's traffic (or equivalently all the entry nodes). The sender is then able to split her traffic across entry nodes and reduce the effectiveness of the intersection attack.

Finally, we look at the fundamental tradeoff between delay and anonymity, show that this issue is necessary to consider in the design of any anonymity system and argue that this choice is currently made in an ad-hoc fashion. Instead, we advocate empirical studies as a way of tailoring the anonymity system to requirements of users and discuss possible problems with this approach.

A Brief History of Onion Routing

Paul Syverson (Naval Research - Washington, USA)

In this talk, I briefly sketched the history of the US Naval Research Loboratory's Onion Routing Program from 1995 to 2005. Onion Routing networks are application-independent systems for low-latency communication resistant to traffic analysis, e.g., web traffic, remote login, IRC chat, etc. The goal of Onion Routing has always been to separate identification and authentication from routing.

There have been three generations of Onion Routing that can be roughly characterized as follows.

Generation 0:

Basic system establishing route/circuit using an "onion" consisting of layers of public-key encryption effectively wrapped around nothing, and used to distribute symmetric keys used for passing data in both directions. Fixed five hop routes. Static network topology. Initial rendezvous service design linking two anonymous parties. Access to hidden location using reply onions. Slides describing the generation 0 design (with some generation 1 elements) available at http://www.onion-router.net/Publications.html#old-slides Also in those slides: rendezvous for IRC chat, applications to cell phones, private location tracking of active badge systems. Official plans to distribute code in order to leverage security of open source leading to and based on wider adoption from the beginning of the first project, before the phrase "open source" had been coined. (First publication release of code was given July 1996.) A publicly accessible prototype generation 0 system ran at NRL continuously from about late 1996 to January 2000. Publishing c. 50K connections per day over its last year from tens of thousands of IP addresses the world over.

Generation 1:

Clients are separated from routers and can be run at locations other than at routers. Variable entrance configurations and variable exit policies added. Mixing added to processing of cells (uniform-size data packets). Variable length routes. Separate crypto modules capable of running on dedicated hardware (and facilitates exportability in the legal environment of the day). Flat distribution of topology and linkstate information via database engine modules. Plan for smoothing of data rates via padding and limiting to reduce effectiveness of traffic analysis. A breakdown of the generation 1 modules is given at http://www.onion-router.net/Archives/TNG.html Several generation 1 test systems were deployed with up to c. a dozen nodes distributed across the US and Canada, but were not made publicly accessible.

Generation 2:

Tor (Tor's Onion Routing). Circuits built incrementally to permit use of ephemeral Diffie-Hellman (forward anonymity). Thus, onions per se are gone, although layering and separation of public-key circuit establishment from symmetric-key data passing remains. Design and implementation of hidden services via rendezvous points. Abandonment of any thought that mixing or padding are currently of use in low-latency systems. Leveraging of now widely used SOCKS protocol allows abandoning of most application-specific proxies that were needed in generation 1. Integrity checking. Congestion control. Multiple application streams on a single circuit. Directory servers to control network membership and maintain node status information. Ability for application traffic to exit circuits at any point (leaky pipes). Slides on the Tor design and project status as of mid 2004 can be found at http://www.onion-router.net/Publications.html#new-slides

The Tor network was the first version of Onion Routing not designed entirely by Naval Research Lab personnel. It was designed jointly with the Free Haven Project (http:
freehaven.net), under contract to NRL. The Tor network became operational in October 2003, and has never been down, as of this writing. It has grown to a present network in October 2005 of about 250 nodes on every continent but Antarctica with an estimated user base in the hundreds of thousands. (The design hides the exact number.) It was named one of the 100 best products of 2005 by PC World. Support for Tor development and maintainance was provided by the EFF from late 2004 to late 2005.

More details about the history of Onion Routing can be found at
http://www.onion-router.net/History.html

This work is the result of input from many people. Primary contributors to generations 0 and 1 Onion Routing besides myself are David M. Goldschlag and Michael G. Reed. Primary contributors to generation 2 Onion Routing are Roger Dingledine and Nick Mathewson.

(Note that the presentation on Onion Routing by Lysyankaya in this Dagstuhl seminar introduces an entirely different use of the term from how we have been using it for the last decade. A central aspect of our Onion Routing is that public-key is used only to establish a circuit. Actual messages are never sent using public-key. Also, our Onion Routing is bidirectional. Once a circuit is established, symmetric-key is used to protect data flowing in both directions.

Finally, Onion Routing derives anonymity primarily from unpredictability of routing. There is no mixing assumed to be taking place.)

Universal Re-encryption for Mixnets

Paul Syverson (Naval Research - Washington, USA)

We introduce a new cryptographic technique that we call *universal re-encryption*. A conventional cryptosystem that permits re-encryption, such as ElGamal, does so only for a player with knowledge of the public key corresponding to a given ciphertext. In contrast, universal re-encryption can be done without knowledge of public keys. We propose an asymmetric cryptosystem with universal re-encryption that is half as efficient as standard ElGamal in terms of computation and storage.

While technically and conceptually simple, universal re-encryption leads to new types of functionality in mixnet architectures. Conventional mixnets are often called upon to enable players to communicate with one another through channels that are *externally anonymous*, i.e., that hide information permitting traffic-analysis. Universal re-encryption lets us construct a mixnet of this kind in which servers hold *no public or private keying material*, and may therefore dispense with the cumbersome requirements of key generation, key distribution, and private-key management. We describe two practical mixnet constructions, one involving asymmetric input ciphertexts, and another with hybrid-ciphertext inputs.

Joint work of: Golle, Philippe; Jakobsson, Markus; Juels, Ari; Syverson, Paul

Provable Anonymity Against Traffic Analysis

Amnon Ta-Shma (Tel Aviv University, IL)

With the advent of peer to peer networks, anonymity is grasped as a desired property of any well designed system for exchanging information between parties. This work focuses on unlinkability, which is one of the possible types of anonymity.

Rackoff and Simon proved the unlinkability of an efficient variant of David Chaum's classic protocol for anonymous data communication, and proved that it is secure against adaptive adversaries. Their protocol is round efficient, but not message efficient (all players have to play at all rounds).

We show that if the attack model is somewhat relaxed (the adversary may control most but not all of the communication links) than one can replace node-mixing with layer mixing, and prove another variant of Chaunm's protocol is anonymous even when honest players play only when they want to send a message. Thus, our protocol is also message efficient.

Our main tool is information Theory. We use this connection to prove the protocol is anonymous even when the adversary has prior information. We are not aware of any previous work achieving that, and we believe the ability to deal with prior information is crucial for achieving anonymity in realistic applications.

Keywords: Anonymity, unlinkability, traffic analysis, node-mixing, layer-mixing

Joint work of: Berman, Ron; Fiat, Amos; Ta-Shma, Amnon

Full Paper:
http://www.springerlink.com/content/qvknd1d83hpur96j/?p=6a646068c5454ca3ba5202be439f4cf4&pi=3

See also: International Conference on Financial, pages 266-280, 2004

Oblivious Information Delivery using Oblivious Signature-Based Envelopes

Gene Tsudik (Univ. of California - Irvine, USA)

Secure, anonymous and unobservable communication is becoming increasingly important due to the gradual erosion of privacy in many aspects of everyday life. This prompts the need for various anonymity- and privacy-enhancing techniques, e.g., group signatures, anonymous e-cash and secret handshakes.

In this work we investigate an interesting and practical cryptographic construct – Oblivious Signature-Based Envelopes (OSBEs) recently introduced by Li, et al. (PODC'03). OSBEs are very useful in anonymous communication since they allow a sender to communicate information to a receiver such that the receiver's rights (or roles) are unknown to the sender. At the same time, a receiver can obtain the information only if it is authorized to access it. This makes OSBEs a natural fit for anonymity-oriented and privacy-preserving applications, such as Automated Trust Negotiation and Oblivious Subscriptions. Our work focuses on the ElGamal signature family: we succeed in constructing practical and secure OSBE schemes for several well-known signature schemes. Our schemes are more efficient than previous techniques and are practical in the true meaning of the word.

Keywords: Anonymity, secret handshakes, untracaeability, privacy, authentication

Joint work of: Tsudik, Gene; Nasserian, Samad

A Flexible Framework for Secret Handshakes or How to Achieve Multi-Party Interactive Anonymous Authentication

Gene Tsudik (Univ. of California - Irvine, USA)

In the society increasingly concerned with the erosion of privacy, privacy-preserving techniques are becoming very important.

Secret handshakes offer anonymous and unobservable authentication and serve as an important tool in the arsenal of privacy-preserving techniques.

Relevant prior research focused on 2-party secret handshakes with one-time credentials, whereby two parties establish a secure, anonymous and unobservable communication channel only if they are members of the same group.

This paper breaks new ground on two accounts: (1) it shows how to obtain secure and efficient secret handshakes with reusable credentials, and (2) it provides the first treatment of multi-party secret handshakes, whereby m>2 parties establish a secure, anonymous and unobservable communication channel if they all belong to the same group. An interesting new issue encountered in multi-party secret handshakes is the need to ensure that all parties are indeed distinct. (This is a real challenge since the parties cannot expose their identities.) We tackle this and other challenging issues in constructing GCD – a flexible secret handshake framework. GCD can be viewed as a "compiler" that transforms three main building blocks: (1) a Group signature scheme, (2) a Centralized group key distribution scheme, and (3) a Distributed group key agreement scheme, into a secure multi-party secret handshake scheme.

The proposed framework lends itself to multiple practical instantiations, and offers several novel and appealing features such as self-distinction and strong anonymity with reusable credentials. In addition to describing the motivation and step-by-step construction of the framework, this work provides a security analysis and illustrates several concrete framework instantiations.

Keywords: Secret handshakes, multi-party protocols, anonymity, untraceability, authentication

Joint work of: Tsudik, Gene; Xu, Shouhuai

Private Approximation of Search Problems

Enav Weinreb (Ben Gurion University, IL)

Many approximation algorithms have been presented in the last decades for hard search problem. In cryptographic applications, it is desired to design algorithms which do not leak unnecessary information. Specifically, we are interested in private approximation algorithms - efficient algorithms whose output does not leak information not implied by the optimal solutions to the search problems. Known approximation algorithms usually leak a lot of information.

In this talk I will discuss how to define private approximation of search problems, I will present some impossibility results for the vertex cover problem, and I will describe some private algorithms that leak little information for MAX-3SAT.

Joint work of: Beimel, Amos; Carmi, Paz; Nissim, Kobbi; Weinreb, Enaw