6

MMIP – Mixed Mobile Internet Protocol

T. Lopatic, C. Eckert and U. Baumgarten
Munich University of Technology, Department of Computer Science
D-80290 Munich, Germany,
{eckertc,baumgaru}@informatik.tu-muenchen.de

Abstract

Mobile IP extends the Internet Protocol to provide mobility features. Unfortunately, Mobile IP introduces new security threats unknown in IP environments. In this paper we emphasize on the problem of ensuring location privacy for mobile users which is not addressed by Mobile IP. To conceal movements of mobile users and to prevent an attacker from tracing users' locations we designed and implemented a simple extension of Mobile IP which is called Mixed Mobile IP. The basic idea of the proposed extension is to separate information about the identity of a host from information about its location. The paper presents the main features of the protocol and sketches our Linux-based prototype implementation.

Keywords

Communication Security, Mobile IP, Anonymity, Mix

1 INTRODUCTION

With the evolving of mobile computing (Forman,1994) the vision of accessing any service anywhere at any time becomes more and more realistic. The first step towards this ultimate goal is to enable mobile hosts to move freely around and attach themselves—preferably via a wireless link—to the network wherever they want. Many applications and services operate based upon protocols of the Internet protocol suite (e.g. IP, TCP, HTTP). Therefore, it is useful to add to it features of mobility. In order to handle mobility without having to change the facilities of IP the *Mobile Internet Protocol* (Mobile IP) was defined (Perkins,1996), which extends IP and adds some administrative components.

Obviously it is desirable to keep mobile IP networks as secure as their stationary counterparts, for which many problems like the threat of impersonation and eavesdropping have been investigated. Several solutions have already been proposed like authentication standards to defeat masquerading and replay attacks or standards for privacy enhancements to prevent eaves-

Communications and Multimedia Security Vol.3 S. Katsikas (Ed.)

dropping attacks by using encryption and to ensure message integrity by computing digital fingerprints of messages. The proposed solutions can equally be applied to mobile IP networks. Unfortunately, some new problems arise within the context of mobility. Among others, location privacy, i.e. preventing the tracing of a mobile user's point of attachment to the network and thus his or her location and movements, is an open problem. Our solution for this threat, the *Mixed Mobile IP*, together with its implementation in a Linux-based environment will be described in this paper.

The rest of the paper is organized as follows. Section 2 presents a brief overview over Mobile IP its security issues and some related work. The design and key features of the Mixed Mobile IP are introduced in section 3. Our Linux implementation is sketched in section 4. The conclusion summarizes the main issues of the paper and gives an outlook on future work.

2 BACKGROUND

2.1 The Mobile IP

In the context of Mobile IP mobile hosts which are attached to a mobile network are referred to as *mobile nodes*. Each mobile node has a *home network*. The home network of a mobile node is the network the mobile node derives its IP address, the *home address*, from. When a mobile node is *at home*, i.e. attached to its home network, the standard IP routing mechanism applies, since the IP address of the node corresponds to the network it is attached to. In the home network a stationary host or router, the *home agent*, must be prepared to provide mobility services to the mobile node. When a mobile node detaches from its home network and starts to move around, any network it may *visit* is referred to as a *foreign network*. In order to enable a network to act as a foreign network a stationary host or router, the *foreign agent*, must be prepared to provide mobility services for visiting mobile nodes. Foreign agents and home agents are also referred to as *mobility agents*.

The home agent intercepts in the home network all datagrams addressed to the mobile node and forwards them to the foreign network which the mobile node is attached to, making use of IP in IP tunneling. The tunnel ends at the *care-of address* of the mobile node, which typically belongs to the foreign agent. There the datagram is decapsulated and forwarded to the mobile node.

Whenever a mobile node MN attaches to a foreign network, it informs its home agent HA of the care-of address FA to use for tunneling by sending a registration request of the form (MN, HA, FA) to the foreign agent, which forwards it to the home agent. After processing the registration request, the home agent returns a registration reply (MN, HA), which is relayed by the foreign agent to the mobile node. If the registration is successful, the mobile

node has established a *mobility binding*, i.e. an association of its home address with a care-of address.

2.2 Security Issues in Mobile IP

Faced with this scenario several security threats immediately strike the eye. *Authentication*: there must be a mutually authenticated binding between the mobile node, the home agent and the foreign agent to prevent malicious users from masquerading and stealing information. *Communication privacy* and *message integrity*: wireless links used to attach mobile hosts to networks are more vulnerable to eavesdropping attacks as well as active attacks than conventional networks are. In addition, the participating mediators of messages, i.e. the mobility agents, are by no means trustworthy. Message contents must be concealed to defend attacks concerning confidentiality. *Anonymity*: traffic analysis attacks are well known threats in conventional network environments. An attacker tries to gain knowledge about a user's communication habits (i.e. with whom, when, how long, how often, how much a user communicates) by observing message exchanges. With Mobile IP the traffic analysis threat aggravates as tracing of users' locations and movements (i.e. when, how long, how often is a user attached to a specific network) which must be defeated as well.

Most of the problems mentioned above are by no means new. During the past years the Internet community has thoroughly studied the security problems and several extensions for the Internet Protocol have already been proposed (e.g. (Atkinson, 1995a)). In (Atkinson, 1995b) a Security Architecture for Internet Protocols has been elaborated which aims to guarantee a certain level of security in Internet communication. The proposed architecture focuses on integrity (i.e. no malicious altering of messages during transfer), authentication and confidentiality (i.e. no information leakage to an unauthorized third party). A lot of the security techniques elaborated for conventional IP can be transfered successfully to Mobile IP. Unfortunately, no solutions are offered to prevent an attacker from tracing a user's movements. Though being aware of the risks of anonymity (Anonymous, 1996) we are convinced that some form of anonymity is clearly desirable and necessary. Hence, we have developed a simple extension of the Mobile IP protocol providing location privacy to mobile hosts: the *Mixed Mobile IP*.

2.3 Related Work

As of this writing, anonymity in mobile networks is a comparably young area of research. However, a substantial amount of research has been done on anonymity in general. The mechanisms employed by the Mixed Mobile IP

owe in general to the principles developed by David Chaum (Chaum,1988) and in particular to his idea of *mixes*. Chaumian mixes constitute intermediate entities placed between the sender and the recipient of a message to conceal the link between them by hiding the relationship between the messages entering and leaving a mix. To prevent an attacker from correlating incoming and outgoing messages by content or length, the encoding of a message is changed and its length adjusted. Any Message M to be forwarded by a mix must have been randomly padded and encrypted (Diffie, 19976) with the public key K_M of the mix by the sender, i.e. consists of $\{M + padding\}^{K_M}$ for any message M*. In order to change the encoding and length of the incoming messages, the mix applies decryption and discards the padding.

Chaumian mixes are typically implemented by anonymous remailers for electronic mail (Gulcu, 1995).

Pfitzmann demonstrated in (Pfitzmann, 1991) that mixes may be modified to make them an appropriate means of defeating traffic analysis in ISDN networks . The mix idea was transfered into a mobility context with the *Non-Disclosure Method* in (Fasbender, 1992). This method aims at achieving the same goal as the Mixed Mobile IP, namely ensuring location privacy for the Mobile IP. However, their solution differs from the Mixed Mobile IP in some key respects since it trades off performance against security.

3 THE MIXED MOBILE INTERNET PROTOCOL

Maintaining location information about mobile hosts in a mobile network—which is an integral part of mobile networking—does not necessarily imply a violation of the users' right of privacy. As long as the location information available cannot be associated with a particular user, privacy is still upheld.

For the Mixed Mobile IP it is assumed that the maintainer of the mobile network, who has control over the foreign agents, is not necessarily trustworthy. Likewise, the network operator providing the stationary network and thus the home agent need not be trusted, either. In addition, it is taken into account that a powerful attacker who is able to perform eavesdropping at arbitrary locations in the network may trace messages while they traverse the network and thus link the sender of a message to its recipient. The Mixed Mobile Internet Protocol amends the Mobile IP in a way that any of the potential opponents named may discover at maximum either the mobile node's identity or its location, but not both. Moreover, conspiracy of two opponents with complementary knowledge is prevented.

A more detailed presentation of the Mixed Mobile IP is given in (Lopatic, 1996).

---

*$\{M\}^K$ denotes encryption of M with a key K.

3.1 Protocol Architecture

In order to make Mixed Mobile IP secure and comfortable to use, the design of the protocol extension adheres to the following principles. *Controllability*: the extension of the Mobile IP must not rely on entities that lie beyond the user's control and which are not trustworthy, e.g. the mobility agents. *Acceptability*: the extension should not impose any restrictions or modifications on the Mobile IP functionality visible to the user. *Compatibility:* the Mixed Mobile IP must interoperate with entities that implement "plain" Mobile IP. *Applicability*: the restricted resources of a mobile host (limited network bandwidth, poor computing power, ...) have to be considered. *Security*: the basic principles of designing secure systems (e.g. need-to-know, open design), transferred into the context of mobile IP networks, must be adhered to.

The Mixed Mobile Internet Protocol introduces new entities termed *mix agents*. Mix agents constitute intermediate agents implementing mix functionality between the foreign agent and the home agent. They are responsible for hiding the foreign agent from the home agent and vice versa, and, from an eavesdropper, the link between them. Figure 1 shows a scenario including two mix agents.

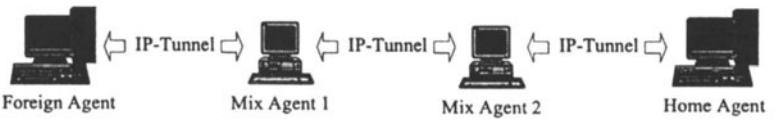

Figure 1 The basic Mixed Mobile IP scenario

The basic idea is to conceal the identity of the mobile node on the left side of the mix agents (to hide it from the foreign agent) whilst covering the node's location on the right side (to hide it from the home agent). Any eavesdropper is limited to obtaining at maximum the same information as either the foreign agent or the home agent. Since it is infeasible to trace a message through the chain of mix agents, he or she is restricted to tapping the flow of messages at only one end of the chain. Consequently, information about the location of a mobile node is completely separated from its identity.

3.2 Grouping Mixes

Obviously, proper design of the mix agents is pivotal to the security of the Mixed Mobile IP. To suit the needs of the mix agents, the original functionality implemented by Chaumian mixes has to be amended for three main reasons.

Resources The mobile node cannot be expected to have sufficient computational power to perform the iterated encryptions and padding.

Replay Protection Defending against an active opponent's attacks requires the foreign agent and the home agent to implement replay protection for messages exchanged with M_1 and M_n, respectively. This is orthogonal to the need for compatibility.

Traffic Correlation The original mix design does not sufficiently cope with attacks that involve correlating the number of messages transmitted to a mix by a certain sender to the number of outgoing messages addressed to a single destination. As in IP networks typically more than one message is sent per time interval to a specific receiver the traffic correlation problem must be solved.

The solution to the first issue is to delegate the expensive operations from the mobile node to the mix agents. Encryption and padding is only done between the mix agents. In contrast to the iterated encryption employed in the original mix model, only a single layer of encryption is applied. Figure 2 presents a scheme illustrating this mechanism for the mix agent M_1 using the same key for messages sent to an internet host (IH).

Figure 2 Tunneling in the Mixed Mobile IP

The solution for the second and third issue is to employ *grouped mixing*, which is performed by *grouping mixes*.

(a) Functionality of Grouping Mixes

The basic idea of grouping mixes is to form groups of potential destinations. Collected messages are forwarded to a destination only if there are messages pending for all other destinations which are members of the same group. All mix agents implement grouped mixing.

Grouping mixes maintain a set of entities $P = \{E_1, E_2, E_3, \ldots\}$, the *pool*, which initially contains the addresses of all other grouping mixes present in the network.* The constant g denotes the *group size* of an individual grouping mix and is chosen by the maintainer. If a message is to be forwarded by the

*The term *entities* refers in this context to mobility agents and mix agents. The initial contents of the pool must be set up manually by the maintainer of the mix.

mix to an entity E, $g-1$ entities X_n, $1 \leq n \leq g-1$ are randomly picked from the pool to form, together with E, a group of g different entities, i.e. a set of the form $G = \{E, X_1, \ldots, X_{g-1}\}$. If a group containing E already exists, no new group needs to be formed. After that, the mix adds E to the pool, if it is not already contained.

To defeat traffic correlation, a grouping mix ensures that for every message forwarded to a member of a group, e.g. E which is contained in G, a message is also sent to each of the remaining entities in the group, e.g. $X_1, \ldots, X_{g-1}$. If there are no messages pending for a group member, a decoy message is generated instead. In this way, increasing traffic to one member leads to increased traffic to all g group members. Therefore, after a message has passed the first grouping mix, a wire tapper correlating incoming and outgoing traffic would have to verify g potential recipients; after the second mix g^2, etc.

The processing of datagrams by mix agents can be split into eleven steps. However, it must be noted that there are other variants of the basic algorithm which provide other trade offs between average response time and network load caused.

Basic Algorithm

1. Decapsulate the datagram received.
2. If the datagram has been sent by another mix agent, decrypt it and remove the padding.
3. If the datagram is not a decoy, store it. Otherwise discard the decoy and form a datagram to a random destination picked from the pool and store the newly generated datagram instead. Go back to step 1 unless a time limit t is reached.
4. Sort the collected datagrams according to their destination address.
5. Mark all existing groups *idle*.
6. Check for each destination address whether it is a member of a group. If it is, mark the corresponding group *busy*. If it is not, form a group by combining it with $g-1$ randomly chosen addresses from the pool and mark the newly created group *busy*.
7. Make sure that for every member of every busy group a datagram will be output by generating decoy datagrams for members that do not have any datagrams pending.
8. Determine the longest datagram to be output for every group and pad all other datagrams belonging to the same group, so that all datagrams to different members of a group have the same size.
9. Where applicable, i.e. if the recipient is a mix agent, encrypt the datagrams to be output.
10. Randomly reorder the datagrams.
11. Output one datagram for each destination address using a tunnel to the recipient and discard the datagrams sent. If all datagrams have been output, go back to step 1. Otherwise return to step 5.

(b) Performance issues

Grouped mixes are on the one hand related to Chaumian mixes and on the other hand to recipient anonymity through *broadcasting*, i.e. sending a message not only to the single intended recipient but to all potential recipients. Considering the latter, it could be argued that grouped mixing is extremely expensive in terms of network bandwidth. Actually this is particularly true for the worst case in which messages are available for only one member of a group, so after the timeout t has passed, $g - 1$ decoys have to be generated, multiplying the amount of messages forwarded by the factor g at each mix agent. However, this scenario is extremely unlikely and practical experience shows that grouped mixing performs reasonably well. In addition, more intelligent grouping mixes could implement *dynamic grouping*, i.e. periodically reorganize the existing groups and form more balanced groups by combining entities that receive roughly the same amount of network traffic, based on previously collected statistical information.

Moreover the very limited bandwidth of wireless links has to be taken into account. Even if a large number of datagrams sent and received by mobile nodes are multiplied, the network bandwidth wasted is still relatively low.

3.3 The Protocol

Before registering, the mobile node selects a chain of two or more mix agents $M_1, \ldots, M_n$, $n \geq 2$ to be employed. Two mix agents M_i and M_j may be identical, as long as there are altogether at least two different mix agents in the chain, i.e. $M_i \neq M_j$ for at least one pair i, j. The index denotes the *position* of a mix agent in the selected chain. The mix agents are required to support a public key cipher with the respective public keys $K_1, \ldots, K_n$ which are complemented by the private keys $K_1^{-1}, \ldots, K_n^{-1}$. It is assumed that the mobile node is aware of the public keys corresponding to its selection of mix agents.

The mix agents are linked to foreign agents and home agents by bidirectional tunnels. The connections between mix agents consist of encrypted tunnels. Accordingly, it must be assumed that each mix agent M_i knows the public key of its successor M_{i+1} and vice versa.

(a) Building the registration request

In addition to the n mix agents, the mobile node chooses a pseudonym P_0 to use as its IP address in the foreign network. In case of collision with another mobile node using the same address—which can be detected during registration—another random P_0 is generated. Then the mobile node forms a registration request (P_0, M_1, FA), pretending to have a home address of P_0 and M_1 as its home agent. The mobile node appends an *anonymity extension* A to its registration request, which results in the tuple (P_0, M_1, FA, A). The Mobile IP defines a standard way of adding extensions, e.g. for authentica-

tion purposes, to registration messages. In this paper extensions are represented by appending a component to the tuple representing the message as in (MN, HA, FA, Extension). The extension adds the following data to the request.

1. $\{M_{i+1}\}^{K_i}$, $1 \leq i < n$, i.e. the address of the mix agents $M_2, \ldots, M_n$, each encrypted with the public key of the preceding mix agent in the chain.
2. $\{\mathrm{HA}\}^{K_n}$, i.e. the address of the home agent encrypted with the public key of the last mix agent.
3. $\{\mathrm{MN}\}^{K_n}$, i.e. the home address of the mobile node encrypted with the public key of the last mix agent.

Accordingly, an anonymity extension A may be represented as a (n+1)-tuple, n denoting the number of mix agents involved, as in

$$\mathrm{A} = (\{M_2\}^{K_1}, \ldots, \{M_n\}^{K_{n-1}}, \{\mathrm{HA}\}^{K_n}, \{\mathrm{MN}\}^{K_n})$$

(b) Routing the registration request

The mobile node sends the request (P_0, M_1, FA, A) it has formed to its foreign agent FA. The foreign agent, believing that M_1 is the node's home agent, forwards the message to M_1. Each of the mix agents $M_1, \ldots M_{n-1}$ processes the request in the same manner. Let M_i, $1 \leq i < n$ be one of these mix agents.

1. M_i extracts $\{M_{i+1}\}^{K_i}$ from the anonymity extension and decrypts it with its private key K_i^{-1}. It therefore obtains M_{i+1}, which is the address of the next mix in the chain.
2. M_i then rotates the components of the anonymity extension by one position to the left. For instance, rotating $(\{M_2\}^{K_1}, \ldots, \{M_n\}^{K_{n-1}}, \{\mathrm{HA}\}^{K_n}, \{\mathrm{MN}\}^{K_n})$ results in $(\{M_3\}^{K_2}, \ldots, \{M_n\}^{K_{n-1}}, \{\mathrm{HA}\}^{K_n}, \{\mathrm{MN}\}^{K_n}, \{M_2\}^{K_1})$, which is what M_1 does. In this way, each mix agent just inspects the first component of the anonymity extension. It does not need to be aware of its position in the chain.
3. M_i copies its address from the *home agent* field of the request to the *care-of address* field and places M_{i+1} in the *home agent* field.
4. M_i replaces P_{i-1} by a pseudonym P_i that it generates itself.
5. M_i creates a *connection* to the neighboring mix agents by storing the tuple $(P_{i-1}, M_{i-1}, P_i, M_{i+1})$. A connection is uniquely identified by the association of P_{i-1} with M_{i-1} and P_i with M_{i+1}.
6. The resulting registration request (P_i, M_{i+1}, M_i, A) is forwarded to M_{i+1}.

The last mix M_n also inserts its address M_n as the new care-of address. However, it obtains HA and MN from the anonymity extension and sets the

home agent field of the request to HA and the *home address* field to MN. This results in a registration request (MN, HA, M_n, A), which is forwarded to the home agent of the mobile node. Consequently, the home agent creates a mobility binding for the node under the assumption that the node's care-of address is M_n and returns a registration reply (MN, HA) to M_n.

As can be easily seen, only the last mix agent M_n and the home agent are aware of the mobile node's home address MN. However, they do not have any information about the location of the mobile node. Contrarily, only the first mix agent M_1 and the foreign agent are informed of the node's care-of address. Yet, they are unaware of the home address nor do they know the home agent.

(c) Routing the registration reply and network traffic

By matching a received registration reply to the connections stored, a mix agent is able to find the appropriate mix agent in the chain, i.e. its predecessor, to forward the message to as well as the corresponding pseudonym to use. In addition, the registration reply is modified in a similar manner as the request.

Analogously, network traffic is routed after the registration has succeeded.

4 IMPLEMENTATION

Implementing the Mixed Mobile IP on top of an existing Mobile IP protocol stack consists of two general tasks. On the one hand the Mobile IP support provided by the mobile node must be extended to comply with the Mixed Mobile IP registration procedure, on the other hand, mix agents have to be created.

4.1 General Issues

Besides providing mix functionality, a mix agent acts as a packet filter. Datagrams have to be filtered to restrict the use of protocols and protocol options which might be abused to reveal the identity of the mobile node, e.g. the ICMP or the source routing IP option. Basically only UDP, TCP and ICMP datagrams are allowed to pass through a mix agent. For UDP and TCP datagrams sent by (or addressed to) the mobile node, the source (or destination) address is rewritten by a mix agent M_i to contain the pseudonym P_i (or P_{i-1}) it (or the preceding mix in the chain) has generated. If necessary, the contents of ICMP messages are adapted as well.

To authenticate registration requests and replies exchanged between mix agents and between a mix agent and a home or foreign agent, the foreign-home authentication extension is used.

4.2 The Mixed Mobile IP for Linux

Our realization of the Mixed Mobile IP is based on version 1.2 of the Mobile IP for Linux created at the National University of Singapore, which was ported by us to the 2.0.28 Linux kernel*. The Singapore Mobile IP implements either mobile node or mobility agent functionality on a single host. To build the Mixed Mobile IP on top of the Mobile IP, we extended the existing registration routines of the mobile node implementation. Furthermore, mix agent functionality was added to the mobility agent code, enabling a home or foreign agent to act as a mix agent at the same time.

Care has been taken not to restrict existing features of the Mobile IP stack by our implementation of the Mixed Mobile IP. Apart from slightly extended configuration files, the Mixed Mobile IP is invisible to the user. On the mobile node, the mix agents to be employed are selected via the configuration file. On the mix agents, various parameters like the group size for grouped mixing may now be set in this file.

Our implementation aims at providing a stable environment for evaluation purposes. As of this writing, some features like encryption have not been implemented yet. However, the implementation is well-suited for experimenting with different algorithms for grouped mixing, observing the effects of grouped mixing on the network load or on response times, etc. It has reached a point at which it is at least as stable as the underlying Mobile IP code.

The Mixed Mobile IP was used in an environment consisting of a mobile node, a foreign agent which also acted as a mix agent, another dedicated mix agent and a home agent running the unmodified Mobile IP. To obtain mix chains of more than two mix agents, an alternating sequence of the first and second mix agent was chosen. This environment performed well enough to permit remote logins from the mobile node to other Internet hosts with response times short enough to allow for working interactively.

5 CONCLUSION

We have presented a method to provide location privacy to users in Mobile IP networks. The solution described, namely the Mixed Mobile IP, is based on the idea of mixes originally developed by David Chaum. However, to fit into the given mobility context, the functionality of mixes had to be amended, resulting in the development of grouped mixing. Grouped mixing is employed by the mix agents introduced by the Mixed Mobile IP to assist in obtaining our design goal. Yet, it must not be missed that grouping mixes are universal—they are not at all limited to the Mobile IP, not even to mobile networks in general. A variant might as well be deployed in stationary networks, e.g. in

*Available from *ftp.nus.sg*. The original implementation required an obsolete 1.3.55 kernel.

the form of proxy servers which provide anonymous access to World-Wide Web pages by separating a HTTP client from the servers it accesses.

As of this writing we are in the process of evaluating the Mixed Mobile IP in terms of response times and overhead in network bandwidth. In addition, our Linux implementation of the protocol forms the base for experiments with different algorithms for grouped mixing.

6 REFERENCES

Anonymous. (1996) Risks of Anonymity. *Communications of the ACM*, Vol. 39 No. 12, p 162

R. Atkinson. (1995a) *IP Authentication Header.* RFC 1826. RFCs may be obtained via FTP from ftp.internic.net.

R. Atkinson. (1995b) *Security Architecture for the Internet Protocol.* RFC 1825

D. Chaum. (1988) Untraceable Electronic Mail, Return Addresses, and Digital Pseudonyms. *Communications of the ACM*, Vol. 24 No. 4, p84–88

W. Diffie and M.E. Hellman. (1976) New Directions in Cryptography. *IEEE Transactions on Information Theory*, IT-22(6), p 644–654,

A. Fasbender et al. (1996) Variable and Scalable Security: Protection of Location Information in Mobile IP. In *46th IEEE Vehicular Technology Society Conference*

G. H. Forman and J. Zahorjan. (1994) The Challenges of Mobile Computing. *IEEE Computer*, 27(4), p 38–47

C. Gulcu. (1995) *The Anonymous E-mail Conversation.* Technical Report, IBM Research Division, Zurich Research Laboratory

T. Lopatic. (1996) The Mixed Mobile Internet Protocol. Master Thesis, Munich University of Technology

C. Perkins. (1996) *IP Mobility Support.* RFC 2002

A. Pfitzmann et al. (1991) ISDN-MIXes - Untraceable Communication with very small Bandwidth Overhead. In D. T. Lindsay and W. L. Price, editors, *Information Security, Proc. IFIP/Sec '91*, pages 245–258. Springer