Revocable anonymous access to the Internet

Joris Claessens[1], Claudia Díaz[1], Caroline Goemans[2],
Bart Preneel[1], Joos Vandewalle[1], Jos Dumortier[2]
[1]COmputer Security and Industrial Cryptography (COSIC)
http://www.esat.kuleuven.ac.be/cosic/
[2]Interdisciplinary Centre for Law and Information Technology (ICRI)
http://www.law.kuleuven.ac.be/icri/
Katholieke Universiteit Leuven, Belgium

March 15, 2002

Abstract

With the worldwide growth of open telecommunication networks and in particular the Internet, also the privacy and security concerns of people using these networks have increased. On the one hand, users are concerned about their privacy, and desire to anonymously access the network. On the other hand, some organizations are concerned about how this anonymous access might be abused. This paper intends to satisfy these seemingly conflicting interests, and presents a solution for revocable anonymous access to the Internet. The paper also presents some legal background and motivation for such a solution.

1 Introduction

Open telecommunication networks and in particular the Internet have been growing substantially during the last years. At the same time, a wide variety of electronic services through these networks has emerged, and the amount of people who are using these services on a frequent basis, is increasing rapidly. On the one hand, and not unexpectedly, people are more and more concerned about their privacy. They are worried about all their (trans)actions being linkable to each other and even worse, to their identity. On the other hand, unfortunately not surprisingly either, some organizations, governments, and also many users are concerned about how anonymous access can be abused by criminals. This paper intends to satisfy these seemingly conflicting interests. The paper presents a solution for revocable anonymous access to the Internet. I.e., this solution provides anonymous access to every user; however when appropriate, and only with the help of a trusted party, the anonymity can be revoked to reveal the real identity of a particular user. The paper also presents some legal background and motivation for such a solution.

1.1 Scope and outline of the paper

The main purpose of this work is to present a technical solution for revocable anonymous access. We do not want to take a pro or contra standpoint with respect to revocability and/or anonymity. A continuous discussion of these issues is however very important in our on-line society. Clearly, the subject of anonymity and revocation includes much more issues than solely technicalities.

The paper starts with a section on the (technical) motivation for a solution for revocable anonymous access to the Internet. Motivation from a legal point of view is discussed in Sect. 3. Section 4 then presents an overview on anonymous communication. Section 5 introduces the concept of revocation, discusses its difficulties, and sets the requirements for a system that provides revocable anonymous access. A solution for revocable anonymous access is proposed in Sect. 6. Finally, concluding remarks and future directions are given.

2 Motivation

This section studies the motivations for a solution for revocable anonymous access to the Internet. We first start with explaining the need for anonymous communication. Then, we give some arguments in favor of revocation.

2.1 Real anonymity: data *and* connection anonymity

For the World Wide Web, many products have been developed to protect users' on-line privacy (and provide extra security): personal firewalls, password managers, form fillers, cookie managers, banner managers, keyword alerts, etc. The latest browsers also include part of this functionality. The focus here is on *data anonymity*, that is, ensuring that the content does not contain any identifiable information. This includes stripping out information (e.g., cookies and other HTTP headers), choosing random usernames, etc.

Electronic payment systems are a particular area of interest with respect to anonymity. Although today's commonly used electronic payment systems are not anonymous, several anonymous payment schemes have been developed (e.g., ecash [18]). Basically, the goal of any anonymous payment scheme is to prevent the bank from being able to link its users to the payments they made, based on the information (at data level!) obtained during the withdrawals and deposits of electronic money.

Any telecommunication network requires users to have a network address during communication. For the Internet, this is an IP address. By definition, the network address is revealed to the communicating party. In practical networks, the network address can be linked to a group of users, sometimes even to one particular user. The network address itself thus constitutes identifiable information. A solution that hides the user's network address provides *connection anonymity*. Such a solution is not obvious to achieve, and is certainly less trivial than a solution for data anonymity.

Real anonymous systems require both data and connection anonymity. For anonymous payments, this has already been indicated in the past by several researchers, e.g., by Simon [20]. If there is no connection anonymity, so-called anonymous coins can be traced back to their originators just by looking at the network addresses.

Moreover, when the whole context of a payment (e.g., ordering of goods by a particular person for a specific amount of money) is identifiable, anonymous coins can be mapped to this specific non-anonymous context. In order to achieve a good level of privacy on the Internet, all IP traffic should be anonymized, and not only HTTP. Web pages can contain non-HTTP links. Even DNS requests should be anonymized, as these requests reveal the intended recipient, even though the actual connection from the originator to the recipient is not traceable. Anonymity should thus be present in all parts of a system. This is very difficult to achieve as network anonymization mechanisms could be circumvented in many indirect ways (e.g., Felten and Schneider [10] demonstrated this by checking if certain web pages are in the browser's cache; Martin and Schulman [15] recently described how anonymity can be totally undermined with JavaScript).

2.2 *Revocable* anonymous communication

In contrast to data anonymity, connection anonymity does not gain interest on a large scale. For example, Zero-Knowledge Systems recently discontinued their Freedom service [24], nevertheless a state-of-the-art solution for connection anonymity. Zero-Knowledge Systems claims this is due to lack of interest, and not to government pressure.

Either way, unconditional anonymity can be misused. In the area of electronic payment systems, a lot of research has already been done on revocable anonymous payment schemes. An overview has been presented in earlier work [7]. The ability to revoke anonymity is required by financial organizations and governments in order to prevent fraud and/or to be able to trace back suspicious payments/withdrawals. Similarly, revocable anonymous communication seems to ensure the balance between the users' right to privacy and the various concerns of governments and organizations.

Generic (and existing) solutions for anonymous communication are not by default included in current operating systems or provided by Internet Service Providers (ISPs). If revocation would be enabled, perhaps solutions for anonymous communication would be offered on a larger scale.

3 From the legal point of view

As stated above, revocable anonymous communication seems to ensure the balance between the users' right to privacy and legitimate concerns of public authorities, organizations and third parties. In this section, we show that from a legal point of view, a solution for revocable anonymous access is very meaningful.

The need for on-line anonymity has been expressly recognized both at U.S. and at European level. U.S. court actions challenging the use of on-line anonymity are often cases of defamation in reply to anonymous on-line postings. U.S. courts have expressly accepted the use of on-line anonymity as a means to exercise the basic right of freedom of speech without fear of retaliation.

At EU level, the use of on-line anonymity is more focused on privacy issues. The EU Data Protection Working Party, official advisor for the EU Commission and the Member States on the EU Directive with regard to processing of personal data, issued a Recommendation in 1997 regarding anonymity on the Internet. The Working Party concludes that the ability to choose to remain anonymously – and consequently to have anonymous access to the Internet – is essential if individuals are to preserve the same protection for their privacy on-line as they currently enjoy off-line. In the same line, the Council of Europe in Strasbourg promotes the use of on-line anonymity as a tool for effective protection of the fundamental right to on-line privacy.

However, those European Institutions agree that anonymity is not appropriate in all circumstances. Restrictions on the use of anonymity can be justified at two levels. Firstly, a legal entity (physical or legal personality) is in principle accountable for his acts. Even if, in principle, according to civil law, anonymously concluded contracts can be legally valid upon condition that no identification formalities are required, shortcomings may occur during the implementation of the contract. This necessitates a breakthrough of the contractual parties' anonymity.

In specific circumstances, the European legislator has expressly imposed identity disclosure, e.g., on service providers offering services of the Information Society, as provided for in the EU Directive of 8 June 2000 on electronic commerce. In other circumstances, the European legislator expressly bans the use of anonymity, e.g., on the use of electronic mail for purposes of direct marketing (spamming), as provided for in the Draft Directive concerning the processing of personal data in the electronic communication sector. Finally, there are circumstances in which the European legislator creates room for "grey areas" of anonymity. This is for example the case for certification service providers who cannot be prevented from indicating in the certificate a pseudonym instead of the signatory's name, according to the EU Directive of 13 December 1999 on electronic signatures.

Secondly, restrictions on the use of anonymity may also be imposed with a view to protect public interest in a democratic society. The fundamental rights to privacy and freedom of expression, guaranteed by article 6 and 10 of the European Convention on Human Rights, are indeed not absolute rights and have to be put in balance with other interests. Interference by public authority may exceptionally take place if explicitly provided by law and necessary in a democratic society. Therefore, the exceptions on the right to privacy and the freedom of expression (and thus on the use of anonymity) touch on public order and on subjective rights of individuals. Public order refers to matters as prevention of national security, territorial integrity of public safety, prevention of disorder or crime whereas subjective rights of individuals will in this context cover the

protection of the reputation or rights of others, preventing the disclosure of information received in confidence.

Moreover, according to the case law of the European Court on Human Rights in Strasbourg, interference by public authority – whether preventive or repressive – should in any case undergo a *proportionality* test. This means that restrictions imposed by public authority should not go beyond the objective to be reached. In this respect, the European Convention on Cybercrime of the Council of Europe, signed on 21 November 2001, adopted a balanced approach regarding retention of traffic data imposed on Internet service providers for purposes of prosecution of criminal offences. The Convention rejects the principle of a *mandatory* retention of traffic data and thus takes the view that traffic data should not be kept by Internet service providers only for law enforcement purposes. Orders to retain traffic data are limited to *specific*, case by case criminal investigations or proceedings.

Unfortunately some countries, such as the UK and Belgium, have introduced the principle of mandatory retention of traffic data in their domestic criminal law, based on the exception of public order. Some regulations go even further and envisage possibilities to prohibit in future the provision of on-line services that hinder the identification of the user. It is not unlikely that those laws will be challenged in court procedures in respect with the proportionality test: can a whole population be put under control for the sake of the easiness of criminal investigations? Retention of traffic data and subsequent revocation of anonymity on a case by case basis would no doubt constitute a more proportional solution.

In the meantime, technical solutions must be put in place both to secure retained data from unauthorised access and to facilitate its disclosure whenever revocation of anonymity is being ordered in criminal investigations. The existence of strong technical guarantees for expeditive revocation of anonymity, whenever ordered, may refrain the legislator to introduce further "anonymity unfriendly" regulations. Today, the challenge for lawyers is to ensure a balanced and harmonised legal framework for conditions and safeguards whenever revocation of anonymity is required.

For further social and legal aspects on anonymity, we refer to van Dellen's "Anonymity Law Survey" [22].

4 Anonymous communication

Before discussing revocation issues, we first explain the concept of anonymous communication, and give an overview of the existing solutions.

4.1 Definition of anonymous communication

In the scope of this paper, we consider real-time, bidirectional (IP-based) communication between an initiator and a responder. During this communication, only the initiator should know with whom (i.e., the responder's IP address)

he is communicating. Other entities in the network should not know a particular initiator is communicating with a particular responder. More precisely, anonymous communication in this paper means communication with *initiator anonymity*, hence 'anonymous access'.

It is important to understand towards whom the initiator is anonymous. We can actually distinguish different 'attack models'. If the adversary is local (e.g, the responder itself looking at the incoming connections), then an intermediate proxy that relays the communication already ensures anonymity (e.g., this is one of the core mechanisms of the Anonymizer [2]). If one considers that the adversary is able to observe the global network, this proxy will not be sufficient. In between these two extremes is the case in which the adversary consists of a number of collaborating local observers. Solutions for anonymous communication mostly rely on the assumption that there exist a number of entities that can be trusted not to collaborate. Ideally, solutions for anonymous communication should not require more trust than that (as a counter example, the intermediate proxy knows the correspondence between initiator and responder and is trusted not to disclose nor log it).

The solution discussed in this paper provides an application-independent building block for anonymous connections. For real initiator anonymity towards the responder, it should be used in combination with a solution for data anonymity. However, the solution can also be deployed in systems where the initiator wants nobody to know who he is communicating with, except for the responder itself. In that case, the initiator can explicitly identify himself towards the responder within the content exchanged over the anonymous channel (which in this case should provide confidentiality too).

Note that a number of solutions for anonymous communication are not useful in our context. For example, TAZ servers and Rewebbers [12] are a solution for anonymous publishing in which the initiator does not know who the responder is (which is the reverse problem of the one we are looking at here).

4.2 Existing solutions for anonymous communication

Two main approaches exist that achieve anonymous communication: Chaum's dc-net [6] and Chaum's mix-net [5]. The first seems more of theoretical interest, the latter forms the basis for almost all practical solutions.

A mix is a network entity that achieves anonymous communication by hiding the correspondence between the messages it receives on its input and messages it forwards on its output. The mix hides the order of arrival of the messages by reordering, delaying and padding traffic. As we consider real-time and bidirectional communication, for example delaying is not really possible. Practical solutions therefore require a chain of mixes in order to provide an adequate level of anonymity. There are two categories of solutions: on the one hand Crowds [17] and Hordes [19], on the other hand Onion Routing [16], PipeNet [8], Freedom [23] and Web MIXes [4]. These solutions represent two different approaches. Both approaches achieve initiator anonymity by setting up a path from initiator to responder through several intermediate entities. Tracing the

path from a particular initiator to a particular responder is made very difficult by hiding the correspondence between the different connections in between the intermediate entities. Encryption of the data that is exchanged in between each two entities is required in order to achieve this. A detailed description can be found in [17] and [16]. We briefly outline the main properties. In addition to these, Hordes has the interesting feature of using multicast for the reply, while Web MIXes deploys a ticket-based authentication system to prevent flooding attacks.

'Crowds'-like anonymous communication. Crowds is intended for HTTP traffic. The approach should be applicable to all IP traffic though. The users themselves are the intermediate entities. The path from initiator to responder is established by the users through random forwarding of the (web) request. This anonymous connection is therefore of random length.

Every user on the path has access to the content of the request. Every user therefore knows the identity of the responder. However, no user knows the identity of the initiator, not even the second on the path (as users do not know their position on the path). Collaborating users can easily detect that they are on the same path. The non-collaborating user immediately preceding the first collaborator, will be the initiator with a certain probability (depending on the number of users, the number of collaborating users, and the fixed probability of forwarding a request). The higher the number of collaborating users, the higher this probability. Thus, Crowds provides an adequate level of anonymity up to a certain maximum number of collaborating users.

'Onion Routing'-like anonymous communication. In Onion Routing the intermediate entities are routers. The path from initiator to responder is established as follows. The initiator prepares a layered request (called onion) that contains information for each router. This info consists of cryptographic key material, the identity of the next hop, and an encrypted onion for the next hop. Once the path has been established, data is encrypted multiple times, and sent through the path of routers. Each router decrypts one layer, and forwards the data to the following one. The anonymous connection is of chosen length and goes through routers chosen by the user.

Onion Routers only know the previous and next hop. None of the routers see the same information. Only the first router knows the initiator, and only the last knows the responder; but even if they are collaborating, they cannot link the two together. Compared to Crowds, Onion Routing provides anonymity against a stronger adversary who is able to observe the network in global.

5 Revocation

Strong anonymous communication schemes make it practically impossible to identify the initiator later on. We say that anonymity is unconditional. As discussed in Sections 2 and 3, revocability may be a required feature in a system for

anonymous communication. In this section, we explain what revocation is, we discuss the difficulties there are to implement revocation for anonymous communication, and we raise some issues regarding trust and revocation. Finally, we list the requirements for a solution for revocable anonymous access to the Internet.

5.1 Concept of revocation

In an unconditional anonymous system, it is impossible under any circumstances to find out the identity behind a particular transaction. In contrast, a revocable anonymous system provides a backdoor with which an identity can be traced back. Revocation should be provided according to some rules: revocation should only be technically possible when a judge or other dedicated trusted party cooperates; this trustee should not be involved in the anonymity service itself; upon revocation, only the identity of the particular targets should be revealed, while all other transactions and/or users remain anonymous.

In our case of (IP based) anonymous access to the Internet, revocation could make sense in a number of situations; for example, tracing of a user who uploaded or downloaded illegal content on a particular web server, tracing of a hacker who broke into a particular host, tracing of users who communicated with a suspicious party, etc. A revocable anonymity service can be intended for users throughout the whole Internet, or it might also be developed for users within a certain organization or ISP.

Revocation is about the ability to trace back. Some equivalences can be drawn between anonymous payment systems and anonymous communication. A revocable anonymous payment system provides *coin tracing*, i.e., it is possible to trace back the electronic coins a particular user has withdrawn. For anonymous communication, the equivalent is *responder tracing*, i.e., it is possible to trace back the identity of the responder a particular initiator has communicated with. A revocable anonymous payment system also provides *owner tracing*, i.e., it is possible to trace back the identity of the user that has spent a particular electronic coin. In a system for anonymous communication, this is equivalent to *initiator tracing*, i.e., it is possible to trace back the initiator that has communicated with a certain responder.

Note that the system discussed in this paper only provides initiator anonymity. It makes sense to define both initiator and responder tracing, as nobody but the initiator knows the identity of the responder it is communicating with. However, as this paper addresses revocable anonymous access, responder tracing is less relevant; in this scenario, the identity of the suspected initiator is known, and thus we cannot speak about anonymous access anymore.

5.2 Difficulties

Upon comparison, revocation of anonymous communication seems to be more difficult than revocation of anonymous cash.

With anonymous cash, it is very clear which data should be stored – i.e., the bank has to store its view of the withdrawal protocols, and the electronic coins it receives. For anonymous communication, it is not immediately clear which data is relevant for revocation: data packets will probably not contain any useful information (e.g., for onion routing, only the onions contain routing information). For anonymous communication, there will typically be many more entities involved than with anonymous cash (and/or they are more distributed). For revocation of anonymous cash, the bank – that is, the responder – is involved: interaction with the bank is required (otherwise, one would not have legitimate cash), and therefore revocation is automatically supported. For revocation of communication, neither initiator nor responder are involved (nor interested!) in revocation. It seems very difficult, if not impossible, to force an initiator not to use other unconditional anonymous channels instead of the revocable one.

Does this mean that a system for revocable anonymous communication needs to be 'hardwired' into computer systems? No, it does not seem to be practically possible to control an initiator's or responder's computing platform (e.g., storing serial numbers and corresponding cryptographic keys into tamper-resistant network interface cards). However, it might be more realistic to assume that the infrastructure in between initiator and responder (e.g., routers) can be controlled. Note that it is always possible to use a computer at for example a cybercafe, to anonymously access the Internet. If revocation of the anonymous communication would just lead to the IP address of that computer, it might give no clue about the identity of the initiator. Thus, a system for revocable anonymous access to the Internet should be carefully designed in order to prevent the situation in which only behaving people will be revocable, while people with malicious intentions will be able to circumvent the system anyway.

5.3 Trust and revocable anonymous communication

Just as in any other security system, trust is a crucial issue in a system for revocable anonymous communication. Who should be trusted, to what extent, and for what purpose, are questions we have to solve.

Depending on the actual system for anonymous communication, there might be one or more entities that know the relationship between an initiator and a responder during communication (e.g., if there is only one intermediate mix). These entities are trusted not to reveal or log this relationship. Obviously, the party dedicated to revocation should be trusted not to misuse its powers. This party should not have access to data records of which the anonymity must not be revoked.

In both Crowds-like and Onion Routing-like systems, applications access the Internet anonymously through a proxy. This proxy knows the initiator's as well as the responder's identity. For individual use, this proxy is under the control of the initiator. However, as indicated by Reed *et al.* [16], for organizational use, there might only be one proxy at the firewall. "This protects the anonymity of connections from observers outside the firewall, but also simplifies enforcement

of and monitoring for compliance with corporate usage policy." Both alternatives should be possible for a revocable system.

In an open environment (e.g., the Internet), it does not seem to be realistic to have just one party that can revoke anonymity on its own. It is better to distribute the capability of revocation among several parties, preferably with a threshold scheme (i.e., a certain minimum number of these parties need to cooperate to be able to revoke). Instead of having separate trusted parties for the revocation purpose only, it might be interesting, in the limit, to distribute the capability of revocation among the anonymous service providers themselves. Upon request of the police, a judge, or the government, the community of providers can then decide to cooperate. Note however that this is in contrast to the requirement that the trustee should not be in involved in the actual anonymity service (see next section).

5.4 Requirements

Based on the observations and discussions made up to this point in the paper, we list here the most important requirements for a system for revocable anonymous access to the Internet.

- The system should provide anonymous access to the Internet at the IP layer. This means that any application is anonymized, and, for example, not only HTTP.
- Revocation should be provided according to the rules: the cooperation of a dedicated trusted party should be required for revocation; this party should not be involved in the actual anonymity service; the trusted party should only be able to revoke the anonymity of the suspected communication and no other.
- Revocation should lead to the identity of the initiator. This identity should be bound to a user, and should not just be the IP address of the originating host. It should not be possible to take over some user's identity or hijack an existing anonymous connection. Revocation should not lead to the identity of a behaving initiator.
- The system should be designed in such a way that only the initiator needs to know about the infrastructure and is required to install the necessary (software) interface to use it. The responder may of course be aware that it receives (some of its) communication through this infrastructure, but does not need to install any special software or hardware.
- The anonymity service providers should not be required to log connection info, they must only securely maintain their cryptographic secrets (needed for the revocable anonymity service). Information needed for revocation should be stored at a central place. This centralized storage should not affect the strong anonymity properties of the service itself.

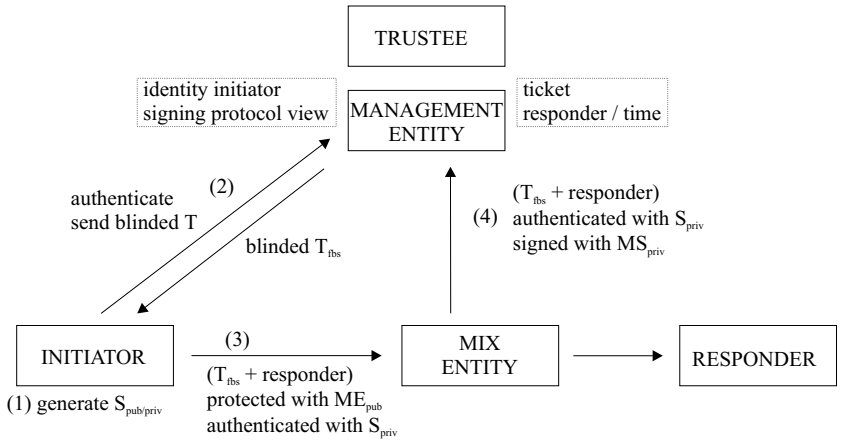

$S_{pub/priv}$	initiator's session key pair
T	ticket = S_{pub} + (optionally) responder's address
T_{fbs}	ticket with fair blind signature by management entity
$ME_{pub/priv}$	communication key pair of the mix
$MS_{pub/priv}$	signing key pair of the mix

Figure 1: Idealized solution with single mix and revocation entities

6 Solution for revocable anonymous access

We first present a basic solution only comprising a single mix entity, a single management entity, and a single trustee. We will then enhance this solution by replacing the single mix entity with two multiple mix systems, a first one based on Onion Routing, and a second one based on Crowds. The model of the basic solution completely remains in all solutions. Management entity and trustee are not explicitly distributed in this paper, but this can easily be done by using standard threshold schemes.

6.1 An 'idealized' solution

The basic solution is visualized in Fig. 1.

Entities. In short, the following parties play a role in the proposed scheme. The *initiator* wants to anonymously access the Internet, and communicate with the *responder*. The *mix entity* provides the anonymity service. The *management entity* issues tickets with which initiators will be able to anonymously access the Internet, but which will also allow anonymity revocation. The management entity is not involved in the anonymity service itself. Revocation is however only possible when the *trustee* cooperates. The trustee is not involved in any other task.

Fair blind signature. A fair blind signature, as defined by Stadler *et al.* [21], is a core mechanism in the solution. A fair blind signature is a blind signature that allows revocation, i.e., the signer (in our case: the management entity) cannot link his view of the signing protocol with the message-signature pair that has been obtained (basically, the signer does not see the message nor the resulting signature). However, after receiving some extra information from a trustee, this link can be made. The trustee should not be involved during the signing process itself. A recent construction of a fair blind signature has been proposed by Abe and Ohkubo [1].

Anonymous access. Setting up a revocable anonymous connection involves the following steps.

1. In order to establish an anonymous connection to a particular responder, the initiator must start by obtaining an approved 'ticket' for that anonymous connection. The initiator generates a session private/public key pair: S_{pub} and S_{priv}. Let the ticket T be S_{pub}. Optionally, the initiator encodes the responder's address as (for example) $IP : port$ and concatenates this with S_{pub}. Let the concatenation then be T. If T includes the responder's address, the session private/public key pair is bound to a particular responder. In the other case, T can be used to anonymously communicate to any responder.

2. The initiator blinds T and establishes a secure connection with the management entity. Within this connection, the initiator authenticates to the management entity, and obtains a fair blind signature on T from the management entity. The authentication should not be based on the initiator's IP address. It can be based on a pseudonym [14], or it can be performed with a real certified identity. The management entity logs the initiator's pseudonym/identity together with its view on the signing protocol. Let T together with the fair blind signature be the ticket T_{fbs}.

3. T_{fbs} and S_{priv} can be stored and do not have to be used immediately. They can even be used from another computer. Note that a ticket does not expire (this could be an option though by having the management entity periodically update its signing key). The initiator should carefully protect its session private key. To establish an anonymous connection with the responder, the initiator encrypts T_{fbs} and the responder's address (if not included in the ticket) with the public encryption key of the mix entity ME_{pub}, and sends the encrypted ticket to the mix entity. The initiator and the mix entity also establish a secure channel through which all communication in the session with the responder will be sent. The initiator should also prove he knows S_{priv}; e.g., he can sign T_{fbs} together with a timestamp, before encrypting it with the public mix encryption key; alternatively, initiator and mix entity could perform a challenge/response protocol.

4. The mix entity decrypts the encrypted ticket with ME_{priv} and verifies the fair blind signature. It also verifies the signature of the initiator. It signs T_{fbs} and the responder's address with the mix entity's signing key MS_{priv}, and sends this to the management entity, together with the authentication proof of the initiator. The management entity logs the ticket, the responder's address, and the mix's signature. The mix entity establishes the anonymous connection with the intended responder, i.e., it forwards all packets sent by the initiator to the responder and vice versa. The mix entity does not log any information.

Revocation. Management entity and trustee must cooperate to revoke the anonymity of a particular session.

- *Initiator tracing.* The management entity looks up the T_{fbs} corresponding to the particular responder at the particular time. The management entity and the trustee then retrieve the signing protocol view and associated identity of the initiator corresponding to the ticket. Note that this requires a type II fair blind signature, as defined in [21].
- *Responder tracing.* The goal is to trace the ticket given the signing protocol view of a particular authenticated initiator. Note that this requires a type I fair blind signature, as defined in [21].

6.2 An Onion Routing-like distributed mix entity

Instead of using a single mix entity, we now introduce a distributed mix entity based on Onion Routing. This increases the level of anonymity, as a chain of mixes is used instead of one intermediate mix. This also decreases the level of trust that needs to be put into the individual mixes, i.e., no single mix will be able to link an initiator to a responder. Only the overall mix entity would be able to perform this, but this would require cooperation of the individual mixes. The solution is shown in Fig. 2.

Anonymous access. Setting up a revocable anonymous connection involves the following steps.

1. Step 1 remains the same as in the basic scheme.
2. Step 2 remains the same as in the basic scheme.
3. To establish an anonymous connection with the responder, the initiator creates an onion. The onion is the same as in the ordinary Onion Routing solution, except for the most inner layer: the data for the last Onion Router includes T_{fbs}, which is concatenated with a timestamp and signed by the initiator using S_{priv}. The initiator sends the onion to the first mix. The anonymous connection is setup by the Onion Routers in the normal way.

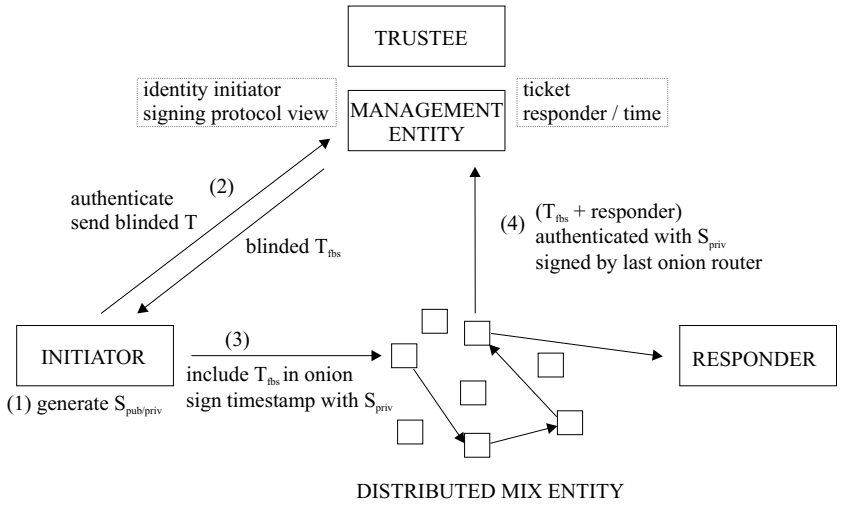

$S_{pub/priv}$	initiator's session key pair
T	ticket = S_{pub} + (optionally) responder's address
T_{fbs}	ticket with fair blind signature by management entity
$ME_{pub/priv}$	communication key pair mix ⇒ onion
$MS_{pub/priv}$	signing key pair mix ⇒ signature by last Onion Router

Figure 2: Solution with Onion Routing-like distributed mix entity

4. The last Onion Router will receive the signed ticket. It verifies the fair blind signature, the signature of the initiator, and the freshness of the ticket. It signs T_{fbs} together with the responder's address, and sends this to the management entity. The management entity logs the ticket, the responder's address, and the signature. The anonymous connection is established by sending back confirmations along the path of Onion Routers.

Revocation. Initiator and responder tracing are performed in a completely similar way as in the basic scheme.

6.3 A Crowds-like distributed mix entity

In this section, the distributed mix entity is based on the Crowds system. The solution is shown in Fig. 3.

Group signature. This solution uses an additional cryptographic tool, the group signature, see for example Ateniese *et al.* [3]. Group signatures allow a member of a group to sign a message, such that everybody can verify that the message is signed by *a* group member; which group member exactly can only be revealed after revocation. A group signature scheme requires a group manager.

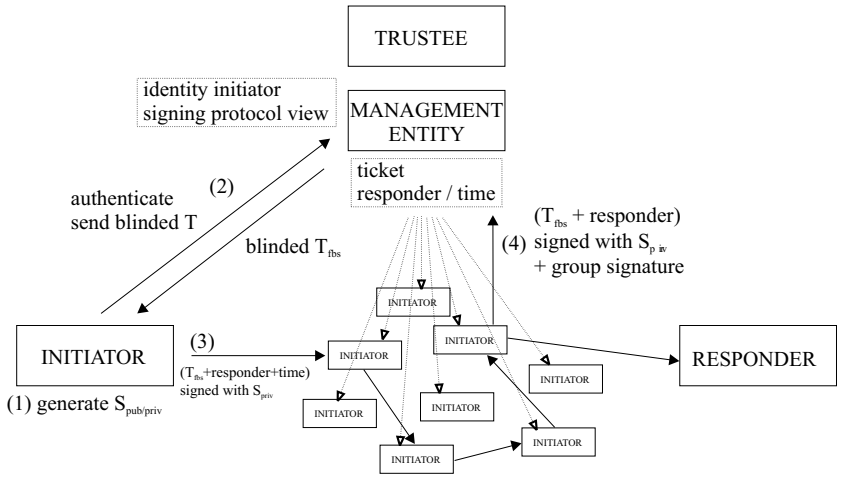

$S_{pub/priv}$	initiator's session key pair
T	ticket = S_{pub} + (optionally) responder's address
T_{fbs}	ticket with fair blind signature by management entity
$ME_{pub/priv}$	communication key pair mix ⇒ Crowds' path protection
$MS_{pub/priv}$	signing key pair mix ⇒ group signature by last Crowds member

Figure 3: Solution with Crowds-like distributed mix entity

Interaction with the group manager is needed when joining a group, and for revocation purposes. Signing and verifying is non-interactive. In our proposed solution, the management entity is the group manager, while the individual initiators form the group members.

Joining a Crowd. The initiators themselves now represent the distributed mix entity. A Crowd of initiators corresponds to a group signature group. The management entity is the group manager of a Crowd. In order to join a Crowd, an initiator must perform a join protocol for the particular group signature scheme, with the management entity.

Anonymous access. Setting up a revocable anonymous connection involves the following steps.

1. Step 1 remains the same as in the basic scheme.
2. Step 2 remains the same as in the basic scheme.
3. To establish an anonymous connection with the responder, the initiator creates a path setup request that contains the ticket, the responder's address, and a timestamp, signed with S_{priv}. This request is randomly

forwarded among the Crowd members. Each Crowd member on the path should verify the signature and the freshness of the request.

4. The Crowd member that decides to forward the request to the responder, sends the ticket and the responder's address, together with a group signature on this data, to the management entity. The management entity logs the ticket, the responder's address, and the group signature. This data is also made available to all Crowd members (e.g., a bulletin board).

5. The last Crowd member sends back a confirmation along the path. This confirmation is signed using a group signature. Each Crowd member on the path should verify the group signature, and should check whether the ticket and the responder's address is known by the management entity. If this is the case, the confirmation is forwarded, otherwise, the path setup is aborted, and the management entity is informed about a potential corrupted Crowd member.

Revocation. Initiator and responder tracing are performed in a completely similar way as in the basic scheme. In addition, a corrupted Crowd member can be traced by revoking the group signature.

6.4 Analysis

We here discuss to what extent the proposed solution meets the requirements for a system for revocable anonymous access.

The solution can be deployed for all IP based communication. However, the amount of data that would have to be logged by the management entity would be enormous. It might therefore be more realistic to only use the system for specific services in which anonymity is particularly important (e.g., surfing to web sites related to healthcare), or for services provided in particular circumstances (e.g., only in cybercafes).

The solution adopts the techniques of one intermediate mix (in the basic scheme), and Onion Routing and Crowds (in the distributed scheme). The proposed revocation enhancement should not decrease the level of anonymity of regular initiators. From an implementation point of view, responders do not need to adapt their configuration.

Revocation can only be done when the management entity cooperates with the trustee. The trustee is not involved in the anonymity service. Only the anonymity of the suspected communication can be revoked and no other. The mix is trusted to send the ticket and responder's address to the management entity. In the case of an Onion Routing-like distributed mix entity, individual Onion Routers still have to be trusted to perform this action. As an initiator can choose the last Onion Router on a path, a collaborating initiator and Onion Router are a realistic threat. Onion Routers should therefore be regularly audited in order to ensure they are doing their job. In the case of a Crowds-like distributed mix entity, we assume that the action of informing the management entity, will be verified by at least one honest Crowd member on the path.

Revocation does not lead to an IP address of a machine, but reveals a meaningful identity of the initiator, whether a pseudonym or a real certified identity.

The distributed scheme provides a better level of initiator anonymity. For Crowds, the group signature (which is equivalent to the signature of the single mix) provides an extra level of anonymity towards the management entity, and especially towards all other entities that have access to this information (i.e., all other Crowd members).

On the negative side, the Onion Routing scheme is still vulnerable to Denial of Service (DoS) attacks: invalid onions (e.g., not containing a valid ticket) will potentially only be detected by the last router on a path. Onion Routers can thus be easily flooded by invalid requests. Also the Crowds scheme is vulnerable to DoS.

The proposed Onion Routing and Crowds like solutions still includes a single and separate management entity and trustee. This should be distributed too. This can be achieved with standard threshold techniques, see for example the fair blind threshold signature scheme by Juang and Lei [13].

6.5 Related solutions

Goldberg [11] described a solution for a pseudonymous communications infrastructure for the Internet. He also made the important observation that anonymity should be built in at the communications layer, as it otherwise makes anonymity at the application layer impossible. Consequently, his solution is based on an Anonymous IP Infrastructure (AIPI) in which anonymous connections can be setup through a chain of Anonymous Internet Proxies (AIPs). Note that this infrastructure is conceptually similar to Onion Routing, and that is the same on which Freedom [23] is based. On top of this infrastructure a pseudonymity layer is added: initiators can choose a unique pseudonym and generate a corresponding public/private key pair. These can be used in the anonymous communications infrastructure: exit AIPs will be able to associate data packets with a pseudonym. The purpose of his solution is to allow initiators to have persistent pseudonyms, so that they can have linkable transactions in (but only in) case this is appropriate – e.g., when reputation based on past transactions is relevant – still without revealing their real identity. Our solution in fact also relies on pseudonyms, which in addition can be revoked to the real identity. The way the pseudonym is obtained is different: in Goldberg's system, pseudonyms are obtained over an anonymous channel, but in visible form; in our system, pseudonyms are obtained over an authenticated channel, but in blinded form.

We also briefly mention a completely other trace back issue here, namely the general problem of determining the path a packet traversed over the Internet. This problem is especially relevant in denial of service attack scenarios in which the IP packets have spoofed source addresses. Solutions are needed here to trace back the origin of the attack, i.e., this will be located on the path the IP packets have traversed (see for example Dean *et al.* [9]). These solutions are typically based on routers marking IP packets while routing them through

the network, in a clever way which will afterwards allow trace back, but only provided the attacker sent enough packets. The latter requirement is easily fulfilled in a denial of service scenario, but would not necessarily be met in a normal communications scenario. The solutions for IP trace back seem therefore not to be suited for our revocation purposes.

7 Conclusion

In today's open telecommunication networks and in particular on the Internet, users are concerned about their privacy. Anonymity is far from guaranteed by default, although it is technically relatively easy to increase the level of anonymity. Note that 'real' anonymity in fact depends on the attack model one considers. The attack model considered in this paper is strong: anonymity should be provided against adversaries that can observe the network in global. Solutions have been proposed in the past that provide unconditional anonymity against this model.

Anonymity can unfortunately be misused by criminals. Clearly, strong solutions for anonymity are therefore not wanted by governments and many organizations. This might be a reason why these solutions are not deployed on a large scale. This paper presented a solution that intends to satisfy these conflicting interests. The proposed solution provides a high level of anonymity for the ordinary user, while at the same time technically guaranteeing good identifiability of misbehaving users. This may refrain legislators from introducing further "anonymity unfriendly" regulations.

Although the proposed solution is conceptually suited for all IP-based communication, it is probably more realistic to be deployed for specific services or in specific locations only. Large ISPs might have their own logging capabilities, and might be trusted to identify their users, and release this information if required.

As already indicated, this paper also presented legal motivations in favor of a solution for revocable anonymous access to the Internet. While revocable anonymity is a quite difficult issue from a technical point of view, it causes an even more challenging legal and social debate.

Acknowledgements

Joris Claessens is funded by a research grant of the Institute for the Promotion of Innovation by Science and Technology in Flanders (IWT). This work was also supported in part by the IWT STWW project on Anonymity and Privacy in Electronic Services (APES), and by the Concerted Research Action (GOA) Mefisto-2000/06 of the Flemish Government.

References

[1] Masayuki Abe and Miyako Ohkubo. Provably Secure Fair Blind Signatures with Tight Revocation. In Colin Boyd, editor, *Advances in Cryptology – ASIACRYPT 2001*, Lecture Notes in Computer Science, LNCS 2248, pages 583–601. Springer-Verlag, December 2001.

[2] Anonymizer. http://www.anonymizer.com/.

[3] Giuseppe Ateniese, Jan Camenisch, Marc Joye, and Gene Tsudik. A Practical and Provably Secure Coalition-Resistant Group Signature Scheme. In M. Bellare, editor, *Advances in Cryptology – CRYPTO 2000*, Lecture Notes in Computer Science, LNCS 1880, pages 255–270. Springer-Verlag, 2000.

[4] Oliver Berthold, Hannes Federrath, and Stefan Köpsell. Web MIXes: A system for anonymous and unobservable Internet access. In Hannes Federrath, editor, *Designing Privacy Enhancing Technologies. Proceedings of the Workshop on Design Issues in Anonymity and Unobservability*, Lecture Notes in Computer Science, LNCS 2009, pages 115–129. Springer-Verlag, July 2000, Proceedings 2001.

[5] David L. Chaum. Untraceable Electronic Mail, Return Addresses, and Digital Pseudonyms. *Communications of the ACM*, 24(2):84–88, February 1981.

[6] David L. Chaum. The Dining Cryptographers Problem: Unconditional Sender and Recipient Untraceability. *Journal of Cryptology*, 1(1):65–75, 1988.

[7] Joris Claessens, Bart Preneel, and Joos Vandewalle. Anonymity Controlled Electronic Payment Systems. In A. Barbé, E.C. van der Meulen, and P. Vanroose, editors, *Proceedings of the 20th Symposium on Information Theory in the Benelux*, pages 109–116, May 1999.

[8] Wei Dai. PipeNet 1.1. http://www.eskimo.com/~weidai/pipenet.txt.

[9] Drew Dean, Matt Franklin, and Adam Stubblefield. An Algebraic Approach to IP Traceback. In *Proceedings of the 2001 Network and Distributed System Security Symposium*, February 2001.

[10] Edward W. Felten and Michael A. Schneider. Timing attacks on Web privacy. In *Proceedings of the 7th ACM Conference on Computer and Communications Security*, pages 25–32, November 2000.

[11] Ian Goldberg. *A Pseudonymous Communications Infrastructure for the Internet*. PhD thesis, University of California at Berkeley, 2000.

[12] Ian Goldberg and David Wagner. TAZ Servers and the Rewebber Network: Enabling Anonymous Publishing on the World Wide Web. *First Monday – Peer-Reviewed Journal on the Internet*, 3(4), April 1998.

[13] Wen-Shenq Juang and Chin-Laung Lei. Fair Blind Signatures Based on Discrete Logarithm. In *Proceedings of National Computer Symposium, Taiwan*, pages C95–C100, 1997.

[14] Anna Lysyanskaya, Ronald L. Rivest, Amit Sahai, and Stefan Wolf. Pseudonym Systems. In H. Heys and C. Adams, editors, *Proceedings of the 6th Annual International Workshop on Selected Areas in Cryptography*, Lecture Notes in Computer Science, LNCS 1758, pages 184–199. Springer-Verlag, 1999.

[15] David Martin and Andrew Schulman. Deanonymizing Users of the SafeWeb Anonymizing Service. Boston University Computer Science Technical Report, February 2002.

[16] Michael G. Reed, Paul F. Syverson, and David M. Goldschlag. Anonymous Connections and Onion Routing. *IEEE Journal on Selected Areas in Communications*, 16(4):482–494, May 1998. Special issue on Copyright and Privacy Protection.

[17] Michael K. Reiter and Aviel D. Rubin. Crowds: Anonymity for Web Transactions. *ACM Transactions on Information and System Security (TISSEC)*, 1(1):66–92, November 1998.

[18] Berry Schoenmakers. Basic Security of the ecash Payment System. In B. Preneel and V. Rijmen, editors, *Computer Security and Industrial Cryptography: State of the Art and Evolution*, Lecture Notes in Computer Science, LNCS 1528, pages 342–356. Springer-Verlag, June 1998.

[19] Clay Shields and Brian Neil Levine. A Protocol for Anonymous Communication Over the Internet. In *Proceedings of the 7th ACM Conference on Computer and Communications Security*, pages 33–42, November 2000.

[20] Daniel R. Simon. Anonymous Communication and Anonymous Cash. In N. Koblitz, editor, *Advances in Cryptology – CRYPTO'96*, Lecture Notes in Computer Science, LNCS 1109, pages 61–73. Springer-Verlag, 1996.

[21] Markus Stadler, Jean-Marc Piveteau, and Jan Camenisch. Fair Blind Signatures. In L.C. Guillou and J.-J. Quisquater, editors, *Advances in Cryptology – EUROCRYPT'95*, Lecture Notes in Computer Science, LNCS 921, pages 209–219. Springer-Verlag, 1995.

[22] Miriam van Dellen. Anonymity Law Survey. http://rechten.kub.nl/anonymity/.

[23] Zero-Knowledge Systems. Freedom Network. http://www.zeroknowledge.com/.

[24] Zero-Knowledge Systems. Zero-Knowledge Systems Discontinues "Freedom Network" Services as of October 11, 2001. Announcement, October 2001.