Accountable Anonymous Communication*

Claudia Diaz[1] and Bart Preneel[2]

[1] K.U.Leuven ESAT-COSIC
Kasteelpark Arenberg 10, B-3001 Leuven-Heverlee, Belgium
http://homes.esat.kuleuven.be/~cdiaz
claudia.diaz@esat.kuleuven.be
[2] K.U.Leuven ESAT-COSIC
Kasteelpark Arenberg 10, B-3001 Leuven-Heverlee, Belgium
http://homes.esat.kuleuven.be/~preneel
bart.preneel@esat.kuleuven.be

Summary. In this chapter we motivate the need for anonymity at the communication layer and describe the potential risks of having traceable communications. We then introduce the legal requirements on data retention and motivate the need for revocability of anonymity upon the request of law enforcement.

We describe the main building blocks for anonymous communication and for anonymity revocation. We explain how these building blocks can be combined in order to build a revocable anonymous communication infrastructure that fulfills both privacy and law enforcement requirements.

1 Introduction

Privacy is increasingly understood as an interdisciplinary subject. Legal, political and social considerations must be taken into account in the design of viable technical solutions that can be implemented at a large scale and accepted by the various players: citizens, governments, companies, etc. Anonymity and identity management technologies are powerful tools to protect privacy. Nevertheless, their potential for abuse is a factor that hinders the development and implementation of privacy enhancing systems at a large scale.

This chapter discusses the requirements that a large scale anonymity infrastructure should comply with in order to be acceptable for all parties. Anonymity infrastructures that protect the privacy of millions of individuals can only be possible if extreme care is taken in balancing the requirements of a multiplicity of interacting entities with sometimes conflicting interests. Otherwise, anonymity systems face the threat of remaining marginal in an environment in which privacy violations become ever more common.

* This work was supported by the IWT SBO project on Advanced Applications for electronic Identity cards in Flanders (ADAPID).

The chapter is structured as follows: we first motivate the need for anonymity and for accountability. In Section 3 we present the requirements for the system. Section 4 describes the building blocks that are used to construct our system. Section 5 describes the proposed model for an accountable anonymity infrastructure. Finally, Section 6 presents the conclusions of this work.

2 Anonymity and Accountability

2.1 Motivation for Communication Anonymity

Anonymity is defined by Pfitzmann and Hansen in [23] as *the state of being not identifiable within a set of subjects, the anonymity set.* This definition implies that, in order to achieve anonymity, we need a large population of users (*anonymity set*) performing actions in such a way that it is not possible to know which action was performed by which user. Users are more anonymous as the anonymity set and the indistinguishability increase (see [12, 26] for practical anonymity metrics).

According to Moore's Law, computer processing power doubles every 18 months. Storage capacity grows even faster, doubling every 13 months, according to Kryder's Law. If no anonymity infrastructure is put in place, all communication can (or will soon) be traced, registered, stored, mined, analyzed and aggregated. Furthermore, the collection of traffic data for law enforcement purposes has become a legal requirement. The 2006 approved EU Directive on Data Retention [15] imposes that all communication providers keep traffic data for law enforcement purposes. It is unclear how stored communication data at communication providers' databases will be secured against abuse for other purposes.

As communication data often leaks information of the content being accessed by users (e.g., `http://www.teensforteens.net/homosexuality/aids_and_homosexuality.html`), this means that personal data can technically (either legally or illegally) be collected (without the consent or awareness of the data subject), in order to build profiles of potential customers, potential employees, potential terrorists, etc. Individuals lose control over their personal data, which implies that they become vulnerable to all kinds of commercial and political manipulation. It is also clear that the large amounts of information available on individuals could be exploited for criminal purposes such as identity theft or targeted crime (e.g., it would be useful for burglars to know who are the wealthiest home-owners in a given area and when they are going on holiday). Privacy should therefore not be considered as contradictory with security: the lack of privacy protection may lead to serious security problems. Moreover, privacy protection is only possible using secure systems.

In order to avoid these privacy violations, we need to hide the communication patterns of Internet users towards untrusted recipients (e.g., web sites)

and external observers (e.g., local eavesdroppers). An anonymous communication infrastructure should therefore be in place in order to protect users' privacy.

2.2 Motivation for Accountability

If a system is deployed for massive use, abuse is unavoidable; moreover, the sense of impunity generated by the impossibility of holding people accountable could further encourage abuse. Without *accountability* mechanisms in place, it is unlikely that an unaccountable system could gather support from public powers and even from many citizens, as security has to be traded with privacy (or, at least, that is a common perception).

Nevertheless, there are strong arguments against trading anonymity with accountability [1]. A similar debate over key escrow for confidentiality took place in the mid 1990s about the Clipper Chip [28]. One of the problems was that the escrow algorithm was secret (i.e., not verifiable) and its hardware implementation was not tamper resistant. This is not the case with the system we propose here, which is based on public protocols and algorithms. The second argument against a key escrow system was that a voluntary system would not solve law enforcement's problems. Indeed, criminals could easily create their own keys, use them to communicate, and not give the escrow information to law enforcement authorities, rendering the whole escrow system useless.

As we can derive from the definition of anonymity (Section 2.1), there is a fundamental difference in the nature of confidentiality and anonymity in communication networks. Confidentiality of the content can be achieved by the communicating partners on their own: when establishing a shared secret, no third entity needs to participate. Even more, one could create a key for encrypting one's own data, without needing external entities. Anonymity is more complex. People act anonymously when their actions cannot be linked to their identities, or more precisely, when there is a set of subjects that could potentially be linked to the action, but there is not enough information to tell which of the subjects relates to the action. While confidentiality can be achieved by those who seek it alone, anonymity needs the cooperation of a group of people, the larger the better. Anonymity is therefore social (as it needs society to work together in order to be achieved), while confidentiality makes sense at the individual level.

While criminals would be able to bridge the key escrow systems using their own keys, they are not able to obtain anonymity on their own. If accountability mechanisms are built in the system, then the potential for abuse sharply decreases. Criminals may then choose not to use the system (exposing themselves to leave traces), or choose an unconditionally anonymous network. If this is the case, the people operating the network may find themselves in trouble, depending on the seriousness of the crime and on the legal framework in which they are operating.

It is still unclear how the EU Directive on Data Retention will affect unconditionally anonymous communication networks. There have been cases in the past in which law enforcement has forced anonymous communication providers (e.g., `anon.penet.fi` or JAP [19]) to either shut down the service, or violate the principle of providing unconditional anonymity to their users by implementing tracing capabilities that were not in the original design. We propose a *checks and balances* model which involves different entities in the anonymity revocation process, in order to ensure that the revocation policy is well understood from the beginning, and only technically possible in specific conditions.

2.3 Related Work on Anonymous Communication

Some of the earliest real-time *anonymous communication* systems were based on trusted or semi-trusted relays (e.g., Anonymizer [2] and SafeWeb). In centralized trust systems, the anonymity depends critically both on the security level and on the integrity of the service provider and its staff.

Pfitzmann et al. proposed in 1991 ISDN Mixes [24], a system to anonymize ISDN telephone conversations. Their design, based on a cascade of relays (mixes), was later adapted for anonymous web browsing and called Web Mixes [3]. A shortcoming of cascade topologies is that they require less effort for attacker to monitor the entry and exit point of the anonymity system. Part of the design has been implemented as a web anonymizing proxy, JAP. The use of multiple intermediate relays between the two ends of the communication improves the trust distribution over the use of a single relay, provided that if some of the relays are honest, the anonymity of the user remains protected. On the other hand, the cascade topology does not have good scalability and availability properties. The JAP design did not consider mechanisms for anonymity revocation; however, upon a law enforcement request for identification of a particular user, an exception had to be made in order to comply to the request.

Onion Routing [16, 17, 25, 27] is a free route mix network topology for unconditionally anonymous communication. Free route mix networks are vulnerable to intersection attacks [4]. The users establish circuits through a number of onion routers of their choice, and distribute symmetric keys to those routers. Data traveling in an established circuit is encrypted in layers, using the symmetric keys distributed to the routers. Tor (*The Onion Router*) [14], an improved second generation of Onion Routing, was proposed and implemented in 2004 (available at `http://tor.eff.org/`). Two years after deployment, it counts hundreds of volunteer nodes and hundreds of thousands of users, making it a very successful anonymous communication network.

Claessens et al. propose in [10] a system for revokable anonymous communication based on blind signatures. They introduce the legal requirements relevant for (revocable) anonymous communication and present a proof-of-concept architecture. Von Ahn at al. [29] propose transformations to add

selective traceability to anonymous systems based on threshold cryptography and group signatures. Kopsell et al. [21] proposed a revocable anonymity system based on threshold group signatures and threshold atomic proxy re-encryption.

All practical low latency anonymous communication systems are vulnerable to adversaries capable of monitoring the entry and exit points: high speed and high volume traffic patterns are too distinct in each connection, making it difficult to hide the correlation of the traffic going in and out [18]. End-to-end full padding solves this problem, but its deployment is very expensive. Whether intermediate solutions, using some cover traffic, can effectively hide the communication patterns, remains an open problem. Also, if all nodes in the anonymous communication path are corrupted by the adversary, then the communication is traceable.

3 Requirements

The system should comply with a basic set of requirements (see [13] for more details) which include:

- Application-independence: it should provide a general purpose low-latency bidirectional communication layer.
- Secure anonymity: untraceability of communication (i.e., the path connecting the communication parties is hard to discover), unlinkability of sessions (i.e., from an adversary's point of view, it is hard to link two sessions as related to the same user), load balancing mechanisms, secure implementation, usability and robustness against attacks.
- Availability: resistance against denial of service attacks and a sufficient number of entry and exit points.
- Scalability: the system must be able to provide service to large numbers of users.

The second set of requirements are an attempt to balance the fundamental right to privacy and the accountability mechanisms needed to make the system acceptable for the public at large and usable at a large scale. Claessens et al. define a complementary set of requirements in [10], where the legal point of view on anonymity and accountability for communication networks is presented.

- *Default privacy protection.* The system must be designed to protect by default the privacy of the users by anonymizing the communication layer. A user remains anonymous unless a judge issues a warrant that demands his identification.
- *Accountability.* The communication system must implement mechanisms that allow for law enforcement, called *identification* and *investigation*. Two types of actions may be considered. First, mechanisms should exist to identify subjects involved in criminal activities acting through the anonymous

system (*post factum*). Second, law enforcement agents should be able to conduct investigations of criminal networks (e.g., money laundering), that is, tracing of the communication of a user under criminal investigation.

- *Transparency.* Clear and public policies and contracts that define rights, obligations and liabilities of all parties, as well as the activities that may lead to identification, discourage abuse in the first place.
- *Trust Distribution* is one of the key aspects of the design of the system. In order to be accepted by all entities, the system must be trusted to provide anonymity for honest users, as well as transparent accountability mechanisms for those who abuse the system for criminal purposes. Trust should be distributed, in order to minimize the possibility of a collusion of entities illegally tracing or identifying a user.
- *Identity management at the user's side.* In order to empower the user in the management of his own identities, all the identity and profile information should be kept under the user's control. Users who wish to obtain personalized services may provide the preferences to their service provider without disclosing their identity (e.g., [11, 20]). Service providers may collect the anonymized data generated by the users' transactions, but they will not be able to link the behavioral data to an identifiable individual, neither to link the user's local pseudonym to other organizations' databases.

4 Building blocks

Here we present various technologies that can be combined in order to implement the proposed anonymity infrastructure. These include: a mix network that provides the anonymous communication functionality; an anonymous credential system to support the identity management and optional revocation; key traitor tracing schemes to support the revocation process; exit policies to distinguish resources with different abuse potential; and secure hardware modules to perform the cryptographic operations.

4.1 Mix network

A *mix* is a router that hides correspondence between inputs and outputs by performing cryptographic operations that provide bitwise unlinkability (semantic security); and by modifying the order of messages to hide timing correlations. A mix network is a network of interconnected mixes.

The core of the anonymity infrastructure is a mix network similar to Tor [14]. Users select a path of nodes in the network and create a circuit through them that later will carry the information of all applications using the anonymous communication infrastructure.

4.2 Anonymous credentials with optional anonymity revocation

Anonymous credential systems [6, 7, 8, 9, 22] allow anonymous yet authenticated and accountable transactions between users and service providers. The basic primitives provided by these systems allow for establishing pseudonyms and issuing, showing, verifying and deanonymizing credentials. All credentials and pseudonyms of a user are generated from a user master secret S_U. A full description of these protocols can be found in Chapter 15.

The anonymous credential infrastructure is a privacy-enhanced pseudonymous PKI which implements zero-knowledge protocols (see Chapter 15). A user U can establish a pseudonym N_I with an issuing organization O_I. U can also obtain a credential C signed by O_I certifying certain attributes. Later on, U may prove these attributes to a verifying organization O_V. In this system, the user may choose which attributes to prove to O_V (note that proving an attribute does not necessarily imply showing its value; for example, a user may prove he is older than 18 without actually showing his age). Multiple credential shows are unlinkable.

If a user is registered with several organizations, the pseudonyms established are not linkable. The anonymous credential system provides protocols that allow the user to prove ownership of multiple credentials.

These systems also implement optional revocation of anonymity for accountability purposes. In this case, the user must provide a verifiable encryption of his pseudonym or identity: he must encrypt the information with the public key of a trusted entity O_D and prove that it can be decrypted by O_D from the protocol transcript.

4.3 Key traitor tracing schemes

Key traitor tracing schemes [5] are public key encryption systems in which there is one public encryption key, and many private decryption keys. They also provide the security feature that if a coalition of private key holders collude to create a new decryption key, then there is an efficient algorithm to trace the new key to its creators.

4.4 Exit policies

Exit policies have been proposed in other systems [14] to provide flexibility for the nodes regarding the resources they want to give access to. In our model, we propose three categories of content, for which different rules apply:

- *Black list:* illegal content sites (e.g., child porn web sites) should not be accessible from the mix network.
- *White list:* there are many services for which unconditional anonymity is desired (see Section 5.1 for some examples). Such locations must be included in a white list, and access to white list resources should not require

the user to provide deanonymization information (encrypted pseudonym). This communication is therefore unconditionally anonymous. Note that in case of abuse, the identity of misbehaving users cannot be recovered.

- *No list:* this refers to sites and applications which have a potential for criminal abuse and may require accountability mechanisms. Users are asked by the exit mix to provide a verifiable encryption of their pseudonym in order to access unlisted content (once for the duration of the tunnel). Note that users should select (at least) two separated exit nodes for *white list* and *no list* requests. If they use the same exit node for both types of requests, their *white list* requests would be technically linkable to their pseudonym (due to the *no list* requests and the linkability of requests in the same tunnel).

Note that exit nodes are the "visible" initiator of the requests they route. If abuse is committed through them, they may face problems with law enforcement. It is therefore in their interest to require an encrypted identity to give access to resources with a potential for abuse (e.g., trade sites). Letting the mixes design their own white lists gives more flexibility to the system. On the other hand, agreeing on common exit policies helps keeping the system simpler. We leave the design of the white and black listing as an open question. One option would be self-regulation by the mixes, and another option would be to regulate the white and black listing by law.

4.5 Secure Hardware Modules

To prevent credentials from being stolen from a user, as well as users sharing their credentials (and thus undermining their value), it is helpful to execute the credential protocols in a secure subsystem that is protected from both other applications on the same platform and from the user. The Trusted Platform Module (TPM), described in Chapter 9 of this book, provides this functionality. Other alternatives include smart cards and secure coprocessors.

5 Proposed accountable anonymity infrastructure

In this section, we present the proposed model for building a large-scale accountable anonymity system. We first introduce the participating entities, and second we describe the protocols.

5.1 Entities

We introduce here the entities playing a role in the system. For each of these entities, we describe their role, what they are trusted for, the secrets they keep, the credentials or public keys they possess and the information they see and log.

Root Authority.

The Root Authority (RA) is a credential issuer. Users registered with this authority with a pseudonym N_I may obtain a credential C_I on N_I signed by the RA. The RA is trusted to keep the civil identity of the user (i.e., an identifier with which the administration can uniquely identify the citizen; e.g., the passport number), together with N_I. Note that users may need to register in person with the RA.

The RA is required as part of the anonymous credential infrastructure that should be in place to support our system, but it is not part of the anonymity infrastructure itself. In order to provide better availability and robustness, the RA may be distributed across several entities.

Mixes.

Mixes are the core of the anonymity service. They act as registration authority for issuing credentials to users, and they anonymize the communication of registered users.

In order to obtain access to the anonymity network, users must first register with a node of their choice. The user establishes a pseudonym N_A with the node, (i) proving ownership of a credential C_I from a trusted RA, (ii) providing a verifiable encryption of his pseudonym with the RA, N_I, and (iii) paying the registration fee (one-show anonymous credentials can implement anonymous coins [6]). The node that registers the user logs N_A and the encrypted N_I. The user then requests from the node a credential C_A issued on N_A and signed by the mix.

In order to prevent abuse, nodes maintain a pseudonym revocation list which lists pseudonyms that should no longer be accepted at entry nodes. Users who abuse the system may have their pseudonym revoked.

Mixes have two public key pairs; one to receive encrypted information (PK_i), and another to digitally sign messages. Mixes also establish shared secret keys with their neighboring nodes for link encryption.

When a mix is the entry node of a tunnel, it must verify the validity of the credential C_A presented by the user, as well as the pseudonym of the user N_A; it should also verify that N_A has not been included in the pseudonym revocation list. C_A may have been issued by the mix itself, or by other mixes of the network. If C_A and N_A are valid, the user can continue to construct the tunnel.

The exit node of a tunnel must follow its exit policy. When the user is accessing white list resources, the exit node does not need to log any information. If the user accesses unlisted resources (for which accountability measures may be needed), then the exit node must ask a verifiable encryption of N_A, keep it for the time established in data retention laws, and log with it a timestamp and the unlisted resources requested by the user.

Nodes, besides being trusted for not logging information which is not strictly required for accountability measures, are trusted to:

- verify the validity of the encryption of N_I when registering a user;
- securely keep the encrypted N_I and N_A for registered users;
- verify the validity of C_A and N_A when acting as entry node;
- verify the validity of the encryption of N_A when acting as exit node and giving access to an unlisted resource;
- securely keep the encrypted N_A and the accesses to unlisted resources for the period of time specified in data retention laws, and delete it afterwards;
- collaborate with law enforcement when required by a judge.

User.

The user is registered with an RA and with the anonymity infrastructure (he may choose the Root Authority and Registration Mix he trusts most), and routes anonymously his communication through the mix network. The user is trusted to securely store his master secret S_U (e.g., in a TPM or in a smart card). S_U is used to establish pseudonyms (N_I, N_A), to obtain credentials (C_I, C_A), and in the credential show protocols.

Resources.

Resources are classified by the nodes of the anonymity network according to the three categories mentioned in Section 4.4: white list, black list and no list. Note that some of these services may as well implement their own conditionally anonymous credentials at the application layer, or even require user authentication.

Judges.

Judges are entities legally authorized to issue a warrant that demands identification of a user accused of criminal activities. They may also request mixes to collaborate in criminal investigations that require tracing of users.

A key traitor tracing scheme is implemented in order to make the identification procedure effective for law enforcement. In this scheme, there is a unique encryption public key PK_J of the judges, but every judge holds a unique private key for decryption. Users verifiably encrypt their pseudonyms N_I and N_A with PK_J when required to provide an encrypted identity as accountability condition in certain protocols.

Judges are trusted to apply the law and protect their private keys against disclosure or abuse.

5.2 Protocols

We outline in this section the protocols for the proposed system. The protocols are based on the anonymous credential protocols described in [6, 7].

Registration with Root Authority.

The user must be registered with an RA which is trusted by the anonymity network. The RA knows the user by a pseudonym N_I, and can link this pseudonym to the real identity of the user. If national electronic ID cards implement anonymous credentials, they could be used as root credentials. In this case, a local authority (e.g., in Belgium the City Hall) would act as trusted Root Authority.

Registration with Anonymity Infrastructure.

The user U chooses one of the mixes, according to his trust preferences, to register for the system. With the registration request the user must:

- prove possession of a root credential C_I and knowledge of the secret S_U;
- provide a verifiable encryption of N_I, encrypted with PK_J;
- pay for the service (anonymous payments may be implemented with one-show anonymous credentials);
- establish a pseudonym N_A with the mix;
- request a credential C_A on N_A signed by the mix, which grants access to the network.

The mix must verify the correctness of the proofs, and issue a credential C_A on N_A for the user. This credential will have a certain validity period, and must be accepted by other nodes in the network. Once the user has established N_A with the mix, he may ask for new credentials once C_A has expired. Note that, during the registration process, the user should be informed of the terms of use of the network, and on the policies and conditions for identification.

Using the Anonymity Infrastructure.

Once the user U has obtained C_A, he can start to use the service. U selects a path through the network for his tunnel, and connects to the entry node, M_1.

U has to prove to M_1 that he has a valid credential C_A issued by a mix of the network. He must also show N_A so the entry mix can check that his pseudonym has not been revoked.

If the verification of C_A and N_A is ok, the user may continue to establish the tunnel, contacting the other nodes in the path. The exit mix of a connection used to access unlisted resources asks a verifiable encryption of N_A, encrypted with PK_J. The exit mix logs the encrypted pseudonym and the IP addresses of the accessed unlisted resources, as well as a timestamp. The mix may keep this information for the time established in the legal framework in which the node is operating. After that period of time, the data must be deleted. The next item explains how the identification of the user is carried out if accountability measures need to be applied.

The user may want to use separate exit nodes (or tunnels) for different requests in very sensitive cases; e.g., he may want to use two exit nodes (or tunnels) to check information on a particular disease and to access the conditions of a life insurance).

Identification.

If the user abuses the anonymous network for criminal purposes (e.g., dealing with child porn, uploading a video which shows the murder of a hostage, or swindling people in eBay), then his identification may be required by a judge.

Note that identification can only be requested when the user has used an unlisted resource (accesses to white list resources are unconditionally anonymous). If that was the case, U has provided to the exit node a verifiable encryption of N_A. As the address of the node is visible in the access to the resource, the judge may request the exit node to provide the encrypted pseudonym of the user who accessed a particular resource at a particular time. As N_A was encrypted with PK_J, the judge can extract the pseudonym N_A of the user in the network, and the name of the issuing mix.

The judge may contact the mix that registered the user, provide N_A and request the encrypted N_I that is kept in the mix. The judge may, again, decrypt this pseudonym and recover the root pseudonym, N_I. The Root Authority may now provide the judge with the identity of the subject.

In this system, four unrelated entities (the issuer node, the exit node, the Root Authority and the judge) need to collaborate in order to identify a user. This is done in order to distribute the trust and to minimize the possibility of collusion of entities for illegal identification.

Investigation.

In some cases (e.g., money laundering, organized crime, terrorism, etc.) law enforcement needs to investigate individuals. In this model, a judge may sign a warrant asking the mixes for cooperation in an investigation on the user known by N_A. When the user contacts the entry node, he must show N_A. The mix can check if the N_A matches the name in the warrant issued by the judge, and if so, it forwards the warrant to the next mix in the path of the user, and logs the necessary information to allow tracing. The next mixes in the path also log the necessary information regarding that tunnel, and forward the warrant to the next node. The exit node receives the warrant and logs the required information. All mixes encrypt the logged information with the judges' public key, PK_J.

The judge may later request the information from the mixes regarding the N_A on the warrant and reconstruct the activities of the user. The entities that need to collaborate in this case are a judge, and all mixes in the tunnel. Note the judge needs to know N_A, which may not be obvious. If the user has been previously identified because of criminal abuse of the network, then N_A is already known to the judge. Otherwise, the judge could use the IP address of the user to ask entry nodes for the pseudonym. This protocol would however require some modifications (as entry nodes do not keep N_A in our model), and falls outside the scope of this chapter.

5.3 Deployment

In order to deploy such an accountable anonymity infrastructure, certain technical and political conditions should exist:

- We are assuming large computing power and high speed links in order to implement the mix network. Today, these resources are still too expensive. The design of secure efficient anonymous communication networks remains a challenge.
- A trusted anonymous credential infrastructure (Root Authority) needs to be in place. The implementation of anonymous credentials in privacy-enhanced national electronic identity cards would provide a solid anonymous credential infrastructure.
- The institutional and legal framework is also important, as the system requires a well coordinated judicial system.
- Finally, there should be social and political support for a large-scale infrastructure. If privacy protection becomes a political priority, or if the privacy awareness of the citizens significantly increases, the support could be gathered.

6 Conclusions

In this chapter, we have proposed an accountable anonymity communication infrastructure. We have motivated the need for implementing both anonymity and accountability mechanisms and argued the differences with the key escrow debate.

We have listed the requirements for a large scale anonymous infrastructure, and proposed a model to comply with them. Our model is built by combining existing technologies and distributes trust among various entities.

The purpose of this work is to provide a solution that combines existing privacy enhancing technologies, rather than presenting new contributions for the building blocks themselves. Here we summarize the main conclusions of this work:

- Both anonymity and accountability requirements should be satisfied in order to gain support for the deployment of large scale anonymity infrastructures that provide privacy protection to the *masses*. Trust should be adequately distributed to minimize the risk of collusion.
- Existing cryptographic primitives and privacy-enhancing technologies can be combined in order to build an anonymous communication infrastructure that is both flexible and accountable.
- The implementation of such a system would allow for user controlled identity management. Organizations may collect pseudonymous data from users, but they may not link the information to the identity of the user unless he willingly provides it.

- Unconditionally anonymous, pseudonymous and authenticated applications may communicate through this system. Anonymous credential and authentication protocols may be implemented at the application layer.
- The system must distribute the ability for revocation in order to be trustworthy for the users.
- The biggest shortcoming of the system is its cost. It needs a large amount of resources in terms of computing power, bandwidth and organizational and administrative overhead. Viable economic models need to be applied in order to deploy the system.
- The technical and political conditions for deployment of an infrastructure of these characteristics are not met today, but they may be met in the future if privacy protection becomes a real concern.

References

1. Abelson H, Anderson R, Bellovin S, Benaloh J, Blaze M, Diffie W, Gilmore J, Neumann P, Rivest R, Schiller J, Schneier B (1997) The Risks of Key Recovery, Key Escrow, and Trusted Third Party Encryption. World Wide Web Journal 2(3):241-257.
2. Anonymizer, http://www.anonymizer.com/
3. Berthold O, Federrath H, Kopsell S (2000) Web MIXes: A system for anonymous and unobservable Internet access. In: Federrath H (ed.) Designing Privacy Enhancing Technologies, LNCS 2009, pp. 115-129. Springer-Verlag.
4. Berthold O, Pfitzmann A, Standtke R (2000) The disadvantages of free MIX routes and how to overcome them. In: Federrath H (ed.) Designing Privacy Enhancing Technologies, LNCS 2009, pp. 30-45. Springer-Verlag.
5. Boneh D, Franklin M (1999) An Efficient Public Key Traitor Tracing Scheme. In: Wiener M (ed.) Advances in Cryptology - CRYPTO'99, LNCS 1666, pp. 338-353. Springer-Verlag.
6. Camenisch J, Lysyanskaya A (2001) An Efficient System for Non-transferable Anonymous Credentials with Optional Anonymity Revocation. In: Pfitzmann B (ed.) Advances in Cryptology - EUROCRYPT'01, LNCS 2045, pp. 93-118. Springer-Verlag
7. Camenisch J, Van Herreweghen E (2002) Design and Implementation of the idemix Anonymous Credential System. In: Atluri V (ed.) Proceedings of the 9th ACM conference on Computer and Communications Security, pp. 21-30. ACM Press.
8. Chaum D (1985) Security without Identification: Transaction Systems to Make Big Brother Obsolete. Communications of the ACM 28(10):1030-1044.
9. Chaum D, Evertse J (1987) A Secure and Privacy-Protecting Protocol for Transmitting Personal Information Between Organizations. In: Odlyzko A (ed.) Advances in Cryptology - CRYPTO'86, LNCS 263, pp. 118-167. Springer-Verlag.
10. Claessens J, Diaz C, Goemans C, Preneel B, Vandewalle J, Dumortier J (2003) Revocable anonymous access to the Internet. Journal of Internet Research 13(4):242-258.

11. Claessens J, Diaz C, Preneel B, Vandewalle J (2002) A Privacy-Preserving Web Banner System for Targeted Advertising. Technical Report 9 p. Katholieke Universiteit Leuven.
12. Diaz C, Seys S, Claessens J, Preneel B (2002) Towards Measuring Anonymity. In: Dingledine R, Syverson P (eds) Designing Privacy Enhancing Technologies, LNCS 2482, pp. 54-68. Springer-Verlag.
13. Diaz C, Naessens V, Nikova S, De Decker B, Preneel B (2004) Tools for Technologies and Applications of Controlled Anonymity. Technical Report, 211 p. Project IWT STWW Anonymity and Privacy in Electronic Services.
14. Dingledine R, Mathewson N, Syverson P (2004) Tor: The Second-Generation Onion Router. In 13th USENIX Security Symposium, pp. 303-320. USENIX.
15. Directive 2006/24/EC of the European Parliament and of the Council (13.4.2006) Official Journal of the European Union.
16. Goldschlag D, Reed M, Syverson P (1996) Hiding Routing Information. In: Anderson R (ed.) Information Hiding, LNCS 1174, pp. 137-150. Springer-Verlag.
17. Goldschlag D, Reed M, Syverson P (1999) Onion Routing. In: Communications of the ACM 42(2):39-41.
18. Hintz A (2002) Fingerprinting Websites Using Traffic Analysis. In: Dingledine R, Syverson P (eds) Designing Privacy Enhancing Technologies, LNCS 2482, pp. 171-178. Springer-Verlag.
19. JAP Anonymity & Privacy, http://anon.inf.tu-dresden.de/
20. Juels A (2001) Targeted Advertising... and Privacy Too. In: Naccache D (ed.) Topics in Cryptology - Proceedings of the Cryptographers' Track at RSA'2001, LNCS 2020, pp. 408-424. Springer-Verlag.
21. Kopsell S, Wendolsky R, Federrath H (2006) Revocable Anonymity. In: Muller G (ed.): Emerging Trends in Information and Communication Security - ETRICS, LNCS 3995, pp. 206-220. Springer-Verlag.
22. Lysyanskaya A, Rivest R, Sahai A, Wolf S (1999) Pseudonym Systems. In: Heys H, Adams C (eds) Proceedings of the 6th Annual International Workshop on Selected Areas in Cryptography, LNCS 1758, pp. 184-199. Springer-Verlag.
23. Pfitzmann A, Hansen M (2000) Anonymity, Unobservability and Pseudonymity: A Proposal for Terminology. In: Federrath H (ed.) Designing Privacy Enhancing Technologies, LNCS 2009, pp. 1-9. Springer-Verlag.
24. Pfitzmann A, Pfitzmann B, Waidner M (1991) ISDN-mixes: Untraceable communication with very small bandwidth overhead. In: Effelsberg W, Meuer H, Muller G (eds) GI/ITG Conference on Communication in Distributed Systems, Informatik-Fachberichte 267, pp. 451-463. Springer-Verlag.
25. Reed M, Syverson P, Goldschlag D (1998) Anonymous Connections and Onion Routing. In: IEEE Journal on Selected Areas in Communications 16(4):482-494.
26. Serjantov A, Danezis G (2002) Towards an Information Theoretic Metric for Anonymity. In: Dingledine R, Syverson P (eds) Designing Privacy Enhancing Technologies, LNCS 2482, pp. 41-53. Springer-Verlag.
27. Syverson P, Tsudik G, Reed M, Landwehr C (2000) Towards an Analysis of Onion Routing Security. In: Federrath H (ed.) Designing Privacy Enhancing Technologies, LNCS 2009, pp. 96-114. Springer-Verlag.
28. The Clipper Chip, http://www.epic.org/crypto/clipper/
29. Von Ahn L, Bortz A, Hopper N, O'Neill K (2006) Selectively Traceable Anonymity. In: Danezis G, Golle P (eds) Designing Privacy Enhancing Technologies, LNCS (pre-proceedings), pp. 199-213. Springer-Verlag.