

A Survey of Anonymous Communication Channels

George Danezis

K.U. Leuven, ESAT/COSIC,

Kasteelpark Arenberg 10, B-3001 Leuven-Heverlee, Belgium.

GEORGE.DANEZIS@ESAT.KULEUVEN.BE

Claudia Diaz

K.U. Leuven, ESAT/COSIC,

Kasteelpark Arenberg 10, B-3001 Leuven-Heverlee, Belgium.

CLAUDIA.DIAZ@ESAT.KULEUVEN.BE

Abstract

We present an overview of the field of anonymous communications, from its establishment in 1981 from David Chaum to today. Key systems are presented categorized according to their underlying principles: semi-trusted relays, mix systems, remailers, onion routing, and systems to provide robust mixing. We include extended discussions of the threat models and usage models that different schemes provide, and the trade-offs between the security properties offered and the communication characteristics different systems support.

Keywords: Anonymous communications, Mix systems, Traffic analysis, Privacy.

1. Introducing Anonymous Communications

Research on anonymous communications started in 1981 with Chaum's seminal paper "Untraceable electronic mail, return addresses, and digital pseudonyms" (Chaum (1981)). Since then, a body of research has concentrated on building, analyzing and attacking anonymous communication systems. In this survey we look at the definition of anonymous communications and the major anonymous communication systems grouped in families according to the key design decisions they are based on.

Data communication networks use addresses to perform routing which are, as a rule, visible to anyone observing the network. Often addresses (such as IP addresses, or Ethernet MACs) are a unique identifier which appear in all communication of a user, linking of all the user's transactions. Furthermore these persistent addresses can be linked to physical persons, seriously compromising their privacy.

Anonymizing the communication layer is thus a necessary measure to protect the privacy of users, and protect computer systems against traffic analysis. They also support anonymization techniques at the application layer, such as anonymous credentials, elections and anonymous cash.

In the remaining of this introductory section, we introduce the terminology of anonymity properties, we present the existing models for quantifying the anonymity, and we explain some limits on anonymity imposed by black box attacks. Section 2 presents anonymity systems based on (cen-

Manuscript received XXX, accepted YYY, published ZZZ

tralized) trusted and semi-trusted relays, and introduces the link between anonymous communications and censorship resistance. Mix-based anonymous communication systems are extensively described in Section 3. Section 4 introduces other proposals for anonymous communication, and Section 5 presents the conclusions of this survey.

1.1 Terminology

Prior to the quantification of anonymity, a set of working definitions for *anonymity* and other related concepts, such as *unlinkability* or *unobservability* were needed.

Pfitzmann and Hansen (2001) proposed a set of working definitions for anonymity, unlinkability, unobservability and pseudonymity. These definitions have since been adopted in most of the anonymity literature. Their authors continue releasing regular updates on the document addressing feedback from the research community¹.

Anonymity. To enable the anonymity of a subject, there always has to be an appropriate set of subjects with potentially the same attributes. Anonymity is thus defined as *the state of being not identifiable within a set of subjects, the anonymity set*.

The *anonymity set* is the set of all possible subjects. With respect to acting entities, the anonymity set consists of the subjects who might cause an action. With respect to addressees, the anonymity set consists of the subjects who might be addressed. Both anonymity sets may be disjoint, be the same, or they may overlap. The anonymity sets may vary over time.

According to the Pfitzmann-Hansen definition of anonymity, the subjects who may be related to an anonymous transaction constitute the *anonymity set* for that particular transaction. A subject carries on the transaction *anonymously* if he cannot be distinguished (by an adversary) from other subjects. This definition of anonymity captures the probabilistic information often obtained by adversaries trying to identify anonymous subjects.

Unlinkability. The [ISO15408 1999] defines unlinkability as follows:

"[Unlinkability] ensures that a user may make multiple uses of resources or services without others being able to link these uses together. [...] Unlinkability requires that users and/or subjects are unable to determine whether the same user caused certain specific operations in the system."

We may differentiate between "absolute unlinkability" (as in the given ISO definition above; i.e., "no determination of a link between uses") and "relative unlinkability" (i.e., "no change of knowledge about a link between uses"), where "relative unlinkability" could be defined as follows:

Unlinkability of two or more Items Of Interest (IOIs; e.g., subjects, messages, events, actions, ...) means that within the system (comprising these and possibly other items), from the attackers perspective, these items of interest are no more and no less related after his observation than they were related concerning his a-priori knowledge.

This means that the probability of those items being related from the attackers perspective stays the same before (a-priori knowledge) and after the attackers observation (a-posteriori knowledge

1. http://dud.inf.tu-dresden.de/Anon_Terminology.shtml

of the attacker). Roughly speaking, providing relative unlinkability of items means that the ability of the attacker to relate these items does not increase by observing the system.

Unobservability. In contrast to anonymity and unlinkability, where not the IOI, but only its relationship to IDs or other IOIs is protected, for unobservability, the IOIs are protected as such. *Unobservability is the state of items of interest (IOIs) being indistinguishable from any IOI (of the same type) at all.*

This means that messages are not discernible from random noise. As we had anonymity sets of subjects with respect to anonymity, we have unobservability sets of subjects with respect to unobservability. Sender unobservability then means that it is not noticeable whether any sender within the unobservability set sends. Recipient unobservability then means that it is not noticeable whether any recipient within the unobservability set receives. Relationship unobservability then means that it is not noticeable whether anything is sent out of a set of could-be senders to a set of could-be recipients. In other words, it is not noticeable whether within the relationship unobservability set of all possible sender-recipient-pairs, a message is exchanged in any relationship.

Pseudonymity. Pseudonyms are identifiers of subjects. We can generalize pseudonyms to be identifiers of sets of subjects. The subject which the pseudonym refers to is the holder of the pseudonym.

Being pseudonymous is the state of using a pseudonym as ID.

We assume that each pseudonym refers to exactly one holder, invariant over time, being not transferred to other subjects. Specific kinds of pseudonyms may extend this setting: a group pseudonym refers to a set of holders; a transferable pseudonym can be transferred from one holder to another subject becoming its holder. Such a group pseudonym may induce an anonymity set: Using the information provided by the pseudonym only, an attacker cannot decide whether an action was performed by a specific person within the set.

Defining the process of preparing for the use of pseudonyms, e.g., by establishing certain rules how to identify holders of pseudonyms, leads to the more general notion of pseudonymity:

Pseudonymity is the use of pseudonyms as IDs.

An advantage of pseudonymity technologies is that accountability for misbehavior can be enforced. Also, persistent pseudonyms allow their owners to build a pseudonymous reputation over time.

1.2 Anonymity Metrics

Most attacks on anonymous communication networks provide the adversary with probabilistic information on the identity of the entities communicating with each other. This is the reason why information theoretical anonymity metrics (Serjantov and Danezis (2002); Diaz et al. (2002)) have been widely adopted to quantify the anonymity provided by a variety of designs.

But before information theoretic anonymity metrics were proposed, there had been some attempts to quantify anonymity in communication networks.

Reiter and Rubin (1998) define the *degree of anonymity* as a probability $1 - p$, where p is the probability assigned by an attacker to potential senders. In this model, users are more anonymous as they appear (towards a certain adversary) to be less likely of having sent a message. This metric considers users separately, and therefore does not capture anonymity properties very well. Consider a first system with 2 users. The first user u_1 appears to be the sender of a message with probability $1/10$, and the second user u_2 with probability $9/10$. Now consider a second system with 1000 users. User u_1 appears as the sender with probability $1/10$, while all the other users are assigned probabilities of having sent the message below 0.001. According to the definition of Reiter and Rubin, the *degree of anonymity* of u_1 would be the same in both systems ($d = 0.9$). However, in the second system, u_1 looks much more likely to be the sender than any other user, while in the first he is the less likely candidate of being so.

Berthold et al. (2000b) define the *degree of anonymity* as $A = \log_2(N)$, where N is the number of users of the system. This metric only depends on the number of users of the system, and therefore does not express the anonymity properties of different systems. The total number N of users may not even be known. Moreover, adversaries may be able to obtain probabilistic information on the set of potential senders, which is not taken into account in this metric.

Information theoretic anonymity metrics were independently proposed in two papers presented at the *2nd Workshop on Privacy Enhancing Technologies*. The basic principle of both metrics is the same. The metric proposed by Serjantov (2002) uses entropy as measure of the *effective anonymity set size*. The metric proposed by Diaz et al. (2002) goes one step further, normalizing the entropy to obtain a *degree of anonymity* in the scale 0..1.

The quantification of anonymity is dependent on the adversary considered. The adversary has certain capabilities and deploys attacks in order to gain information and find links between subjects and items of interest. Most of these attacks lead to a distribution of probabilities that assign subjects a certain probability of being linked to the items of interest. In this respect, a clear and detailed formulation of the attack model considered is a required step to measure the anonymity provided towards that attacker.

The information theoretic concept of entropy (Shannon (1948)) provides a measure of the uncertainty of a random variable. Let X be the discrete random variable with probability mass function $p_i = Pr(X = i)$, where i represents each possible value that X may take with probability $p_i > 0$. In this case, each i corresponds to a subject of the anonymity set; i.e., p_i is the probability of subject i being linked to the item of interest.

The entropy describes thus the information (measured in bits) contained in the probability distribution that describes the links between a set of subjects (the anonymity set) and an item of interest. In Serjantov and Danezis (2002), the entropy is proposed as a measure of the effective anonymity set size. If the entropy is normalized by the maximum the system could provide (if it was perfect and leaked no information) for a given number of users, we obtain a degree of anonymity (Diaz et al. (2002)) that gives a measure of the anonymity provider's performance.

1.3 Limits and black box attacks

No matter how good the anonymity performed by the network, persistent communication between two users will eventually be detected just by observing the edges of the network and correlating the

activity at the two ends of the communication. Such *intersection attacks*, are presented in Berthold et al. (2000b), who try to extract the sender of a stream of messages by intersecting the sender anonymity sets of consecutive messages sent by a user. This attack model is also considered in Kesdogan et al. (2002a); Agrawal et al. (2003) and Wright et al. (2003). The statistical variant of this attack is the statistical disclosure attack presented in Danezis (2003b); Danezis and Serjantov (2004); Mathewson and Dingledine (2004).

In Wright et al. (2002) the authors present a set of attacks that can be performed by a group of subverted network nodes. Against mix networks, they calculate the number of routes to be chosen between a sender and a receiver before the full path has been entirely populated by subverted nodes. They also examine the effect that fixed or variable length paths have on this probability. Similar results are found for Crowds and DC-nets. In Wright et al. (2003) they extend their analysis to considering a subset of network nodes that simply log all traffic, and provide bounds on how quickly an intersection attack can be performed. Despite these studies being set in the frame of particular systems, like DC-nets and Crowds, they in fact explore fundamental limitations for any systems that select trusted parties at random from a larger set of potentially corrupt nodes to provide security.

2. Trusted and semi-trusted relays

We start presenting anonymous communications by introducing systems that rely on one central trusted node to provide security. We will see that they provide a varying, but usually low, degree of anonymity protection against traffic analysis and active attacks.

2.1 The Anon.penet.fi relay

Johan Helsingius started running a trusted mail relay, `anon.penet.fi`, providing anonymous and pseudonymous email accounts in 1993. The technical principle behind the service was a table of correspondences between real email addresses and pseudonymous addresses, kept by the server. Email to a pseudonym would be forwarded to the real user. Email from a pseudonym was stripped of all identifying information and forwarded to the recipient. While users receiving or sending email to a pseudonym would not be able to find out the real email address of their anonymous correspondent, it would be trivial for a local passive attacker or the service itself to uncover the correspondence by correlating the timing of incoming and outgoing email traffic.

While protecting from a very weak threat model, the service was finally forced to close down through legal attacks. In 1996 the “Church of Spiritual Technology, Religious Technology Center and New Era Publications International Spa” reported that a user of `anon.penet.fi` sent a message to a newsgroup infringing their copyright. Johan Helsingius, the administrator of `anon.penet.fi`, was asked to reveal the identity of the user concerned. The details of the case, that put enormous strain on the service, can be found in the press releases of the 23 September 1996 (Helsingius (1996b,c)) or the information center set up around this case (Newman (1997)). Reputation attacks were also experienced, when unfounded reports appeared in mainstream newspapers about the service being used to disseminate child pornography (Helsingius (1996a)).

The service finally closed in August 1996 since it could no longer guarantee the anonymity of its users. The closure was quite significant for the privacy and anonymity research community. In the

initial judgment the judge had ruled that “a witness cannot refrain from revealing his information in a trial” (Helsingius (1996c)), even though an appeal was lodged on the grounds of privacy rights protected by the Finnish constitution, and the fact that the information might be privileged, as is the case for journalistic material.

The concept that even non-malicious relay operators could be forced under legal or other compulsion, to reveal any information they have access to, provided a new twist to the conventional threat models. Honest relays or trusted nodes could under some circumstances be forced to reveal any information they held concerning the origin or destination of a particular communication. Minimizing the information held by trusted parties is therefore not just protecting their users, but also the services themselves.

2.2 Anonymizer & SafeWeb

Anonymizer² is a company set up by Lance Cottrell, also author of the Mixmaster remailer software, that provides anonymous web browsing for subscribed users. The Anonymizer product acts as a web proxy through which all web requests and replies are relayed. The web servers accessed, should therefore not be able to extract any information about the address of the requesting user. Special care is taken to filter out any “active” content, such as javascript or Java applets, that could execute code on the user’s machine, and then signal back identifying information.

As for `anon.penet.fi`, the anonymity provided depends critically on the integrity of the Anonymizer company and its staff. The service is less vulnerable to legal compulsion attacks, since no long-term logs are required to be kept, that could link users with resources accessed. Unlike email, users always initiate web requests, and receive the replies, and all records can be deleted after the request and reply have been processed. Records can be made unavailable to seize just a few seconds after the provision of the anonymous service to the user.

SafeWeb was a company that provided a very similar service to Anonymizer. The two main differences in their initial products, was that SafeWeb allowed the traffic on the link from SafeWeb to the user to be encrypted using SSL (Dierks and Allen (1999)), and “made safe” active content in pages by using special wrapper functions. Unfortunately their system of wrappers did not resist a set of attacks devised by Martin and Schulman (2002). Simple javascript attacks turned out to be able to extract identifying information from the users.

In the absence of any padding or mixing, a passive attacker observing the service would also be able to trivially link users with pages accessed, despite the use of SSL. This vulnerability was studied in danezis; Bissias et al. (2005); Sun et al. (2002); Cheng and Avnur; Hintz (2002). This line of research established that an adversary is capable of compiling a library of ‘traffic signatures’ for all interesting web-pages that may be accessed. The signatures can then be compared with the traffic characteristics of the encrypted SSL connection to infer the page that was accessed.

The key weaknesses come down to the shape of traffic, which is inadequately padded and concealed. Browsers request resources, often HTML pages, that are also associated with additional resources (images, style sheets, ...). These are downloaded through an encrypted link, yet their size is apparent to an observer, and can be used to infer which pages are accessed. There are

2. <http://www.anonymizer.com/>

many variants of this attack: some attempt to build a profile of the web-site pages and guess for that which pages are being accessed while others use these techniques to beat naive anonymizing SSL proxies. In the latter case, the attacker has access to the cleartext input streams and he tries to match them to encrypted connections made to the proxy.

Note that latent structure and contextual knowledge are again of great use to extract information from traffic analysis: in danezis, it is assumed that users will mostly follow links between different web resources. A Hidden Markov Model is then used to trace the most likely browsing paths a user may have taken given only the lengths of the resources that can be observed. This provides much faster and more reliable results than considering users that browse at random, or web-sites that have no structure at all.

2.2.1 Censorship resistance & access to anonymous communications

The threat model that SafeWeb wished to protect against was also very peculiar. The company was partly funded by the United States Central Intelligence Agency (CIA), and attempted to secure funding from The Voice of America and the Internet Broadcasting Bureau in order to “help Chinese and other foreign Web users get information banned in their own company (*sic*)”³ (Singer (2001)). This claim explicitly links anonymous communications with the ability to provide censorship resistance properties. The link has since then become popular, and often anonymity properties are seen as a pre-requisite for allowing censorship resistant publishing and access to resources. No meticulous requirements engineering study has even been performed that proves (or disproves) that claim, and no cost benefit analysis has ever been performed to judge if the technical cost of an anonymity system would justify the benefits in the eyes of those interested in protecting themselves against censorship. Furthermore no details were ever provided, besides hearsay claims, about groups using this particular technology in a hostile environment, and their experiences with it. The latter would be particularly interesting given the known vulnerabilities of the product at the time.

The first paper to make a clear connection between censorship resistant storage and distribution is Anderson's Eternity service (Anderson (1996)). Serjantov has also done some interesting work on how to use strong anonymous communications to provide censorship resistance (Serjantov (2002)). The system presented is, for good technical and security reasons, very expensive in terms of communication costs and latency. Peer-to-peer storage and retrieval systems such as Freenet (Clarke et al. (2000)), FreeHaven (Dingledine et al. (2000)) and, more recently, GnuNet (Bennett and Grothoff (2003)) also claimed to provide anonymous communications. Attacks against some anonymity properties provided by GnuNet have been found (Kügler (2003)). Since the design of the three mentioned systems changes frequently it is very difficult to assess the security, or the anonymity they provide at any time. Feamster et al. (2002, 2003) have looked at different aspects of web censorship resistance by making use of steganography to send high volumes of data, and what they called “URL hopping” to transmit requests. Finally, aChord (Hazel and Wiley (2002)) presents concisely the requirements of a censorship resistant system, and attempts to build one based on a distributed hash table primitive.

3. The Freudian slip confusing “country” with “company”, and the way this goal can be understood in two opposing ways might be interpreted as quite telling of the nature of the CIA's interest in this product.

Aside from complete systems, many isolated mechanisms have been proposed to bypass restrictive firewalls that attempt to prevent access to an anonymous communication system. The Tor (Dingledine et al. (2004a)) system, provides a mode that tunnels everything over TCP Port 80 which is often not filtered since it is usually reserved for HTTP (Web) traffic. Anonymizer relies on providing people behind national firewalls (notoriously in China and Iran) with network addresses unknown to the firewall that are not being filtered (Leyden (2006)). This results in an arms race between the providers of fresh addresses, that extensively rely on spam for distribution, and the authorities that seek to detect them and block them. A similar architecture (Köpsell and Hilling (2004)), that relies on volunteers donating their non-blocked network address to help those behind filtering firewalls, has been described and implemented by the JAP⁴ project.

Other studies have looked at censorship and Internet filtering in China (Walton (2001)), and the specific capabilities of the national firewall (Clayton et al. (2006)). It was discovered that it simply sends TCP resets to force communicating parties, with compliant TCP/IP stacks, to drop their connection. Modified clients that ignore such resets were able to carry on communicating. Finally, two studies, one German (Dornseif (2003)) and one British (Clayton), have looked at the effectiveness of Internet Service Providers filtering out web sites that are known to contain child pornography. In their study of the live BT content blocking system *cleanfeed* (Clayton) they discovered that forbidden content could be trivially accessed. Furthermore, the blocking mechanism had the unwanted feature that it could be used as an oracle for interested parties to discover sites with illegal material.

2.3 Type I “Cypherpunk” remailers

Type I remailers, first developed by Eric Hughes and Hal Finney (Parekh (1996)), are nodes that relay electronic mail, after stripping all identifying information and decrypting it with their private key. The first code-base was posted to the cypherpunks mailing list, which gave the remailers their nickname. The encryption was performed using the Pretty Good Privacy (PGP) public key encryption functions. The encoding scheme was also designed to be performed manually, using standard text and email editing tools. Many remailers could be chained together, in order for users not to rely on a single remailer to protect their anonymity.

Reply blocks were supported to allow for anonymous reply addresses. The email address of the user would be encrypted using the remailer’s public key, and inserted in a special header. If a user wished to reply to an anonymous email, the remailer would decrypt it and forward the contents.

The type I remailers offer better resistance to attacks than the simple `anon.penet.fi` relay. No database that links real user identities with pseudonyms is kept. The addressing information required to reply to messages is included in the messages themselves, in an encrypted form.

The encryption used when the messages are forwarded through the network prevents the most trivial passive attacks based on observing the exact bit patterns of incoming messages and linking them to outgoing ones. However it leaks information, such as the size of the messages. Since PGP, beyond compressing the messages, does not make any further attempts to hide their size, it is trivial to follow a message in the network just by observing its length. The reply blocks provided are also a source of insecurity. They can be used many times and an attacker could encode

4. <http://anon.inf.tu-dresden.de/>

an arbitrary number of messages in order to mount an attack to find their destination. Since all replies encoded with the same reply block would contain an identical sequence of bytes, this attack is even easier than the statistical disclosure attacks (Danezis (2003b); Danezis and Serjantov (2004)). The attack can then be repeated to trace any number of hops.

Despite these drawbacks, type I remailers became very popular. This is due to the fact that their reply capabilities allowed the use of Nym Servers. Their reply block feature, that is not present in the later type II Mixmaster software, is both essential to build nym servers, but also insecure even against passive adversaries. This has prompted the design of Mixminion, a type III remailer, that is extremely resistant to traffic analysis, and provides secure single-use reply blocks.

2.4 Crowds

Crowds was developed by Reiter and Rubin (1998) at the AT&T Laboratories. It aims to provide a privacy preserving way of accessing the web, without web sites being able to recognize who is browsing. Each user contacts a central server and receives the list of participants, the “crowd”. A user then relays her web requests by passing it to another randomly selected node in the crowd. Upon receiving a request each node tosses a biased coin and decides if it should relay it further through the crowd or send it to the final recipient. Finally, the reply is sent back to the user via the route established as the request was being forwarded through the crowd.

Crowds is a landmark in anonymity research since its security relies on the adversary not being able to observe the links. Instead, the adversary is assumed to only control a fraction of nodes in each crowd, and the ultimate web server. Although this threat model was initially believed to be unrealistic and fragile, it was later realized that it can be achieved using simple link encryption and padding.

A system which also relies for its security on the inability of the adversary to intercept all communication channels was presented by Katti et al. (2005). They conceptually ‘slice’ each message into many parts, using a secret sharing scheme, and send them out in the network using different channels to an intermediary: if the adversary fails to observe one of the channels they cannot reconstruct the message or the final address to which it is destined. The scheme can be simplified, by considering the secret sharing scheme over a partially secure set of channels, as a primitive encryption mechanism, and the intermediary as a trusted relay that is to decrypt the message and forward it.

Crowds is one of the first papers to address quantitatively how colluding nodes would affect the anonymity provided by the system. It is clear that after the first dishonest node in the path of a request no further anonymity is provided, since the clear text of all requests and replies is available to intermediate nodes. Therefore, given a certain fraction of colluding attacker nodes it is possible to measure the anonymity that will be provided (Diaz et al. (2002)).

Crowds also introduces the concept of *initiator anonymity*: a node that receives a request cannot know if the previous node was the actual requester or was just passing the request along. This property is quite weak and two independent groups have found attacks that identify the originator of requests (Wright et al. (2002); Shmatikov (2002)). They discovered that if a client repeatedly requests a particular resource, they can eventually be linked: The attack relies on the intuition that the true initiator of the repeated request will be the predecessor of a corrupt node more often than

a random node in the crowd. Therefore, for each resource accessed it is sufficient to count how many times each node is seen to be accessing it, and select the node corresponding to the most requests as the most likely initiator. This attack sensitized the anonymity community to the problem of protecting *persistent* relationships instead of simple single message or request exchanges.

Despite the difficulty of securing initiator anonymity, a lot of subsequent systems such as achord (Hazel and Wiley (2002)) and MorphMix (Rennhard and Plattner (2002)), try to achieve it.

2.5 Nym servers

Nym servers (Mazières and Kaashoek (1998)) store an anonymous reply block, and map it to a pseudonymous email address. When a message is received for this address it is not stored, but immediately forwarded anonymously using the reply block to the owner of the pseudonym. In other words, Nym Servers act as a gateway between the world of conventional email and the world of anonymous remailers. Since they hold no identifying information, and are simply using anonymous reply blocks for routing purposes, they do not require users to trust them in order to safeguard their anonymity. Over the years, special software has been developed to support complex operations such as encoding anonymous mail to go through many remailers, and managing Nym Server accounts. Mathewson (2005) presents a contemporary design for a Nym server called *Underhill* that uses the state of the art in remailer technology, and *NymBaron* is its current implementation (Fousse and Reinhard (2006)).

Nym servers are also associated with pseudonymous communications. Since the pseudonymous identity of a user is relatively persistent it is possible to implement reputation systems, or other abuse prevention measures. For example, a nym user might at first only be allowed to send out a small quantity of email messages, that increases over time, as long as abuse reports are not received by the nym server operator. Nym servers and pseudonymous communications offer some hope of combining anonymity and accountability.

At the same time, it is questionable how long the true identity of a pseudonymous user can be hidden. If all messages sent by a user are linked between them by the same pseudonym, one can try to apply author identification techniques to uncover the real identity of the user. Rao and Rohatgi (2000) in their paper entitled "Can Pseudonymity Really Guarantee Privacy?" show that the frequency of function words⁵ in the English language can be used in the long term to identify users. A similar analysis could be performed using the sets of correspondents of each nym, to extract information about the user. Mathewson and Dingledine (2004) have noted that statistical disclosure attacks are very effective at linking pseudonyms with their corresponding users, when those are based on remailer systems.

The shortcomings of remailer based systems have prompted a line of research that looks at alternative techniques to provide receiver anonymity. Techniques from Private Information Retrieval (PIR) have been suggested. PIR is concerned with a family of techniques that allow clients to query a database, without the database or any third party being able to infer which record was retrieved. PIR in the context of receiver anonymity can be implemented either using secure hardware (Asonov and Freytag; Kesdogan et al. (2002b)), or distributed servers as in Sassaman et al.

5. Function words are specific English words used to convey ideas, yet their usage is believed to be independent of the ideas being conveyed. For example: *a, enough, how, if, our, the, ...*

(2005); Kissner et al. (2004). Private Search, a simplified PIR construction (Ostrovsky and III (2005)) was made more efficient by Bethencourt et al. (2006); Danezis and Diaz (2006), and has the potential to be used in efficient receiver anonymity systems.

Interestingly, Ishai et al. (2006) also show that given a strong anonymous channel, one can construct an efficient Private Information Retrieval system.

3. Mix systems

The type I remailer, presented in section 2.3, is the insecure version of a whole body of research that we shall call mix systems and mix networks. This section presents secure constructions based on these ideas.

3.1 Chaum's original mix

The first, and most influential, paper in the field of anonymous communications was presented in Chaum (1981). Chaum introduced the concept of a "mix" node that hides the correspondences between its input messages and its output messages in a cryptographically strong way.

The work was done in the late seventies, when RSA public key encryption was relatively new. For this reason the paper might surprise today's reader by its use of raw RSA, the direct application of modular exponentiation for encryption and decryption, along with an ad-hoc randomisation scheme. Nonces are appended to the plaintext before encryption in order to make two different encryptions output different ciphertext.

The principal idea is that messages to be anonymized are relayed through a node, called a mix. The mix has a well-known RSA public key, and messages are divided into blocks and encrypted using this key. The first few blocks are conceptually the "header" of the message, and contain the address of the next mix. Upon receiving a message, a mix decrypts all the blocks, strips out the first block that contains the address of the recipient, and appends a block of random bits (the junk) at the end of the message. The length of the junk is chosen to make messages size invariant. The most important property that the decryption and the padding aim to achieve is *bitwise unlinkability*. An observer, or an active attacker, should not be able to find the link between the bit pattern of the encoded messages arriving at the mix and the decoded messages departing from the mix. The usage of the word encoded and decoded instead of encrypted and decrypted serves to highlight that the former operations are only used to achieve unlinkability, and not confidentiality, as may be understood to be the aim of encryption. Indeed, modifying RSA or any other encryption and decryption functions to provide unlinkability against passive or active attackers is a problem studied in depth in the context of the design of Mixminion.

Pfitzmann and Pfitzmann (1990) show that Chaum's scheme does not provide the necessary unlinkability properties. The RSA mathematical structure can be subject to active attacks that leak enough information during decryption to link ciphertexts with their respective plaintexts. Further *tagging attacks* are possible, since the encrypted blocks, using RSA are not in any way dependent on each other, and blocks can be duplicated or simply substituted by known ciphertexts. The output message would then contain two blocks with the same plaintext or a block with a known plaintext, respectively. Once again, the use of RSA in the context of a hybrid cryptosystem, in

which only the keys are encrypted using the public key operations, and the body of the message using a symmetric cipher were not very well studied at the time.

A further weakness of Chaum's scheme is its direct application of RSA decryption, which is also used as a signature primitive. An active attacker could substitute a block to be signed in the message and obtain a signature on it. Even if the signature has to have a special form, such as padding, that could be detected, a blinding technique could be used to hide this structure from the mix. It would be unfair to blame this shortcoming on Chaum, since he himself invented RSA blind signatures only a few years later (Chaum (1983)).

The second function of a mix is to actually mix together many messages, to make it difficult for an adversary to follow messages through it, on a first-in, first-out basis. Therefore a mix batches a certain number of messages together, decodes them as a batch, reorders them in lexicographic order and then sends them all out. Conceptually, while bitwise unlinkability makes sure that the contents of the messages do not allow them to be traced, mixing makes sure that the output order of the messages does not leak any linking information.

In order to make the task of the attacker even more difficult, dummy messages are proposed. Dummy messages are generated either by the original senders of messages or by mixes themselves. As far as the attacker is concerned, they are indistinguishable in length and content to normal messages, which increases the difficulty in tracing the genuine messages. We will call the actual mixing strategy, namely the batching and the number of dummy messages included in the inputs or outputs, the *dynamic aspects* of mixing.

Chaum notes that relying on just one mix would not be resilient against subverted nodes, so the function of mixing should be distributed. Many mixes can be chained to make sure that even if just one of them remains honest some anonymity would be provided. The first way proposed to chain mixes is the *cascade*. Each message goes through all the mixes in the network, in a specific order. The second way proposed to chain mixes is by arranging them in a fully connected *network*, and allowing users to pick arbitrary routes through the network. Berthold, Pfitzmann, and Standtke argue in Berthold et al. (2000b) that mix networks do not offer some properties that cascades offer. They illustrate a number of attacks to show that if only one mix is honest in the network the anonymity of the messages going through it can be compromised. These attacks rely on compromised mixes that exploit the knowledge of their position in the chain; or multiple messages using the same sequence of mixes through the network.

Along with the ability for a sender to send messages anonymously to a receiver, Chaum presents a scheme by which one can receive messages anonymously. A user that wishes to receive anonymous email constructs an *anonymous return address*, using the same encoding as the header of the normal messages. She creates blocks containing a path back to herself, and recursively encrypts the blocks using the keys of the intermediate mixes. The user can then include a return address in the body of a message sent anonymously. The receiver simply includes the return address as the header of his own message and sends it through the network. The message is routed through the network as if it was a normal message.

The reply scheme proposed has an important feature. It makes replies in the network indistinguishable from normal messages. In order to securely achieve this, it is important that both the encoding and the decoding operation provide bitwise unlinkability between inputs and outputs.

This is necessary, because replies are in fact encoded when processed by the mix. The resulting message, after it has been processed by all the mixes in the chain specified by the return address, is then decoded with the keys distributed to the mixes in the chain. Both the requirement for decryption to be as secure as encryption, and for the final mix to know the decryption keys to recover the message, means that raw RSA cannot be used. Therefore, a hybrid scheme is proposed that simply encrypts a symmetric key in the header along with the address of the next mix in the chain, that can be used to encrypt or decrypt the message. Since the keys are encoded in the return address by the user, they can be remembered by the creator of the reply block and used to decrypt the messages that are routed using them. Return addresses were also discussed in the Babel system (citebabel) and implemented in the cypherpunk type I remailers. Unfortunately, other deployed systems like Mixmaster did not support them at all.

Chaum's suggestion that a receipt system should be in place to make sure that each mix processes correctly messages, has become a branch of anonymity research in itself, namely mix systems with verifiable properties. We will give an overview of these systems in section 3.9. A system was also proposed to support pseudonymous identities that was partly implemented as the Nym Server described in section 2.3.

3.2 ISDN mixes, Real Time mixes and Web mixes

In 1991, Pfizmann et al. (1991) designed a system to anonymize ISDN telephone conversations. This design could be considered practical, from an engineering point of view, since it met the requirements and constraints of the ISDN network. Later the design was generalized to provide a framework for real-time, low-latency, mixed communications in (Jerichow et al. (1998)). Finally, many of the design ideas from both ISDN and Real Time mixes were adapted for anonymous web browsing and called Web Mixes (Berthold et al. (2000a)). Part of the design has been implemented as a web anonymizing proxy, JAP⁶. All three designs were the product of what could be informally called the Dresden anonymity community (although early research started in Karlsruhe), and the main ideas on which these systems are based are better illustrated by presenting them together.

A major trend in all three papers is the willingness to secure anonymous communication, even in the presence of a very powerful adversary. It is assumed that this adversary would be able to observe all communications on the network (global passive), modify the communications on the links by delaying, injecting or deleting messages, and control all but one of the mixes. While other designs, such as Mixmaster and Babel (that will be presented next), opted for a free route network topology, ISDN, Real Time and Web mixes always use cascades of mixes, making sure that each message is processed by all mixes in the same order. This removes the need for routing information to be passed along with the messages, and also protects the system from a whole set of intersection attacks presented in Berthold et al. (2000b). The debate between the pros and cons of cascade topologies has continued throughout the years, with debates (such as Díaz et al. (2004a)) as well as work exploring the advantages of different topologies (Danezis (2003a); Dingledine et al. (2004c)).

The designs try never to compromise on security, and attempt to be efficient. For this reason, they make use of techniques that provide bitwise unlinkability with very small bandwidth overheads

6. <http://anon.inf.tu-dresden.de/>

and few asymmetric cryptographic operations. *Hybrid encryption with minimal length* encrypts the header, and as much as possible of the plaintext in the asymmetrically encrypted part of the message. A stream cipher is then used to encrypt the rest of the message. This must be performed for each mix that relays the message.

Furthermore, it is understood that some protection has to be provided against active tagging attacks on the asymmetrically encrypted header. A block cipher with a globally known key is used to transform the plaintext before any encryption operation. This technique allows the hybrid encryption of long messages with very little overhead. It is interesting to notice that while the header is protected against tagging attacks, by using a known random permutation, there is no discussion about protecting the rest of the message encrypted using the stream cipher. Attacks in depth could be used, by which a partially known part of the message is XORed with some known text, in order to tag the message in a way that is recognizable when the message is decrypted. As we will see Mixmaster protects against this using a hash, while Mixminion makes sure that if modified, the tagged decoded message will contain no useful information for the attacker.

From the point of view of the dynamic aspects of mixing, ISDN, Real Time and Web mixes also introduce some novel techniques. First the route setup messages are separated from the actual data traveling in the network. In ISDN mixes, the signaling channel is used to transmit the onion encoded message that contains the session keys for each intermediary mix. Each mix then recognizes the messages belonging to the same stream, and uses the session key to prime the stream cipher and decode the messages. It is important to stress that both "data" messages and "route setup" messages are mixed with other similar messages. It was recognized that all observable aspects of the system such as route setup and end, have to be mixed.

In order to provide anonymity for both the initiator and the receiver of a call, rendezvous points were defined. An initiator could use an anonymous label attached to an ISDN switch in order to be anonymously connected with the actual receiver. This service is perhaps the circuit equivalent of a Nym server that can be used by message-based systems. It was also recognized that special cases, such as connection establishment, disconnection and busy lines could be used by an active attacker to gain information about the communicating party. Therefore a scheme of *time slice* channels was established to synchronize such events, making them unobservable to an adversary. Call establishment, as well as call ending have to happen at particular times, and are mixed with, hopefully many, other such events. In order to create the illusion that such events happen at particular times, real or cover traffic should be sent by the users' phones through the cascade for the full duration of the time slice. An even more expensive scheme requires users to send cover traffic through the cascade back to themselves all the time. This would make call initiation, call tear-down and even the line status unobservable. While it might be possible to justify such a scheme for ISDN networks where the lines between the local exchange and the users are not shared with any other parties, it is a very expensive strategy to implement over the Internet in the case of Web mixes.

Overall, the importance of this body of work is the careful extension of mixes to a setting of high-volume streams of data. The extension was done with careful consideration for preserving the security features in the original idea, such as the unlinkability of inputs and outputs and mixing all the relevant information. Unfortunately, while the ideas are practical in the context of telecommunication networks, where the mix network is intimately linked with the infrastructure, they are less so

for widely deployed modern IP networks. The idea that constant traffic can be present on the lines, and that the anonymity can be guaranteed, but be relatively low, is not practical in such contexts. Onion routing, presented in section 3.6, provides a more flexible approach that can be used as an overlay network, but it is at the same time open to more attacks. These techniques may become increasingly relevant if fixed rate traffic, such as streaming data and VoIP, require anonymization.

3.3 Babel and Mixmaster

Babel (Gülcü and Tsudik (1996)) and Mixmaster (Möller et al. (2003)) were designed in the mid-nineties, and the latter has become the most widely deployed remailer. They both follow a message-based approach, namely they support sending single messages, usually email, through a fully connected mix network.

Babel offers sender anonymity, called the “forward path” and receiver anonymity, through replies traveling over the “return path”. The forward part is constructed by the sender of an anonymous message by wrapping a message in layers of encryption. The message can also include a return address to be used to route the replies. The system supports bidirectional anonymity by allowing messages to use a forward path, to protect the anonymity of the sender, and for the second half of the journey they are routed by the return address so as to hide the identity of the receiver.

While the security of the forward path is as good as in the secured original mix network proposals, the security of the return path is slightly weaker. The integrity of the message cannot be protected, thereby allowing tagging attacks, since no information in the reply address, which is effectively the only information available to intermediate nodes, can contain the hash of the message body. The reason for this is that the message is only known to the person replying using the return address. This dichotomy will guide the design of Mixminion, since not protecting the integrity of the message could open a system to trivial tagging attacks. Babel reply addresses and messages can also be used more than once, while messages in the forward path contain a unique identifier and a timestamp that makes detecting and discarding duplicate messages efficient.

Babel also proposes a system of intermix detours. Messages to be mixed could be “repackaged” by intermediary mixes, and sent along a random route through the network. It is worth observing that even the sender of the messages, who knows all the symmetric encryption keys used to encode and decode the message, cannot recognize it in the network when this is done.

Mixmaster has been an evolving system since 1995 (Möller et al. (2003)). It is the most widely deployed and used remailer system.

Mixmaster supports only sender anonymity, or in the terminology used by Babel, only anonymizes the forward path. Messages are made bitwise unlinkable by hybrid RSA and EDE 3DES encryption, while the message size is kept constant by appending random noise at the end of the message. In version two, the integrity of the RSA encrypted header is protected by a hash, making tagging attacks on the header impossible. In version three, the noise to be appended is generated using a secret shared between the remailer, and the sender of the message, included in the header. Since the noise is predictable to the sender, it is possible to include in the header a hash of the whole message therefore protecting the integrity of the header and body of the message. This trick makes replies impossible to construct: since the body of the message would not be known to the creator of the anonymous address block, it is not possible to compute in the hash.

Beyond the security features, Mixmaster provides quite a few usability features. It allows large messages to be divided in smaller chunks and sent independently through the network. If all the parts end up at a common mix, then reconstruction happens transparently in the network. So large emails can be sent to users without requiring special software. Recognizing that building robust remailer networks could be difficult (and indeed the first versions of the Mixmaster server software were notoriously unreliable) it also allowed messages to be sent multiple times, using different paths. It is worth noting that no analysis of the impact of these features on anonymity has ever been performed.

Mixmaster also realizes that reputation attacks, by users abusing the remailer network, could discredit the system. For this reason messages are clearly labeled as coming from a remailer and black lists are kept up-to-date with email addresses that do not wish to receive anonymous email. While not filtering out any content, for example not preventing death threats being transmitted, at least these mechanisms are useful to make the network less attractive to email spammers.

3.4 Mixminion: the Type III Remailer

Mixminion (Danezis et al. (2003)) is the state of the art anonymous remailer. It allows for a fixed size message, of about 28 kbytes, to be anonymously transported over a set of remailers, with high latency. Mixminion supports sender anonymity, receiver anonymity via single-use reply blocks (SURBs), and bi-directional anonymity by composing the two mechanisms. This is achieved by mixing the message through a string of intermediate Mixminion remailers. These intermediate remailers do not know their position on the path of the message, or the total length of the path (avoiding partitioning attacks as described by Berthold et al. (2000b)). Intermediate remailers cannot distinguish between messages that benefit from sender anonymity and anonymous replies.

Mixminion's first key contribution concerns the cryptographic packet format. Transported messages are divided into two main headers and a body. Each main header is further divided into sub-headers encrypted under the public keys of intermediary mixes. The main objective of the cryptographic transforms is to protect messages from tagging (Pfizmann and Pfizmann (1990); Pfizmann (1994)): an active adversary or corrupt node may modify a message, in the hope that they will be able to detect the modification after the message has been mixed. This would allow an adversary to trace the message and compromise anonymity. Mixmaster solves this problem by including an integrity check in the header read by each intermediary: if tampering is detected the message is dropped at the first honest mix. Mixminion cannot use a similar mechanism, because of the need to support indistinguishable routing of anonymous replies. Instead, it relies on an all-or-nothing encryption of the second header and the body of the message, which is very fragile. Tampering cryptographically results in the address of the final receiver and the message being destroyed. The cryptographic format was designed to be well understood, and as a result it is quite conservative and inefficient.

The Minx packet format aims to provide the same properties as Mixminion at a lower computational cost and overhead (Danezis and Laurie (2004)). It relies on a single pass of encryption in IGE mode, that propagates ciphertext errors forward. As a result, modifying the message results again in all information about the final receiver and the message being destroyed. Since all messages look random, no partial information is ever leaked through tampering.

Mixminion uses a TCP based transport mechanism, that can accommodate link padding. Messages are transferred between remailers using a TLS protected tunnel, with an Ephemeral Diffie-Hellman based key exchange to provide forward security. This renders any material gathered by a passive adversary useless, since it cannot be presented to a remailer operator for decryption after the ephemeral keys are deleted. It also detects active adversaries that try to corrupt data traveling on the network. Therefore an adversary must be running malicious nodes to attack the network.

Two proposals have been put forward to strengthen the forward security and compulsion resistance of Mixminion mixing. The first, in Danezis (2002), assumes that any communication leaves a trail of keys on intermediate mixes that can be used to decrypt future communications. Once a key is used, it is deleted or updated using a one-way function. Since subsequent messages may be dependent on previous messages for decoding, a mix that honestly deletes keys cannot decrypt intercepted messages upon request. Furthermore, an adversary needs to intercept and request the decryption of many messages in order to retrieve the key material necessary to decode any particular target message. The second technique (Danezis and Clulow (2005)) relies on the fact that the genuine receiver of an anonymous reply can pretend to be a relay, and pass the message to another pre-determined node. This assumes a peer-to-peer remailer system, and may be an incentive to run a Mixminion server.

The implementation of Mixminion brought to the surface many practical questions. Since the transport of Mixminion messages is unreliable, it is important to implement mechanisms for re-transmissions and forward error correction. Such mechanisms are not trivial to implement and may lead to traffic analysis attacks. In order to be secure, all clients must know the full network of remailers. This has proved to be a scalability problem, and a distributed directory service had to be specified in order to distribute this information. Robust mechanisms for vetting nodes, and ensuring their honest behavior are still elusive. Practical deployment and integration into familiar clients has also been a challenge.

3.5 Foiling flooding attacks: Tickets, Stop-and-Go, Red-Green-Black-Mixes and mixing strategies

As we saw above, Babel and Mixmaster implement a traditional mix network model. They also both extend the original idea of mixing batches of messages together to feeding back messages in a pool, in the case of Mixmaster, or to delaying a fraction of messages an additional round, in the case of Babel. Such mix strategies, along with others, are susceptible to an $(n - 1)$ attack, in which the adversary sends one message to be traced to an empty mix, followed by adversary messages. When the mix flushes, the only message that cannot be recognized is the one to be traced, which compromises anonymity.

In Serjantov et al. (2002) the attack is explained as acting in two phases: a trickle and a flood. In the first instance the adversary tries to flush all genuine messages from the mix, before injecting the target message and flooding it. A more rigorous analysis of how many different mix strategies are susceptible to such an attack is provided by O'Connor (2005).

The simple strategy proposed to counter such attacks is admission control, through authentication and ticketing systems (Berthold and Langos (2002)). If each user is properly authenticated when sending a message, flooding can be detected and foiled. This solution is not fully satisfactory

though, since corrupt mixes may also inject messages. Having to authenticate may also reduce the perception of security offered by the system.

Diaz and Serjantov (2003) introduced a model for representing the mixing strategies of pool mixes. This model allows for easy computation of the anonymity provided by the mixing strategy towards active and passive adversaries. It was noted that $n - 1$ attacks on pool mixes were favored by the deterministic dependency of the number of messages forwarded in any round and the number of messages kept in the pool for future rounds. The adversary could use this knowledge to optimize his efforts in terms of time and number of messages generated and have 100% certainty on the detection of the target at the output of the mix. In order to increase the effort and the uncertainty of the attacker, they propose randomizing the number of messages forwarded, as a binomial distribution of the number of messages contained in the pool. The randomization can be done almost for free: at the time of forwarding, the mix, instead of choosing a fixed number of random messages from the pool, flips a biased coin for each message.

The first effect of the randomization is that the attacker succeeds only probabilistically, and the effort of the attacker increases as he tries to increase his probability of success. Díaz and Preneel (2004) analyzes the robustness of various combinations of mixing and dummy generation strategies towards $n - 1$ attacks. It is shown that the combination of binomial mixing and randomized dummy generation strategies sets a lower bound on the anonymity of the target message. The adversary is able to significantly reduce the anonymity set of the message but he does not uniquely identify the message at the output of the mix. The protection offered to the message is proportional to the amount of dummies generated by the mix. A detailed analysis of the result and costs of deploying $n - 1$ attacks is presented in Serjantov (2004).

Stop-and-Go mixes (Kesdogan et al. (1998)) (sg-mix) present a mixing strategy, that is not based on batches but delays. It aims at minimizing the potential for $(n - 1)$ attacks. Each packet to be processed by an sg-mix contains a delay and a time window. The delay is chosen according to an exponential distribution by the original sender, and the time windows can be calculated given all the delays. Each sg-mix receiving a message, checks that it has been received within the time window, delays the message for the specified amount of time, and then forwards it to the next mix or final recipient. If the message was received outside the specified time window it is discarded. This security feature was, however, not implemented in the practical implementation of sg-mixes Reliable, which inter-operated with the pool mixes of the Mixmaster network. A practical comparison on the anonymity provided by both the pool and sg nodes of the Mixmaster network towards passive adversaries is presented in Díaz et al. (2004b). This paper shows that, even in very low traffic conditions, the pool nodes provide a high anonymity set to the messages they route at the expense of longer delays. The Reliable node, which does not adapt the delay to the traffic load, provides no anonymity in extreme cases.

A very important feature of sg-mixes is the mathematical analysis of the anonymity they provide. Assuming that the messages arriving to the mix follow a Poisson distribution, it is observed that each mix can be modeled as a $M/M/\infty$ queue, and a number of messages waiting inside it follow the Poisson distribution. The delays can therefore be adjusted to provide the necessary anonymity set size.

The time window is used in order to detect and prevent $(n - 1)$ attacks. It is observed that an adversary needs to flush the sg-mix of all messages, then let the message to be traced through

and observe where it goes. This requires the attacker to hold the target message for a certain time, necessary for the mix to send out all the messages it contains and become empty. The average time that the message needs to be delayed can be estimated, and the appropriate time window can be specified to make such a delayed message be rejected by the mix.

A different solution to the $(n - 1)$ attack, the *rgb-mix* by Danezis and Sassaman (2003), is based on a controlled level of cover traffic. In their scheme, each mix in the network sends 'red' heartbeat messages back to itself through the mix network. If at some point such messages stop arriving it may mean that the mix is subject to the first phase of the $(n - 1)$ attack. The mix then responds by injecting 'green' cover traffic to confuse the adversary. The key property that makes this scheme secure is the inability of the adversary to tell apart genuine messages, to be blocked, and heartbeat messages that need to be let through for the mix not to introduce additional cover traffic. Under normal operating conditions the traffic overhead of this scheme is minimal, since additional traffic is only introduced as a response to attack.

3.6 Onion routing

Onion routing (Goldschlag et al. (1996); Reed et al. (1998); Goldschlag et al. (1999); Syverson et al. (2000)) is the equivalent of mix networks, but in the context of circuit-based routing. Instead of routing each anonymous packet separately, the first message opens a circuit through the network, by labeling a route. Each message having a particular label is then routed on this predetermined path. Finally, a message can be sent to close the path. Often, we refer to the information traveling in each of these labeled circuits as an anonymous stream.

The objective of onion routing is to make traffic analysis harder for an adversary. It aims first at protecting the unlinkability of two participants who know each other from third parties, and secondly, at protecting the identities of the two communicating parties from each other. Furthermore, onion routing notes that ISDN mixes are not easily implementable over the Internet, and aims to distribute the anonymous network and adapt it to run on top of TCP/IP.

The first message sent through the network is encrypted in layers, that can only be decrypted by a chain of onion routers using their respective private keys. This first message contains key material shared between the original sender and the routers, as well as labels and addressing information about the next node. As with Chaum's mixes, care is taken to provide bitwise unlinkability, so that the path that the first message takes is not trivial to follow just by observing the bit patterns of messages. Loose routing is also proposed, according to which routers relay streams through paths that are not directly specified in the original path opening message. The hope was that such a scheme would increase the anonymity provided.

Data traveling in an established circuit is encrypted using the symmetric keys distributed to the routers. Labels are used to indicate which circuit each packet belongs to. Different labels are used on different links, to ensure bitwise unlinkability, and the labels on the links are encrypted using a secret shared key between pairs of onion routers. This prevents a passive observer from knowing which packets belong to the same anonymous stream, but does not hide this information from a subverted onion router.

Onion routing admits to being susceptible to a range of attacks. It has become clear that in the absence of heavy amounts of cover traffic, patterns of traffic are present that could allow an

attacker to follow a stream in the network and identify the communicating parties. Such attacks have been called timing attacks. While they are often cited in the literature (Raymond (2000)), details of how they work and how effective they are have only been presented relatively recently.

Unlike ISDN mixes, onion routing does not perform mixing on the requests for opening or closing channels. While it might be plausible that enough data would be available to mix properly, it is very unlikely that the anonymity of circuit-setup messages can be maintained. Therefore, an attacker could follow such messages and compromise the anonymity of the correspondents. Furthermore, very little mixing is done in the system generally, because of the real-time performance that is assumed to be needed. Onion routing aims at providing anonymous web browsing, and therefore would become too slow if proper mixing was to be implemented. Therefore, a mixing strategy that is very close to first-in first-out for each stream is implemented. This provides only minimal mixing, and as a result a lot of attacks against onion routing focus on its weak dynamic features.

In order to make deployment easier, it was recognized that some onion routers might wish to only serve particular clients. The concept of *exit policies* was developed to encapsulate this, allowing routers to advertise which section of the network they were configured to serve. Onion routers are also free to peer with only a subset of other routers, with which they maintain long standing connections.

Zero Knowledge, a Canadian company, designed the Freedom network that follows quite closely the architecture of onion routing. The principal architect of the network was Ian Goldberg (Goldberg (2000)) who published with others a series of technical papers describing the system at various levels of detail (Boucher et al. (2000); Back et al. (2001a)).

3.7 Tor: The new generation onion router

The onion routing project was revived in 2004, with the design and implementation of a second generation onion router called Tor (Dingledine et al. (2004b)). Tor relays arbitrary TCP streams over a network of relays, and is particularly well tuned to work for web traffic, with the help of the Privoxy⁷ content sanitizer.

Tor uses a traditional network architecture: a list of volunteer servers is downloaded from a directory service. Then, clients can create paths by choosing three random nodes, over which their communication is relayed. Instead of an 'onion' being sent to distribute the cryptographic material, Tor uses an iterative mechanism. The client connects to the first node, then it request this node to connect to the next one. The bi-directional channel is used at each stage to perform an authenticated Diffie-Hellman key exchange. This guarantees forward secrecy and compulsion resistance: only short term encryption keys are ever needed. This mechanism was first described in Cebolla (Brown (2002)), and is not covered by the Onion Routing patent (Reed et al. (2001)).

One notable difference between Tor and previous attempts at anonymizing streams of traffic, is that it does not claim to offer security against even passive global observers. A set of traffic analysis techniques by Danezis (2004); Levine et al. (2004); Serjantov and Sewell (2003); Zhu et al. (2004); Wang et al. (2005); Zhu and Bettati (2005) have been developed throughout the years to trace streams of continuous traffic traveling in a low latency network. A separate but

7. <http://www.privoxy.org/>

related thread of research has been developed in the intrusion detection community, that tries to uncover machines used as stepping stones for attack (Wang and Reeves (2003); Blum et al. (2004)). These attacks have been shown difficult to foil, unless the latency of the traffic is high, or a lot of cover traffic is injected – both of which are very expensive. Tor instead opts for getting security though being highly usable and cheap to operate (Back et al. (2001b); Dingledine and Mathewson (2005)). As a result, an adversary who can observe a stream at two different points, can trivially realize it is the same traffic.

Its vulnerability against passive adversaries has made Tor fragile against previously unexplored attacks. First, a realistic assessment of the probability a single party can observe multiple points on the path is necessary. It turns out that the topology of the Internet is such that many, seemingly unrelated networks, are interconnected through hubs, or long distance links that can be observed cheaply by a single ISP entity (Feamster and Dingledine (2004)). A second possible path for attack in Murdoch and Danezis (2005) uses indirect network measurements to perform traffic analysis, and does away with the assumption that a passive adversary needs local access to the communication to perform traffic analysis. An attacker relays traffic over all routers, and measures their latency: this latency is affected by the other streams transported over the router. Long term correlations between known signals injected by a malicious server and the measurements are possible. This allows an adversary to trace a connection up to the first router used to anonymize it.

Tor also provides a mechanism for 'hidden servers'. A hidden server opens an anonymous connection and uses it to advertise a contact point. A client that wants to contact the server, goes to the contact point and negotiates a separate anonymous channel used to relay the actual communication. An attack against this early architecture was demonstrated by Øverlier and Syverson (2006). The intuition behind this attack is that an adversary can open multiple connections to the hidden server, sequentially or in parallel, and control the flow of traffic towards the server. The adversary needs to control one corrupt router, and wait until for one of the connections his router is chosen by the server as the first node for the fresh anonymous path. Then the adversary effectively controls two nodes on the anonymous path, one of which is next to the real server – and the anonymity provided to the server is completely compromised. The idea of consistently using a more trusted 'valet' router as the first node into the Tor network was proposed as a countermeasure against this attack by verlier and Syverson (2006).

3.8 Peer-to-peer mix networks

In Chaum's original work it is assumed that if each participant in the mix network also acts as a mix for others, this would improve the overall security of the network. Recent interest in peer-to-peer networking has influenced some researchers to further examine such networks with large, but transient, numbers of mixes.

Freedman and Morris (2002) designed *Tarzan*, a peer-to-peer network in which every node is a mix. A peer initiating the transport of a stream through the network would create an encrypted tunnel to another node, and ask that node to connect the stream to another peer. By repeating this process a few times, it is possible to have an onion encrypted connection, relayed through a sequence of intermediate nodes.

An interesting feature of Tarzan is that the network topology is somewhat restricted. Each node maintains persistent connections with a small set of other nodes, forming a structure called a *mimics*. Routes of anonymous messages are selected in such a way that they will go through mimics and between mimics in order to avoid links with insufficient traffic. A weakness of the mimics scheme is that the selection of neighboring nodes is done on the basis of a network identifier or address which, unfortunately, is easy to spoof in real-world networks.

The original Tarzan design only required each node to know a random subset of other nodes in the network. This is clearly desirable due to the very dynamic nature of peer-to-peer networks, and the volatility of nodes. On the other hand, Danezis and Clayton (2006) found some attacks against this strategy in a preliminary version of Tarzan (Freedman et al. (2002)). The attack relies on the fact that the network is very large, and nodes have a high churn rate. As a result any single node only knows a small subset of other nodes. An adversary node, included on the anonymous path, can tell that the originator of the connection knew three nodes: the corrupt node itself, its successor and predecessor. It turns out that those three nodes identify uniquely the originator of the stream with very high probability. The final version of Tarzan requires each node to know all others in order to fix this attack, which is clearly less practical.

Rennhard and Plattner (2002) introduced *MorphMix*, which shares a very similar architecture and threat model with Tarzan. A crucial difference is that the route through the network is not specified by the source but chosen by intermediate nodes, observed by witnesses specified and trusted by the user. While the attack by Danezis and Clayton does not apply to route selection, variants might apply to the choice of witness nodes.

MorphMix realises that leaving the intermediate nodes to choose the route through the network might lead to *route capture* or, in other words, the first subverted mix on the path choosing only other subverted mixes. For this reason, MorphMix includes a *collusion detection* mechanism that monitors for any cliques in the selection of nodes in the path. This prevents subverted nodes from routinely running attacks on the network but does not provide security in every case. Tabriz and Borisov (2006) presented an attack on the collusion resistance mechanism of MorphMix.

3.9 Robust & verifiable mix constructions

Chaum's original mix network design included a system of signed receipts to assure senders that their messages have been properly processed by the network. A whole body of research was inspired by this property and has attempted to create mix systems which are robust against subverted servers denying service, and that could offer a proof of their correct functioning alongside the mixing. Such systems have been closely associated with voting, where both universal verifiability of vote delivery and privacy are of great importance.

Most of the proposed schemes use the idea of a mix cascade. For this reason, no information is usually communicated between the sender of a message and intermediate mixes in the cascade. It is assumed that routing information is not necessary, since mixes process messages in a fixed order. The first scheme to take advantage of this was the *efficient anonymous channel and all/nothing election scheme* proposed by Park et al. (1993). In this system, messages are an El Gamal ciphertext of fixed length, independently of the number of mixes they go through. Furthermore, the scheme uses a *cut and choose* strategy, which makes it all-or-nothing, meaning

that if any of the ciphertexts is removed, then no result at all is output. This property assures that partial results do not affect a re-election. Birgit Pfitzmann found two attacks against this proposal (Pfitzmann (1994)). The first attack is very similar to Pfitzmann and Pfitzmann (1990), and makes use of characteristics that are invariant at the different stages of mixing because of the El Gamal cryptosystem. An active attack is also found, where the input El Gamal ciphertext is blinded, by being raised to a power, which results in the final output also being raised to this power. This is a chosen ciphertext attack with which a lot of systems will struggle, and eventually fail to eliminate. Birgit Pfitzmann also notes that the threat model assumed is somehow weaker than the one proposed by Chaum. A dishonest sender is capable of disrupting the whole network, which is worse than a single mix, as it is the case in Chaum's paper. Birgit did not propose any practical countermeasures to these attacks, since any straightforward fix would compromise some of the interesting features of the system.

In parallel with Birgit Pfitzmann's work, Kilian and Sako (1995) proposed a *receipt-free mix-type voting scheme*. They attempt to add universal verifiability to Park et al. (1993), which means that all senders will be able to verify that all votes were taken into account, not simply their own. They also highlight that many verifiable mix schemes provide at the end of mixing a receipt, that could be used to sell or coerce one's vote, and attempt to make their system *receipt-free*. They do this by forcing each mix to commit to their inputs and outputs, and prove in zero knowledge that they performed the decryption and shuffle correctly. Unfortunately, Michels and Horster (1996) show that the scheme is not receipt-free if a sender collaborates with a mix, and that the active attacks based on blinding proposed by Birgit Pfitzmann could be used to link inputs to outputs.

In order to avoid disruption of the system if a subset of mixes is subverted, Ogata et al. (1997) proposed a *fault tolerant anonymous channel*. This uses a threshold cryptosystem to make sure that a majority of mixes can decode messages even if a minority does not collaborate. Two systems are proposed, one based on El Gamal and the other based on the r^{th} residue problem. A zero knowledge proof of correct shuffling is also proposed for the r^{th} residue problem.

In 1998, Abe (1998) presented a mix system that provided universal verifiability and was efficient, in the sense that the verification work was independent from the number of mix servers. This scheme shows an attack on Kilian and Sako (1995), that uses the side information output for the verification to break the privacy of the system. It then presents a mix system that works in two phases, El Gamal re-encryption and then threshold decryption. The first phase is proved to be correct before the second can proceed, and then a proof of correct decryption is output at the end of the second stage.

The systems that provide universal verifiability based on proofs of permutations, and zero knowledge proofs are computationally very expensive. Jakobsson (1998) designs the Practical Mix, and tries to reduce the number of expensive operations. In order to prove the correctness of the shuffle, novel techniques called *repetition robustness* and *blinded destructive robustness* are introduced. The network works in two phases: first, the ciphertexts are El Gamal blinded, and then, the list of inputs is replicated. Each of the replicated lists is decoded by all mixes, which results in lists of blinded plaintexts. The resulting lists are sorted and compared. If all elements are present in all lists then no mix has tampered with the system and the unblinding and further mixing can proceed. Otherwise, the sub-protocol for cheater detection is run. While being very efficient, the Practical Mix has not proved to be very secure, as shown by Desmedt and Kurosawa (2000). They

show that one subverted mix in the practical mix can change ciphertexts, and still not be detected. They then introduce a new mix design, in which verification is performed by subsets of mixes. The subsets are generated in such a way that at least one is guaranteed not to contain any subverted mixes.

In an attempt to further reduce the cost of mixing, Jakobsson (1999) introduced the Flash Mix, that uses re-encryption instead of blinding to keep the number of exponentiations down. As in the practical mix, mixing operates in many phases, and uses *repetition robustness* to detect tampering. Furthermore, two dummy messages are included in the input, that are de-anonymized after all mixes have committed to their outputs, in order to make sure that attacks such as Desmedt and Kurosawa (2000) do not work. An attack against Flash mixing was found in Mitomo and Kurosawa (2000) and fixed by changing the unblinding protocol.

A breakthrough occurred when Furukawa and Sako (2001) and Neff (2001) proposed efficient general techniques to universally verify the correctness of a shuffle of El Gamal ciphertexts. The first provides proof that the matrix used was a permutation matrix, and the second uses verifiable secret exponent multiplication to improve its efficiency.

Even though the above techniques are more efficient than any other previously known, they are still not efficient enough to scale for elections, with millions of participants. For this reason, Golle et al. (2002) proposed optimistic mixing, a mix that works quickly if there is no attack detected, but provides no result if an error occurs. In this case, it provides a fall back mechanism for a more robust technique such as Neff (2001) to be used. Each mix in the chain outputs a “proof” of permutation, that could be faked by tampering with the ciphertexts. This is detected by making the encryption plaintext-aware. The second decryption, revealing the votes, is only performed if all outputs of mixing are well-formed. A series of attacks were found against this scheme by Wikström (2003b, 2002). The first two attacks are closely related to Pfitzmann (1994) and can break the anonymity of any user. The second attack is related to Desmedt and Kurosawa (2000) and can break the anonymity of all users and compromise the robustness. Finally, attacks based on improperly checking the El Gamal elements are also applicable, and further explored in Wikström (2003a).

A serious drawback of traditional robust mix cascades is that each mix has to wait for the output of the previous mix before processing messages. This means that the latency increases with the number of mixes, and that most of the time mixes perform no computations. Golle and Juels (2004b) present a technique that allows for universally verifiable parallel mixing in four steps of communication and the equivalent of two steps of mixing. Their techniques drastically reduce the latency of mixing, but Borisov (2005) shows that when multiple input messages are known to the adversary, the anonymity provided by this technique is far from optimal.

A hopeful line of research looks at extending robust cascades into general mix networks. These may use non deterministic decoding and routing protocols, possibly implemented using the new universal re-encryption primitive suggested in Golle et al. (2004), and improved for space efficiency by Fairbrother (2004). Sadly, its mathematical properties make universal re-encryption malleable and many systems that use it (Klonowski et al. (2005,); Gomulkiewicz et al. (2004); Klonowski et al. (2004); Lu et al. (2005b,a)) were found to be vulnerable to tagging attacks in Danezis (2006).

Finally, a fundamentally new way of looking at robust mixing is presented in Adida and Wikström (2005): mixing is seen as a computation to be outsourced to a third party. Yet, this third party should gain no information about the actual shuffle. Two protocols that implement such an algorithm are presented, based on Paillier and BGN homomorphic encryption. The third party accepts a set of ciphertexts, and performs in the obfuscated mixing algorithm to produce a re-encrypted and shuffled set of outputs. Despite only public keys being used, neither the third party, nor any observer, can link inputs and outputs.

3.10 Provable security for mix-networks

While most designs for robust mix nets use pure El Gamal encryption, some provide solutions for hybrid schemes. Ohkubo and Abe (2000) present a hybrid mix without ciphertext expansion. Jakobsson and Juels (2001) also present a scheme that is resistant to any minority coalition of servers. Möller (2003) proves the security properties of a mix packet format. Further work in proving packet formats correct was presented in Camenisch and Lysyanskaya (2005). Other packet formats attempt to provide specialized properties: Golle (2004) allows a mix to prove that a particular output was an input to the mix, clearing the operator from any suspicion that they injected the message.

Reputation based schemes have also been used to increase the reliability of mix networks in Dingledine et al. (2001), and mix cascades in Dingledine and Syverson (2002). Both these papers present how *statistics pages* compiled in the Mixmaster system using *pingers* (Palfrader) can be replaced with a more robust system to determine which nodes are reliable and which are not. Users can then choose reliable nodes, or the system can exclude unreliable ones from directories.

An option that universal verifiability to be implemented on generic mix networks, is the randomized partial checking presented in Jakobsson et al. (2002). In this scheme, all mixes commit to their inputs and outputs and then they are required to disclose half of all correspondences. This assures that if a mix is dropping messages it will be quickly detected. Privacy is maintained by pairing mixes, and making sure that the message is still going through enough secret permutations. For safety, it was proposed that mixes are paired, and when one in a pair is required to disclose the correspondences the other keeps it secret, in order to ensure that enough mixing is performed for each message. In Gomulkiewicz et al. (2003) it is shown, using path coupling tools and graph theory, that such caution is not necessary, since messages will mix with high probability after $\log N$ steps, even if correspondences are revealed at random.

A separate line of research attempts to prove the mixing, and hence privacy properties instead of robustness. Such systems are usually expensive, since they rely on extensive amounts of cover traffic to provably ensure that no information about the actual traffic patterns is leaked. Systems in this tradition are Rackoff and Simon (1993); Berman et al. (2004). The latter proves that in a random network of communications, one could embed a very large number of possible sub-networks of a certain butterfly-like form, and show that, at each step, messages are mixed together with high probability. Interestingly, Klonowski and Kutylowski (2005) prove that traditional mix networks mix all input messages after $\log N$ rounds, despite the presence of adversaries that reveal to each other the path of messages.

Some mix strategies are designed on purpose to be fragile. If a single message gets deanonymized through compulsion, then all the messages get deanonymized (Reiter and Wang (2004)), or a secret of the mix operator can easily be inferred (Golle et al. (2006)). This provides operators with incentives to resist compulsion, and is meant to make it disproportionate (in jurisdictions where this is a concern) to request even one message to be traced.

4. Other systems

A number of other anonymous communication systems have been proposed through the years. Chaum (1988) presents the dining cryptographers' network, a multi-party computation that allows a set of participants to have perfect (information theoretic) anonymity. The scheme is very secure but impractical, since it requires a few broadcasts for each message sent and is easy to disrupt for dishonest users. A modification of the protocol in Waidner and Pfitzmann (1989), guarantees availability against disrupting nodes. Herbivore (Goel et al. (2003)) uses DC-nets as part of a two-level anonymity system: users form small cliques that communicate within them using DC-nets. Finally, in Golle and Juels (2004a) asymmetric techniques are described that make DC-nets robust against disruption.

P5, by Sherwood et al. (2002), uses broadcast-trees to achieve anonymity. Buses by Beimel and Dolev (2003) use the metaphor of a bus route that travels over all nodes carrying messages. This is in fact a broadcast, and trade-offs between longer routes and more routes are discussed from an anonymity and latency perspective.

Traffic Analysis Prevention (TAP) systems, attempt to provide third party anonymity, given a collaborating set of senders, receivers and relays. Timmerman (1999) describes adaptive traffic masking techniques, and a security model to achieve *traffic flow confidentiality* (Timmerman (1997)). The information theoretic approach to analysing TAP systems is presented by Newman et al. (2003). They study how much protection is offered overall to the traffic going through a TAP system, by creating a rigorous mathematical model of traffic analysis, rerouting and cover traffic. This builds on their previous work in Venkatraman and Newman-Wolfe (1994). The research group at the Texas A&M University has a long-term interest in traffic analysis prevention of real time traffic (Guan et al. (2001)). Similarly, Jiang et al. (2000) present TAP techniques to protect wireless packet radio traffic.

5. Conclusions

Anonymous communications, despite being first proposed over 25 years ago, has become since 2000 an extremely active field of research. It is also increasingly relevant since systems that are the direct result of this research, like Tor, JAP and Mixminion, are being deployed and used to protect the privacy of thousands of people.

Anonymous communications research has also matured to the point that new systems must imperatively take into account the existing literature and ensure that they are not weak under known attacks and models. The aims of this survey has been to present a road map of the most important systems-concepts, and the key refinements they have been subject to.

As in any mature field new designs will inevitably have to mix and match from elements already present, in older systems, to best match their environment. Designs tailored to peer-to-peer systems or telephony are a prime example of this. Those systems are also a prime example of the care that researcher must exert then mixing and matching ideas: anonymous communications are fragile, and even simple modifications may lead to traffic analysis attacks.

A set of key observations must be in the minds of designers and researchers looking at anonymous communications in the future.

The *concepts of anonymity* in communication networks is a well understood problem. Definitions and metrics that express the anonymity properties of communications are available, and used to evaluate systems. Despite all security efforts, an upper limit on the anonymity that a system can provide is given by black box attacks: no matter how good the anonymity system is, effective attacks can be deployed in the long term by observing the edges of the anonymous communication network. As a result we say that the use of anonymous communication can be secured only tactically (for short periods) and not strategically or in the long term.

Concerning *trust models* the earliest anonymous communication systems relied on one central trusted server. This model has proven weak against compulsion, denial of service, traffic analysis carried by a local eavesdropper, or maliciousness of the central node. Centralized trust models have been abandoned in favor of models where trust is distributed over a set of servers. Trusting a set of unknown, random nodes presents some serious limitations as well, particularly against attackers able to introduce a large number of corrupted nodes in the system (sybil attacks).

Solutions for *high-latency applications* such as email have significantly evolved since the first schemes were proposed. The loose delay requirements allow for the design of secure solutions, providing a reasonably high resistance to attacks and anonymity level.

On the other hand, *low-latency* constraints seriously limit the anonymity that can be provided against powerful adversaries. Currently deployed solutions are vulnerable against attackers who have access to both ends of the communication. In particular, the variability of HTTP traffic makes it hard to conceal the correlation of input and output at the edges of the network using black box attacks.

There has been a link between anonymous communications and *censorship resistance* research, as solutions for one problem have been applied to the other. More research is needed to determine whether anonymity is the best tactic to distribute content that may be censored, or whether it adds cost that may be limiting the distribution even further.

Finally, anonymous communication networks can be subject to a wide range of *attacks*. The most used attacker model is the global attacker (with access to all communication lines, passive or active); and attackers capable of controlling a subset of the nodes in the network (Crowds model). The attacks against which anonymity networks are most vulnerable include traffic analysis, flooding, compulsion, and attacks on the cryptographic protocols (such as tagging attacks).

Know-how in attacking anonymous communication grows at the same, or even faster rate, as our ability to design secure systems. As more systems are deployed further attacks are uncovered, making use of the implementation environment and the actual usage of the anonymity system.

Anonymity design has proved to be a non trivial problem, but so far we have only scraped the surface of anonymity engineering, deployment and operations problems.

References

- Masayuki Abe. Universally verifiable MIX with verification work independent of the number of MIX servers. In K. Nyberg, editor, *Advances in Cryptology (Eurocrypt'98)*, volume 1403 of LNCS, pages 437–447, Helsinki, Finland, 31 May–4 June 1998. Springer-Verlag.
- Ben Adida and Douglas Wikström. Obfuscated ciphertext mixing. *Cryptology ePrint Archive*, Report 2005/394, November 2005.
- Dakshi Agrawal, Dogan Kesdogan, and Stefan Penz. Probabilistic treatment of mixes to hamper traffic analysis. In *IEEE Symposium on Security and Privacy*, pages 16–27, Berkeley, CA, USA, May 2003. IEEE Computer Society.
- Ross Anderson. The eternity service. In *1st International Conference on the Theory and Applications of Cryptology (Pragocrypt '96)*, pages 242–252, Prague, Czech Republic, September/October 1996. Czech Technical University Publishing House.
- D. Asonov and J.C. Freytag. Almost optimal private information retrieval. *Privacy Enhancing Technology 2002*.
- Adam Back, Ian Goldberg, and Adam Shostack. Freedom systems 2.1 security issues and analysis. White paper, Zero Knowledge Systems, Inc., May 2001a.
- Adam Back, Ulf Möller, and Anton Stiglic. Traffic analysis attacks and trade-offs in anonymity providing systems. In Ira S. Moskowitz, editor, *Information Hiding workshop (IH 2001)*, volume 2137 of LNCS, pages 245–257. Springer-Verlag, April 2001b.
- Amos Beimel and Shlomi Dolev. Buses for anonymous message delivery. *Journal of Cryptology*, 16(1):25–39, 2003.
- Krista Bennett and Christian Grothoff. GAP – practical anonymous networking. In Roger Dingledine, editor, *Privacy Enhancing Technologies workshop (PET 2003)*, volume 2760 of LNCS, pages 141–160. Springer-Verlag, March 2003.
- Ron Berman, Amos Fiat, and Amnon Ta-Shma. Provable unlinkability against traffic analysis. In Ari Juels, editor, *Proceedings of Financial Cryptography (FC '04)*. Springer-Verlag, LNCS 3110, February 2004.
- Oliver Berthold, Hannes Federrath, and Stefan Köpsell. Web MIXes: A system for anonymous and unobservable Internet access. In H. Federrath, editor, *Designing Privacy Enhancing Technologies*, volume 2009 of LNCS, pages 115–129. Springer-Verlag, July 2000a.
- Oliver Berthold and Heinrich Langos. Dummy traffic against long term intersection attacks. In *Designing Privacy Enhancing Technologies, Proceedings of PET'02*, pages 110–128. Springer-Verlag, LNCS 2482, 2002.

- Oliver Berthold, Andreas Pfizmann, and Ronny Standtke. The disadvantages of free MIX routes and how to overcome them. In H. Federrath, editor, *Designing Privacy Enhancing Technologies*, volume 2009 of LNCS, pages 30–45. Springer-Verlag, July 2000b.
- John Bethencourt, Dawn Xiaodong Song, and Brent Waters. New constructions and practical applications for private stream searching (extended abstract). In *S&P*, pages 132–139. IEEE Computer Society, 2006. ISBN 0-7695-2574-1.
- George Dean Bissias, Marc Liberatore, , and Brian Neil Levine. Privacy vulnerabilities in encrypted HTTP streams. In *5th Workshop on Privacy Enhancing Technologies (PET2005)*, 2005.
- Avrim Blum, Dawn Xiaodong Song, and Shobha Venkataraman. Detection of interactive stepping stones: Algorithms and confidence bounds. In Erland Jonsson, Alfonso Valdes, and Magnus Almgren, editors, *RAID*, volume 3224 of *Lecture Notes in Computer Science*, pages 258–277. Springer, 2004. ISBN 3-540-23123-4.
- Nikita Borisov. An analysis of parallel mixing with attacker-controlled inputs. In *Proceedings of Privacy Enhancing Technologies workshop (PET 2005)*, May 2005.
- Philippe Boucher, Adam Shostack, and Ian Goldberg. Freedom systems 2.0 architecture. White paper, Zero Knowledge Systems, Inc., December 2000.
- Zach Brown. Cebolla – pragmatic IP anonymity. In *Ottawa Linux Symposium*, June 2002.
- Jan Camenisch and Anna Lysyanskaya. A formal treatment of onion routing. In Victor Shoup, editor, *Proceedings of CRYPTO 2005*, pages 169–187. Springer-Verlag, LNCS 3621, August 2005.
- David Chaum. Untraceable electronic mail, return addresses, and digital pseudonyms. *Communications of the ACM*, 24(2):84–88, February 1981.
- David Chaum. Blind signatures for untraceable payments. In David Chaum, Ronald L. Rivest, and Alan T. Sherman, editors, *Advances in Cryptology (Crypto'82)*, pages 199–203, New York and London, 1983. Plenum Press.
- David Chaum. The dining cryptographers problem: Unconditional sender and recipient untraceability. *Journal of Cryptology*, 1:65–75, 1988.
- Heyning Cheng and Ron Avnur. Traffic analysis of ssl encrypted web browsing. URL citeseer.ist.psu.edu/656522.html.
- Ian Clarke, Oskar Sandberg, Brandon Wiley, and Theodore W. Hong. Freenet: A distributed anonymous information storage and retrieval system. In H. Federrath, editor, *Designing Privacy Enhancing Technologies*, number 2009 in LNCS, pages 46–66, Berkeley, CA, USA, July 2000. Springer-Verlag.
- R. Clayton. Failures in a Hybrid Content Blocking System.
- Richard Clayton, Steven J. Murdoch, and Robert N. M. Watson. Ignoring the great firewall of china. In George Danezis and Philippe Golle, editors, *Privacy Enhancing Technologies workshop (PET 2006)*, LNCS. Springer-Verlag, 2006.

- George danezis. Traffic analysis of the http protocol over tls. <http://www.cl.cam.ac.uk/~gd216/TLSanon.pdf>.
- George Danezis. Forward secure mixes. In Jonsson Fisher-Hubner, editor, *Nordic workshop on Secure IT Systems (Norsec 2002)*, pages 195–207, Karlstad, Sweden, November 2002.
- George Danezis. Mix-networks with restricted routes. In Roger Dingledine, editor, *Privacy Enhancing Technologies workshop (PET 2003)*, volume 2760 of LNCS, pages 1–17, Dresden, Germany, March 2003a. Springer-Verlag.
- George Danezis. Statistical disclosure attacks. In Gritzalis, Vimercati, Samarati, and Katsikas, editors, *Security and Privacy in the Age of Uncertainty, (SEC2003)*, pages 421–426, Athens, May 2003b. IFIP TC11, Kluwer.
- George Danezis. The traffic analysis of continuous-time mixes. In *Proceedings of Privacy Enhancing Technologies workshop (PET 2004)*, LNCS, May 2004.
- George Danezis. Breaking four mix-related schemes based on universal re-encryption. In *Proceedings of Information Security Conference 2006*. Springer-Verlag, September 2006.
- George Danezis and Richard Clayton. Route fingerprinting in anonymous communications. In *IEEE International Conference on Peer-to-Peer Computing (P2P2006)*. IEEE, 2006.
- George Danezis and Jolyon Clulow. Compulsion resistant anonymous communications. In Mauro Barni, Jordi Herrera-Joancomartí, Stefan Katzenbeisser, and Fernando Pérez-González, editors, *Information Hiding*, volume 3727 of *Lecture Notes in Computer Science*, pages 11–25. Springer, 2005. ISBN 3-540-29039-7.
- George Danezis and Claudia Diaz. Improving the decoding efficiency of private search. Cryptology ePrint Archive, Report 2006/024, 2006. <http://eprint.iacr.org/>.
- George Danezis, Roger Dingledine, and Nick Mathewson. Mixminion: Design of a Type III Anonymous Remailer Protocol. In *IEEE Symposium on Security and Privacy*, Berkeley, CA, 11-14 May 2003.
- George Danezis and Ben Laurie. Minx: A simple and efficient anonymous packet format. In *Workshop on Privacy in the Electronic Society (WPES 2004)*. ACM, October 2004.
- George Danezis and Len Sassaman. Heartbeat traffic to counter $(n - 1)$ attacks. In *workshop on Privacy in the Electronic Society (WPES 2003)*, Washington, DC, USA, November 2003.
- George Danezis and Andrei Serjantov. Statistical disclosure or intersection attacks on anonymity systems. In Jessica J. Fridrich, editor, *Information Hiding*, volume 3200 of *Lecture Notes in Computer Science*, pages 293–308. Springer, 2004. ISBN 3-540-24207-4.
- Yvo Desmedt and Kaoru Kurosawa. How to break a practical mix and design a new one. In Bart Preneel, editor, *Advances in Cryptology (Eurocrypt 2000)*, volume 1807 of LNCS, pages 557–572, Bruges, Belgium, 14-18 May 2000. Springer-Verlag.

- Claudia Díaz, George Danezis, Christian Grothoff, Andreas Pfiztmann, and Paul F. Syverson. Panel discussion - mix cascades versus peer-to-peer: Is one concept superior? In Martin and Serjantov (2005), pages 242–242. ISBN 3-540-26203-2.
- Claudia Díaz and Bart Preneel. Reasoning about the anonymity provided by pool mixes that generate dummy traffic. In *Proceedings of 6th Information Hiding Workshop (IH'04)*, pages 309–325. Springer, LNCS 3200, 2004.
- Claudia Díaz, Len Sassaman, and Evelyne Dewitte. Comparison between two practical mix designs. In *Proceedings of 9th European Symposium on Research in Computer Security (ESORICS'04)*, pages 141–159. Springer-Verlag, LNCS 3193, 2004b.
- Claudia Diaz and Andrei Serjantov. Generalising mixes. In Roger Dingledine, editor, *Privacy Enhancing Technologies workshop (PET 2003)*, volume 2760 of LNCS, pages 18–31, Dresden, Germany, March 2003. Springer-Verlag.
- Claudia Diaz, Stefaan Seys, Joris Claessens, and Bart Preneel. Towards measuring anonymity. In Roger Dingledine and Paul Syverson, editors, *Privacy Enhancing Technologies workshop (PET 2002)*, volume 2482 of LNCS, pages 54–68, San Francisco, CA, USA, 14-15 April 2002. Springer-Verlag.
- T. Dierks and C. Allen. Rfc 2246: The tls protocol version 1.0. <http://www.ietf.org/rfc/rfc2246.txt>, January 1999.
- Roger Dingledine, Michael J. Freedman, David Hopwood, and David Molnar. A Reputation System to Increase MIX-net Reliability. In Ira S. Moskowitz, editor, *Information Hiding workshop (IH 2001)*, volume 2137 of LNCS, pages 126–141. Springer-Verlag, 25-27 April 2001.
- Roger Dingledine, Michael J. Freedman, and David Molnar. The free haven project: Distributed anonymous storage service. In H. Federrath, editor, *Designing Privacy Enhancing Technologies*, volume 2009 of LNCS, pages 67–95, Berkeley, CA, USA, July 2000. Springer-Verlag.
- Roger Dingledine and Nick Mathewson. Anonymity loves company: Usability and the network effect. In "Security and Usability", an O'Reilly Media book, August 2005.
- Roger Dingledine, Nick Mathewson, and Paul Syverson. Tor: The second-generation onion router. In *Proceedings of the 13th USENIX Security Symposium*, August 2004a.
- Roger Dingledine, Nick Mathewson, and Paul Syverson. Tor: The second-generation onion router. In *Proceedings of the 13th USENIX Security Symposium*, August 2004b.
- Roger Dingledine, Vitaly Shmatikov, and Paul F. Syverson. Synchronous batching: From cascades to free routes. In Martin and Serjantov (2005), pages 186–206. ISBN 3-540-26203-2.
- Roger Dingledine and Paul Syverson. Reliable MIX Cascade Networks through Reputation. In Matt Blaze, editor, *Financial Cryptography (FC '02)*, volume 2357 of LNCS. Springer-Verlag, March 2002.
- Roger Dingledine and Paul F. Syverson, editors. *Privacy Enhancing Technologies, Second International Workshop, PET 2002, San Francisco, CA, USA, April 14-15, 2002, Revised Papers*, volume 2482 of *Lecture Notes in Computer Science*, 2003. Springer. ISBN 3-540-00565-X.

- M. Dornseif. Government mandated blocking of foreign web content. In *von Knop, J., Haverkamp, W., Jessen, E. (eds.): Security, E-Learning, E-Services: Proceedings of the 17. DFN-Arbeitstagung über Kommunikationsnetze, Dusseldorf, 2003.*
- Peter Fairbrother. An improved construction for universal re-encryption. In Martin and Serjantov (2005), pages 79–87. ISBN 3-540-26203-2.
- N. Feamster, M. Balazinska, G. Harfst, H. Balakrishnan, and D. Karger. Infranet: Circumventing web censorship and surveillance. In Dan Boneh, editor, *USENIX Security Symposium*, pages 247–262, San Francisco, CA, 5-9 August 2002.
- Nick Feamster, Magdalena Balazinska, Winston Wang, Hari Balakrishnan, and David Karger. Thwarting web censorship with untrusted messenger discovery. In Roger Dingledine, editor, *Privacy Enhancing Technologies workshop (PET 2003)*, volume 2760 of LNCS, pages 125–140, Dresden, Germany, March 2003. Springer-Verlag.
- Nick Feamster and Roger Dingledine. Location diversity in anonymity networks. In *Proceedings of the Workshop on Privacy in the Electronic Society (WPES 2004)*, Washington, DC, USA, October 2004.
- Laurent Fousse and Jean-Ren Reinhard. Nymbaron: A type iii nymserver. On-line, 2006. <http://www.komite.net/laurent/soft/nymbaron/>.
- Michael J. Freedman and Robert Morris. Tarzan: A peer-to-peer anonymizing network layer. In Vijayalakshmi Atluri, editor, *ACM Conference on Computer and Communications Security (CCS 2002)*, pages 193–206, Washington, DC, November 2002. ACM.
- Michael J. Freedman, Emil Sit, Josh Cates, and Robert Morris. Introducing tarzan, a peer-to-peer anonymizing network layer. In Peter Druschel, M. Frans Kaashoek, and Antony I. T. Rowstron, editors, *International workshop on Peer-to-Peer Systems (IPTPS)*, volume 2429 of LNCS, pages 121–129, Cambridge, MA, March 2002. Springer-Verlag.
- Jun Furukawa and Kazue Sako. An efficient scheme for proving a shuffle. In Joe Kilian, editor, *Advances in Cryptology (Crypto 2001)*, volume 2139 of LNCS, pages 368–387, Santa Barbara, CA, USA, 19-23 August 2001. Springer-Verlag.
- Sharad Goel, Mark Robson, Milo Polte, and Emin Gun Sirer. Herbivore: A Scalable and Efficient Protocol for Anonymous Communication. Technical Report 2003-1890, Cornell University, Ithaca, NY, February 2003.
- Ian Goldberg. *A Pseudonymous Communications Infrastructure for the Internet*. PhD thesis, UC Berkeley, December 2000.
- David M. Goldschlag, Michael G. Reed, and Paul F. Syverson. Hiding routing information. In Ross J. Anderson, editor, *Information Hiding*, volume 1174 of LNCS, pages 137–150, Cambridge, U.K., 1996. Springer-Verlag.
- David M. Goldschlag, Michael G. Reed, and Paul F. Syverson. Onion routing. *Communications of the ACM*, 42(2):39–41, 1999.

- Philippe Golle. Reputable mix networks. In *Proceedings of Privacy Enhancing Technologies workshop (PET 2004)*, volume 3424 of LNCS, May 2004.
- Philippe Golle, Markus Jakobsson, Ari Juels, and Paul Syverson. Universal re-encryption for mixnets. In *Proceedings of the 2004 RSA Conference, Cryptographer's track*, San Francisco, USA, February 2004.
- Philippe Golle and Ari Juels. Dining cryptographers revisited. In *Proceedings of Eurocrypt 2004*, May 2004a.
- Philippe Golle and Ari Juels. Parallel mixing. In *Proceedings of the 11th ACM Conference on Computer and Communications Security (CCS 2004)*. ACM Press, October 2004b.
- Philippe Golle, XiaoFeng Wang, Markus Jakobsson, and Alex Tsow. Detering voluntary trace disclosure in re-encryption mix networks. In *Proceedings of the 2006 IEEE Symposium on Security and Privacy*, pages 121–131, Oakland, CA, May 2006. IEEE CS.
- Philippe Golle, Sheng Zhong, Dan Boneh, Markus Jakobsson, and Ari Juels. Optimistic mixing for exit-polls. In Yuliang Zheng, editor, *Advances in Cryptology (Asiacrypt 2002)*, volume 2501 of LNCS, pages 451–465, Queenstown, New Zealand, 1-5 December 2002. Springer-Verlag.
- Marcin Gomulkiewicz, Marek Klonowski, and Mirosław Kutylowski. Rapid mixing and security of chaum's visual electronic voting. In *Proceedings of ESORICS 2003*, October 2003.
- Marcin Gomulkiewicz, Marek Klonowski, and Mirosław Kutylowski. Onions based on universal re-encryption – anonymous communication immune against repetitive attack. In Chae Hoon Lim and Moti Yung, editors, *Information Security Applications, 5th International Workshop, WISA 2004*, volume 3325 of *Lecture Notes in Computer Science*, pages 400–410, Jeju Island, Korea, August 23–25 2004. Springer.
- Yong Guan, Xinwen Fu, Dong Xuan, P. U. Shenoy, Riccardo Bettati, and Wei Zhao. Netcamo: camouflaging network traffic for qos-guaranteed mission critical applications. *IEEE Transactions on Systems, Man, and Cybernetics, Part A* 31(4):253–265, 2001.
- Ceki Gülcü and Gene Tsudik. Mixing E-mail with Babel. In *Network and Distributed Security Symposium — NDSS '96*, pages 2–16, San Diego, California, February 1996. IEEE.
- Steven Hazel and Brandon Wiley. Achord: A variant of the chord lookup service for use in censorship resistant peer-to-peer publishing systems. In Peter Druschel, M. Frans Kaashoek, and Antony I. T. Rowstron, editors, *Peer-to-Peer Systems, First International workshop, IPTPS 2002*, volume 2429 of LNCS, Cambridge, MA, USA, 7-8 March 2002. Springer-Verlag. URL citeseer.nj.nec.com/hazel02achord.html.
- Johan Helsingius. Johan helsingius closes his internet remailer. <http://www.penet.fi/press-english.html>, August 1996a.
- Johan Helsingius. Johan helsingius gets injunction in scientology case privacy protection of anonymous messages still unclear. <http://www.penet.fi/injunc.html>, September 1996b.

- Johan Helsingius. Temporary injunction in the anonymous remailer case. <http://www.penet.fi/injunc1.html>, September 1996c.
- Andrew Hintz. Fingerprinting websites using traffic analysis. In Dingledine and Syverson (2003), pages 171–178. ISBN 3-540-00565-X.
- Yuval Ishai, Eyal Kushilevitz, Rafail Ostrovsky, and Amit Sahai. Cryptography from anonymity. Cryptology ePrint Archive, Report 2006/084, 2006. <http://eprint.iacr.org/>.
- Markus Jakobsson. A practical mix. In Kaisa Nyberg, editor, *Advances in Cryptology - EURO-CRYPT '98*, volume 1403 of LNCS, pages 448–461, Espoo, Finland, 31 May – 4 June 1998. Springer-Verlag.
- Markus Jakobsson. Flash Mixing. In *Principles of Distributed Computing - PODC '99*. ACM Press, 1999.
- Markus Jakobsson and Ari Juels. An optimally robust hybrid mix network. In *Principles of Distributed Computing (PODC 2001)*, pages 284–292, Newport, Rhode Island, USA, 26-29 August 2001. ACM.
- Markus Jakobsson, Ari Juels, and Ronald L. Rivest. Making mix nets robust for electronic voting by randomized partial checking. In Dan Boneh, editor, *USENIX Security Symposium*, pages 339–353, San Francisco, CA, USA, 5-9 August 2002. USENIX.
- Anja Jerichow, Jan Müller, Andreas Pfitzmann, Birgit Pfitzmann, and Michael Waidner. Real-Time MIXes: A Bandwidth-Efficient Anonymity Protocol. *IEEE Journal on Selected Areas in Communications*, 16(4):495–509, 1998.
- Shu Jiang, Nitin H. Vaidya, and Wei Zhao. Routing in packet radio networks to prevent traffic analysis. In *IEEE Information Assurance and Security Workshop*, June 2000.
- Sachin Katti, Dina Katabi, and Katay Puchala. Slicing the onion: Anonymous communication without pki. In *HotNets*, 2005.
- Dogan Kesdogan, Dakshi Agrawal, and Stefan Penz. Limits of anonymity in open environments. In Fabien A. P. Petitcolas, editor, *Information Hiding workshop (IH 2002)*, volume 2578 of LNCS, pages 53–69, Noordwijkerhout, The Netherlands, 7-9 October 2002a. Springer-Verlag.
- Dogan Kesdogan, Mark Borning, and Michael Schmeink. Unobservable surfing on the world wide web: Is private information retrieval an alternative to the mix based approach? In Dingledine and Syverson (2003), pages 224–238. ISBN 3-540-00565-X.
- Dogan Kesdogan, Jan Egner, and Roland Büschkes. Stop-and-Go MIXes: Providing probabilistic anonymity in an open system. In David Aucsmith, editor, *Information Hiding workshop (IH 1998)*, volume 1525 of LNCS, pages 83–98, Portland, Oregon, USA, 14-17 April 1998. Springer-Verlag.
- Joe Kilian and Kazue Sako. Receipt-free MIX-type voting scheme — a practical solution to the implementation of a voting booth. In Louis C. Guillou and Jean-Jacques Quisquater, editors, *Advances in Cryptology (Eurocrypt 1995)*, volume 921 of LNCS, pages 393–403, Saint-Malo, France, 21-25 May 1995. Springer-Verlag.

- Lea Kissner, Alina Oprea, Michael K. Reiter, Dawn Xiaodong Song, and Ke Yang. Private keyword-based push and pull with applications to anonymous communication. In Markus Jakobsson, Moti Yung, and Jianying Zhou, editors, *ACNS*, volume 3089 of *Lecture Notes in Computer Science*, pages 16–30. Springer, 2004. ISBN 3-540-22217-0.
- Marek Klonowski and Mirosław Kutylowski. Provable anonymity for networks of mixes. In *Proceedings of Information Hiding Workshop (IH 2005)*, June 2005.
- Marek Klonowski, Mirosław Kutylowski, Anna Lauks, and Filip Zagorski. Universal re-encryption of signatures and controlling anonymous information flow. In *WARTACRYPT '04 Conference on Cryptology*, Bedlewo/Poznan, July 1–3 2004.
- Marek Klonowski, Mirosław Kutylowski, and Filip Zagorski. Anonymous communication with on-line and off-line onion encoding. In Peter Vojts, Mria Bielikov, Bernadette Charron-Bost, and Ondrej Skora, editors, *SOFSEM 2005: Theory and Practice of Computer Science, 31st Conference on Current Trends in Theory and Practice of Computer Science*, Lecture Notes in Computer Science, pages 229–238, Liptovsk Jn, Slovakia, January 22–28 2005,. 3381.
- Stefan Köpsell and Ulf Hilling. How to achieve blocking resistance for existing systems enabling anonymous web surfing. In *Proceedings of the Workshop on Privacy in the Electronic Society (WPES 2004)*, Washington, DC, USA, October 2004.
- Dennis Kügler. An Analysis of GUnet and the Implications for Anonymous, Censorship-Resistant Networks. In Roger Dingledine, editor, *Privacy Enhancing Technologies workshop (PET 2003)*, volume 2760 of *LNCS*, pages 161–176, Dresden, Germany, March 2003. Springer-Verlag.
- Brian N. Levine, Michael K. Reiter, Chenxi Wang, and Matthew K. Wright. Timing attacks in low-latency mix-based systems. In Ari Juels, editor, *Proceedings of Financial Cryptography (FC '04)*. Springer-Verlag, LNCS 3110, February 2004.
- John Leyden. Anonymizer looks for gaps in great firewall of china. *The Register*, April 3 2006.
- Tianbo Lu, Binxing Fang, Yuzhong Sun, and Xueqi Cheng. Performance analysis of wongoo system. In *CIT*, pages 716–723. IEEE Computer Society, 2005a. ISBN 0-7695-2432-X.
- Tianbo Lu, Binxing Fang, Yuzhong Sun, and Li Guo. Some remarks on universal re-encryption and a novel practical anonymous tunnel. In Xicheng Lu and Wei Zhao, editors, *ICCNMC*, volume 3619 of *Lecture Notes in Computer Science*, pages 853–862. Springer, 2005b. ISBN 3-540-28102-9.
- David Martin and Andrew Schulman. Deanonimizing users of the safeweb anonymizing service. Technical Report 2002-003, Boston University Computer Science Department, February 2002.
- David Martin and Andrei Serjantov, editors. *Privacy Enhancing Technologies, 4th International Workshop, PET 2004, Toronto, Canada, May 26-28, 2004, Revised Selected Papers*, volume 3424 of *Lecture Notes in Computer Science*, 2005. Springer. ISBN 3-540-26203-2.
- Nick Mathewson. Underhill: A proposed type 3 nymserver protocol specification. On-line, 2005. <http://svn.conuropsis.org/nym3/trunk/doc/nym-spec.txt>.

- Nick Mathewson and Roger Dingledine. Practical traffic analysis: Extending and resisting statistical disclosure. In *Proceedings of Privacy Enhancing Technologies workshop (PET 2004)*, LNCS, May 2004.
- David Mazières and M. Frans Kaashoek. The Design, Implementation and Operation of an Email Pseudonym Server. In *ACM Conference on Computer and Communications Security (CCS'98)*, pages 27–36, San Francisco, CA, USA, 3-5 November 1998. ACM Press.
- Markus Michels and Patrick Horster. Some remarks on a receipt-free and universally verifiable mix-type voting scheme. In Kwangjo Kim and Tsutomu Matsumoto, editors, *Advances in Cryptology (Asiacrypt '96)*, volume 1163 of LNCS, pages 125–132, Kyongju, Korea, 3-7 November 1996. Springer-Verlag.
- Masashi Mitomo and Kaoru Kurosawa. Attack for flash MIX. In Tatsuaki Okamoto, editor, *Advances in Cryptology (Asiacrypt 2000)*, volume 1976 of LNCS, pages 192–204, Kyoto, Japan, 3-7 December 2000. Springer-Verlag.
- Bodo Möller. Provably secure public-key encryption for length-preserving chaumian mixes. In Marc Joye, editor, *Topics in Cryptology CT-RSA 2003*, volume 2612 of LNCS, pages 244–262, San Francisco, CA, USA, 13-17 April 2003. Springer-Verlag.
- Ulf Möller, Lance Cottrell, Peter Palfrader, and Len Sassaman. Mixmaster Protocol — Version 2. Draft, July 2003.
- Steven J. Murdoch and George Danezis. Low-cost traffic analysis of Tor. In *Proceedings of the 2005 IEEE Symposium on Security and Privacy*. IEEE CS, May 2005.
- C. Andrew Neff. A verifiable secret shuffle and its application to e-voting. In Pierangela Samarati, editor, *ACM Conference on Computer and Communications Security (CCS 2002)*, pages 116–125. ACM Press, November 2001.
- Richard E. Newman, Ira S. Moskowitz, Paul Syverson, and Andrei Serjantov. Metrics for traffic analysis prevention. In *Privacy Enhancing Technologies Workshop*, Dresden, Germany, March 2003.
- Ron Newman. The Church of Scientology vs. anon.penet.fi, March 1997. <http://www.xs4all.nl/~kspaink/cos/rnewman/anon/penet.html>.
- Luke O'Connor. On blending attacks for mixes with memory. In *Proceedings of Information Hiding Workshop (IH 2005)*, June 2005.
- Wakaha Ogata, Kaoru Kurosawa, Kazue Sako, and Kazunori Takatani. Fault tolerant anonymous channel. In Yongfei Han, Tatsuaki Okamoto, and Sihan Qing, editors, *Information and Communication Security, First International Conference (ICICS'97)*, volume 1334 of LNCS, pages 440–444, Beijing, China, 11-14 November 1997. Springer-Verlag.
- Miyako Ohkubo and Masayuki Abe. A Length-Invariant Hybrid MIX. In Tatsuaki Okamoto, editor, *Advances in Cryptology (Asiacrypt 2000)*, volume 1976 of LNCS, pages 178–191, Kyoto, Japan, 3-7 December 2000. Springer-Verlag.

- Rafail Ostrovsky and William E. Skeith III. Private searching on streaming data. In Victor Shoup, editor, *CRYPTO*, volume 3621 of *Lecture Notes in Computer Science*, pages 223–240. Springer, 2005. ISBN 3-540-28114-2.
- Lasse Øverlier and Paul Syverson. Locating hidden servers. In *Proceedings of the 2006 IEEE Symposium on Security and Privacy*. IEEE CS, May 2006.
- Peter Palfrader. Echolot: a pinger for anonymous remailers. <http://www.palfrader.org/echolot/>.
- Sameer Parekh. Prospects for remailers: where is anonymity heading on the internet? *First Monday*, 1(2), August 5 1996. On-line journal <http://www.firstmonday.dk/issues/issue2/remailers/>.
- Choonsik Park, Kazutomo Itoh, and Kaoru Kurosawa. Efficient anonymous channel and all/nothing election scheme. In Tor Helleseth, editor, *Advances in Cryptology (Eurocrypt '93)*, volume 765 of *LNCS*, pages 248–259, Lofthus, Norway, 23-27 May 1993. Springer-Verlag.
- Andreas Pfitzmann and Marit Hansen. Anonymity, Unobservability and Pseudonymity – A Proposal for Terminology. In Hannes Federrath, editor, *Designing Privacy Enhancing Technologies*, pages 1–9. Springer-Verlag, LNCS 2009, 2001.
- Andreas Pfitzmann, Birgit Pfitzmann, and Michael Waidner. ISDN-mixes: Untraceable communication with very small bandwidth overhead. In Wolfgang Effelsberg, Hans Werner Meuer, and Günter Müller, editors, *GI/ITG Conference on Communication in Distributed Systems*, volume 267 of *Informatik-Fachberichte*, pages 451–463. Springer-Verlag, February 1991.
- Birgit Pfitzmann. Breaking efficient anonymous channel. In Alfredo De Santis, editor, *Advances in Cryptology (Eurocrypt '94)*, volume 950 of *LNCS*, pages 332–340, Perugia, Italy, 9-12 May 1994. Springer-Verlag.
- Birgit Pfitzmann and Andreas Pfitzmann. How to break the direct RSA-implementation of MIXes. In Jean-Jacques Quisquater and Joos Vandewalle, editors, *Advances in Cryptology (Eurocrypt '89)*, volume 434 of *LNCS*, pages 373–381, Houthalen, Belgium, 10-13 April 1990. Springer-Verlag.
- Charles Rackoff and Daniel R. Simon. Cryptographic defense against traffic analysis. In *ACM Symposium on Theory of Computing (STOC'93)*, pages 672–681. ACM, 1993.
- Josyula R. Rao and Pankaj Rohatgi. Can pseudonymity really guarantee privacy? In *Proceedings of the 9th USENIX Security Symposium*, pages 85–96. USENIX, August 2000.
- Jean-François Raymond. Traffic Analysis: Protocols, Attacks, Design Issues, and Open Problems. In Hannes Federrath, editor, *Designing Privacy Enhancing Technologies*, volume 2009 of *LNCS*, pages 10–29. Springer-Verlag, July 2000.
- 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.

- Michael G. Reed, Paul F. Syverson, and David M. Goldschlag. Onion routing network for securely moving data through communication networks. United States Patent 6,266,704, 2001.
- Michael Reiter and Aviel Rubin. Crowds: Anonymity for web transactions. *ACM Transactions on Information and System Security (TISSEC)*, 1(1):66–92, 1998.
- Michael Reiter and XiaoFeng Wang. Fragile mixing. In *Proceedings of the 11th ACM Conference on Computer and Communications Security (CCS 2004)*. ACM Press, October 2004.
- Marc Rennhard and Bernhard Plattner. Introducing MorphMix: Peer-to-Peer based Anonymous Internet Usage with Collusion Detection. In *Workshop on Privacy in the Electronic Society (WPES 2002)*, Washington, DC, USA, November 2002.
- L. Sassaman, B. Cohen, and N. Mathewson. The pynchon gate: a secure method of pseudonymous mail retrieval. *Proceedings of the 2005 ACM workshop on Privacy in the electronic society*, pages 1–9, 2005.
- Andrei Serjantov. Anonymizing censorship resistant systems. In Peter Druschel, M. Frans Kaashoek, and Antony I. T. Rowstron, editors, *International Peer-To-Peer Systems workshop (IPTPS 2002)*, volume 2429 of LNCS, pages 111–120, Cambridge, MA, USA, 7-8 March 2002. Springer-Verlag.
- Andrei Serjantov. *On the Anonymity of Anonymity Systems*. PhD thesis, University of Cambridge, 2004.
- Andrei Serjantov and George Danezis. Towards an information theoretic metric for anonymity. In Roger Dingledine and Paul Syverson, editors, *Privacy Enhancing Technologies workshop (PET 2002)*, volume 2482 of LNCS, pages 41–53, San Francisco, CA, USA, 14-15 April 2002. Springer-Verlag.
- Andrei Serjantov, Roger Dingledine, and Paul Syverson. From a trickle to a flood: Active attacks on several mix types. In Fabien A. P. Petitcolas, editor, *Information Hiding workshop (IH 2002)*, volume 2578 of LNCS, pages 36–52, Noordwijkerhout, The Netherlands, 7-9 October 2002. Springer-Verlag.
- Andrei Serjantov and Peter Sewell. Passive attack analysis for connection-based anonymity systems. In *European Symposium on Research in Computer Security (ESORICS 2003)*, Gjøvik, Norway, 13-15 October 2003.
- Claude Shannon. A mathematical theory of communication. *The Bell System Technical Journal*, 27:379–423:623–656, 1948.
- Rob Sherwood, Bobby Bhattacharjee, and Aravind Srinivasan. P5: A protocol for scalable anonymous communication. In *IEEE Symposium on Security and Privacy*, page 58, Berkeley, California, USA, 12-15 May 2002. IEEE Computer Society.
- Vitaly Shmatikov. Probabilistic analysis of anonymity. In *Computer Security Foundations workshop (CSFW-15 2002)*, pages 119–128, Cape Breton, Nova Scotia, Canada, 24-26 June 2002. IEEE Computer Society.

- Michael Singer. CIA Funded SafeWeb Shuts Down. http://siliconvalley.internet.com/news/article.php/3531_926921, November 20 2001.
- Qixiang Sun, Daniel R. Simon, Yi-Min Wang, Wilf Russell, Venkata N. Padmanabhan, and Lili Qiu. Statistical identification of encrypted web browsing traffic. In *IEEE Symposium on Security and Privacy*, pages 19–30, 2002.
- Paul F. Syverson, Gene Tsudik, Michael G. Reed, and Carl E. Landwehr. Towards an analysis of onion routing security. In Hannes Federrath, editor, *Designing Privacy Enhancing Technologies*, volume 2009 of *LNCS*, pages 96–114, Berkeley, CA, USA, 25-26 July 2000. Springer-Verlag.
- Parisa Tabriz and Nikita Borisov. Breaking the collusion detection mechanism of morphmix. In George Danezis and Philippe Golle, editors, *Privacy Enhancing Technologies workshop (PET 2006)*, LNCS. Springer-Verlag, 2006.
- Brenda Timmerman. A security model for dynamic adaptive traffic masking. In *New Security Paradigms Workshop*, pages 107–115, Langdale, Cumbria, UK, 1997. ACM.
- Brenda Timmerman. Secure dynamic adaptive traffic masking. In *New Security Paradigms Workshop*, pages 13–24, Ontario, Canada, September 1999. ACM.
- B. R. Venkatraman and Richard E. Newman-Wolfe. Performance analysis of a method for high level prevention of traffic analysis using measurements from a campus network. In *Proceeding of the IEEE/ACM Tenth Annual Computer Security Applications Conference*, pages 288–297, Orlando, FL, December 5-9 1994. IEEE CS Press.
- Michael Waidner and Birgit Pfizmann. The dining cryptographers in the disco — underconditional sender and recipient untraceability with computationally secure serviceability. In Jean-Jacques Quisquater and Joos Vandewalle, editors, *Advances in Cryptology (Eurocrypt '89)*, volume 434 of *LNCS*, page 690, Houthalen, Belgium, 10-13 April 1989. Springer-Verlag.
- G. Walton. Chinas golden shield: corporations and the development of surveillance technology in the Peoples Republic of China. *Montreal: International Centre for Human Rights and Democratic Development*, URL (consulted 29 October 2001): <http://www.ichrdd.ca/frame.iphtml>, 2001.
- Xinyuan Wang, Shiping Chen, and Sushil Jajodia. Tracking anonymous peer-to-peer voip calls on the internet. In *Proceedings of the ACM Conference on Computer and Communications Security*, pages 81–91, November 2005.
- Xinyuan Wang and Douglas S. Reeves. Robust correlation of encrypted attack traffic through stepping stones by manipulation of interpacket delays. In Sushil Jajodia, Vijayalakshmi Atluri, and Trent Jaeger, editors, *ACM Conference on Computer and Communications Security*, pages 20–29. ACM, 2003. ISBN 1-58113-738-9.
- Douglas Wikström. How to break, fix, and optimize “optimistic mix for exit-polls”. Technical Report T2002-24, Swedish Institute of Computer Science, SICS, Box 1263, SE-164 29 Kista, SWEDEN, 2002.
- Douglas Wikström. Elements in $Z_p^* \setminus G_q$ are dangerous. Technical Report T2003-05, Swedish Institute of Computer Science, SICS, Box 1263, SE-164 29 Kista, SWEDEN, 2003a.

- Douglas Wikström. Four practical attacks for “optimistic mixing for exit-polls”. Technical Report T2003-04, Swedish Institute of Computer Science, SICS, Box 1263, SE-164 29 Kista, SWEDEN, 2003b.
- Matthew Wright, Micah Adler, Brian Neil Levine, and Clay Shields. An analysis of the degradation of anonymous protocols. In *Network and Distributed Security Symposium (NDSS '02)*, San Diego, California, 6-8 February 2002.
- Matthew Wright, Micah Adler, Brian Neil Levine, and Clay Shields. Defending anonymous communication against passive logging attacks. In *IEEE Symposium on Security and Privacy*, page 28, Berkeley, CA, USA, 11-14 May 2003. IEEE Computer Society.
- Ye Zhu and Riccardo Bettati. Un-mixing mix traffic. In *Proceedings of Privacy Enhancing Technologies workshop (PET 2005)*, May 2005.
- Ye Zhu, Xinwen Fu, Bryan Graham, Riccardo Bettati, and Wei Zhao. On flow correlation attacks and countermeasures in mix networks. In *Proceedings of Privacy Enhancing Technologies workshop (PET 2004)*, LNCS, May 2004.
- Lasse Verlier and Paul Syverson. Valet services: Improving hidden servers with a personal touch. In George Danezis and Philippe Golle, editors, *Privacy Enhancing Technologies workshop (PET 2006)*, LNCS. Springer-Verlag, 2006.