Privacy in the Internet Architecture
MobilityFirst white paper draft
(work-in-progress)

Janne Lindqvist
WINLAB
Rutgers University

January 10, 2012

Contents

Chapter 1

Introduction

Privacy issues on the Internet have been studied for many angles since the network became popular. Today, several themes such as privacy on social networks [39], web browsing and cookies (e.g. Do Not Track [66]) have received considerable attention. In contrast, in this paper, we focus on trying to identify fundamental privacy *choke-points* in the Internet *architecture*. First, we will introduce three fundamental concepts related to privacy: *confidentiality of content, communication participants and location*, and proceed to discuss how privacy can be violated by observers in different locations. The following discussion has been partially adapted from "Practical Privacy Enhancing Technologies for Mobile Systems" [52] by the author.

1.1 Privacy Concepts

Confidentiality of Communication Content The concept of *confidentiality of communication content* is rather self-explanatory; we are concerned is to protect communication. For example, the voice in a phone conversation, the body of an email message or the data transmitted during a web browsing session. Confidentiality of communication content and the right to privacy are commonly associated concepts, and in the context of computer communication are often considered to merely mean data encryption.

Confidentiality of Communication Participants Confidentiality of communication participants means concealing the identities of the involved parties. We refer to this concept also as *identity privacy*. In practice, this kind of confidentiality is achieved with *pseudonymity* or *anonymity*. In short, pseudonymity means that, for example, a person has an identifier that cannot be directly attributed to that person. Whereas anonymity means that the person cannot be identified within an anonymity set [72], which could be "everybody who goes to a local coffee shop", "everybody at a given time at a local coffee shop", "everybody who use computers at Rutgers University" or "all Internet users globally".

The reason why both confidentiality of communication content and communication participants are equally important in privacy protection is summarized in the following quote from a US Supreme Court Justice:

> "The evil incident to invasion of privacy of the telephone is far greater than that involved in tampering with the mails. Whenever a telephone line is tapped, the privacy of persons at both ends of the line is invaded, and all conversations between them upon any subject, and although proper,

> confidential and privileged, may be overheard. Moreover, the tapping of one man's telephone line involves the tapping of the telephone of every other person whom he may call or who may call him. As a means of espionage, writs of assistance and general warrants are but puny instruments of tyranny and oppression when compared to wire-tapping." – Justice Louis Brandeis, dissenting opinion in *Olmstead vs. United States* (277 US 438, 1928, pp. 475–476), cited from a secondary source [30].

Confidentiality of Location The last privacy concept to be introduced is about protecting or hiding the location of the communicating party. Providing location privacy does not necessarily require protection of the confidentiality of communication content or the participants. This means that we may know in detail who is talking to whom and the content of their communication, but not their physical or actual network location. However, many anonymity systems (e.g. Tor [32]) provide also location privacy while providing sender and receiver anonymity. Location privacy can be provided with different granularities, for example, locally, on the level of access point usage, on the level of network access (network prefix), or by the country from which the user currently is contacting the network. Naturally, finer or coarser-grained granularities can also be defined.

As a concrete example of location privacy protection let us consider GSM networks. Roughly speaking, the GSM number consists of the country code, operator code and finally the number of the particular subscription (note that e.g. US is exception to this and GSM numbers can be "local"). Thus, an outside observer by looking at the subscriber's number can identify where the subscription was bought, but not much more. The GSM network operator, however, knows all the time where the user resides because of the access points near the mobile phone. GSM networks fail to provide complete location privacy against outside observers, too. The dial tone is different in many countries, which already gives a clue where the user might reside. Furthermore, when the mobile phone is switched off and somebody tries to call the user, the caller hears recorded message from the operator along the lines of "the number you have tried to reach.." This message is usually given in different languages depending on the country, and thus, gives a good hint about the country the user resides in at the moment. Maybe somebody wants to visit a person's house when the person is not present. Maybe the other person wants to be reachable from the same number, but not want to reveal the fact that he or she has left the country.

1.2 Privacy Violations

In this paper, we will refer an actor that commits or tries to commit privacy violations against a user as an *attacker*.

Attacks: Questions Given what we have discussed above on what we would want to protect with respect to privacy, another way to look at is to ask questions that what kind of questions an attacker could ask. These include

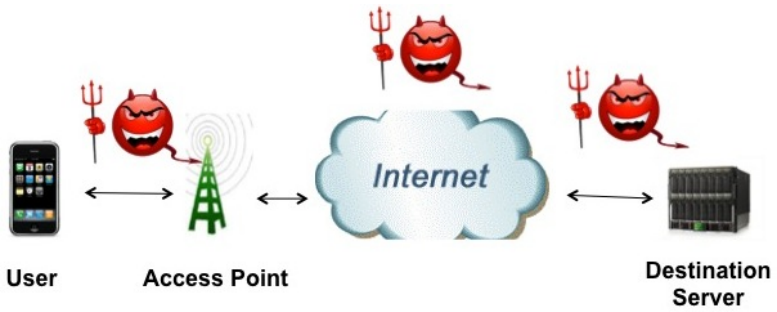

Figure 1.1: Attackers against user privacy can reside in several different locations: in the local access network, some where towards the path to the destination, or nearby the destination. More powerful attackers can also be at several locations at the same time.

- Who are you?
  - Have I seen you before?
- Who do you talk to?
  - Did you talk to them before?
- What are you talking about?
- What is your location?
  - Have you been here before?

Note that these questions are connected knowing places you go can tell who you are e.g. home/work pairs have been shown highly likely to be unique [38] even on the scale of zip codes.

For a more detailed coverage of privacy violations in the context of US legalization and court cases, we refer the reader to Daniel J. Solove's "Taxonomy of Privacy" [70]. For a more complete treatment of related terminology, we refer the reader to "Anonymity, Unlinkability, Undetectability, Unobservability, Pseudonymity, and Identity Management – A Consolidated Proposal for Terminology" by Pfitzmann and Hansen [61].

Attacker Capabilities and Locations We have above discussed briefly in the form of questions that how an attacker can violate the privacy of a user. For the purposes of inspecting choke-points in the Internet architecture, we also need to define where the attacker can ask those questions.

Fundamentally, there are three different kind of locations that the attacker can reside in: 1. near the user (e.g. in the access network, WiFi network, hotel Ethernet), 2. beyond the access network (e.g. 1 or more hops beyond the access network) and 3. nearby the destination server. Note that the third location falls also into the second category, however, the distinction is here important. An attacker may be interested in only e.g.

who is accessing a particular server, and therefore a good way to gain that knowledge is to introduce passive tapping point near the access network of that server. The attackers in the second category may reside in operators networks, a centralized points for a country's Internet access or even in Internet Exchange (IX) points [57].

The easiest attack to carry is passive listening in a local radio network (e.g. WiFi). For example, with WiFi no special equipment is needed, attackers can use e.g. *tcdump*, a diagnostic tool available on all unix-type system to capture all traffic. Furthermore, WiFi allows passive attacks against location privacy by outsiders at the same location even when 1) the attackers are not in the same network and 2) the data traffic is encrypted [53]. However, this attack scenario is naturally constrained to that particular location. However, e.g. operators of nation-wide WiFi networks could carry out this attack in several different locations at the same time. More powerful adversaries can launch active attacks, for example, manipulate protocols messages, repeat old messages or impersonate some of the communication parties.

Chapter 2

Architectural Privacy Choke-Points

In this chapter, we will consider the choke-points for privacy in the current Internet architecture. We will also give an overview of existing approaches and proposals for the problems. In general, we can divide the problems in to two categories *data plane* and *control plane*.

2.1 Data Plane Choke-Points

In the data plane, there are at least four different points, which can be used to violate user privacy. These are *naming*, *name lookup (or service discovery)*, *authentication and key exchange*, and *data transit*.

2.1.1 Naming - Identifiers in Networking Protocols

Network protocols on all layers of the protocol stack use identifiers for establishing and multiplexing connections. Unfortunately, these identifiers introduce privacy problems. The identifiers might not directly reveal the identity of the communication participants, but persistent identifiers can be used to (re)locate the users. The attacker tries to find answers to questions such as: "Did this traffic sample come from device D?" as discussed in the previous chapter.

The basic naming available in today's Internet Architecture can be summarized as IP addresses and short-hands for IP addresses: DNS names. Introduction of WWW and HTTP (later also SIP) gave us more complex naming schemes, but we will leave them out of this discussion for now. It is worth noting that IP addresses today server both as *device identifiers* and also as location identifiers for routing purposes. Thus, finding out the IP address of a connection, can be used to both identify and locate a communication participant.

Similar problems with identifier usage also exist with the next generation Internet Protocol: IPv6. The IPv6 [28] addressing architecture [43] allows hosts to configure their addresses automatically without involving a server. The original specifications [75] of the IPv6 stateless address autoconfiguration created a lot of controversy. The host could choose the interface identifier portion of the address by using the EUI-64 encoding of the MAC address. Creating the address this way meant that there would be a persistent unique identifier for the host wherever it might travel. This would have made the privacy of the mobile user much worse compared to the situation with the current Internet Protocol (IPv4) [62]. The privacy problem was fixed five years later in privacy extensions for the IPv6 stateless address autoconfiguration [58] while the autoconfiguration mechanism specification was also updated. The privacy extensions provided a way to choose the

interface identifier in random. However, Escudero-Pascual [34] argues that the privacy extensions present another kind of problem for privacy, because an attacker can observe that the privacy extensions are used: it might be interesting in some situations to see which of the clients are to keep hidden from observation. It should be also noted that even the current specifications [76] of the IPv6 stateless address autoconfiguration today allow it to be implemented without the corresponding updated [59] privacy extensions. Interestingly, the use of IP addresses as Home Address or Care of Address in Mobile IP and Mobile IPv6 can also introduce location privacy problems [34, 49] and, for example, the security design of IPv6 did not even consider the location privacy [11]. Finally, even the IPsec architecture ESP [46] uses identifiers called SPIs that can be used to identify the mobile host, although not persistently [8].

Even application layer identifiers (or names) can be harmful for the location privacy of the user. The user names used in the Session Initiation Protocol [67], which is used for VoIP signaling, for example, can be used to identify the mobile host and the user. Further, it has been shown that dynamic DNS is a very effective way to track the user remotely [40]. Dynamic DNS servers allow the user to register their current IP address for a particular domain name. This enables the users to be reachable with a human-readable name, even though they would be changing their places, and thus, frequently networks and IP addresses. Unfortunately, the IP addresses can be effectively mapped with geolocation services and researchers have even developed an automatic tool for dynamic DNS-based tracking [40].

There are also other kinds of practical, but limited, mechanisms that provide additional privacy for mobile users without the need for infrastructure support. For example, Aura & Zugenmaier [12] proposed that the mobile host should acquire a new IPv6 address for every new TCP flow. This introduces the equivalent privacy for IPv6 that the Network Address Translation provides for IPv4.

Anonymity Systems Anonymity systems research can be said to have started with Chaum's seminal paper on network mixes [20]. The article published in *Communications of the ACM* has been inspiring researchers for over three decades. Anonymity systems can generally be categorized as *low-latency* or *high-latency*. The systems can be further categorized by their type, for example, as *mixes* [20], *onion routing* [37, 74], *dc-nets* [22] or by the infrastructure support they require: a single server or remailer, multiple servers or *peer-to-peer* systems.

We start our review of anonymity systems with the simplest form of infrastructure support: a single server in the hands of a trusted third party. The single server acts, for example, as an email remailer. Users send emails to the server including information on the final recipient, and the remailer assigns pseudonyms for both parties to allow subsequent communication. To enhance the system, the traffic to the server can be protected using encryption, such as PGP. However, in addition to the problem that the server needs to be trusted, the architecture is brittle, as the story of the famous Finnish anonymous remailer *anon.penet.fi* demonstrates [42]. In 1995 and 1996, the Church of Scientology wanted to find out who were behind some pseudonyms relayed by the

anon.penet.fi server. These issues finally led to the shutdown of the system, since it could not guarantee the privacy of their user because of explicit real address and pseudonym mappings.

A more advanced design using multiple servers as a cascade was proposed by Chaum already in 1981 [20] as noted above. The mix network provides sender and receiver anonymity using public key cryptography, where the users have digital pseudonyms denoted by their public key. The mix(es) provide anonymity by introducing bitwise unlinkability: the input to the mix cannot be correlated to the output of the mix. Chaum also proposed a way to arrange the mixes into a network and the use of dummy traffic to further complicate traffic analysis. As the mix network relies on public key cryptography, it is computationally secure. In 1988, Chaum proposed the dining cryptographers network (dc-net) that, in contrast to mixes, provides information theoretic security, but is unfortunately not very practical (and never deployed in any form) due to the amount of messages it needs in every round [22]. Furthermore, despite its information theoretic security, dc-net is also vulnerable to many traffic analysis attacks, as discussed later.

Perhaps the most important anonymity system that has been developed from Chaum's mix nets is *onion routing* [36, 37, 73, 74]. Whereas the original mix design can be understood as a packet anonymizer, the onion routing system provides the equivalent of circuit switching for mixes. The anonymity of onion routing is established by distributed trust: the clients contact the onion router network that is build as an overlay on the Internet, and hide their path by layers of encryption. As a result the onion router was a success and in 1999 the network served more than one million Web connections per month in twenty countries [36]. The onion routing project has been developed further, and the current de facto way of achieving sender and receiver anonymity on the Internet is the second-generation onion router succinctly known as Tor [32]. The Tor network is currently estimated to have several hundred thousand users and over a thousand active volunteer servers relaying the traffic [33]. Although the low-latency network is primarily used for Web browsing, it can be used for accessing a number of different connection-oriented services, such as IRC or SSH. One interesting feature that onion routing (and Tor) provide is the hidden server [74]. A user can set up a server (SSH, web, etc.) that has a name as used in the DNS system to resolve IP addresses. The difference is that this hidden server name is resolvable only in the onion routing system, and the IP address is thus hidden from its users. A hidden server establishes an anonymous connection to a rendezvous point in the onion routing network, and a client that wishes to contact the server establishes another anonymous connection via the rendezvous point. Conceptually, this rendezvous mechanism could be seen as the anonymous equivalent of a Mobile IP home agent or a SIP relay server.

There are also many other designs for providing anonymity on the Internet. Among the most succesful ones were the Crowds [64, 65], designed for solely providing privacy for Web browsing, and the Freenet peer-to-peer network [23, 24] for anonymous file storage.

Limits of Anonymity Systems and Privacy The privacy that anonymity systems can provide is limited at least by the anonymity set [72]. In other words, anonymity is about hiding in the crowds. Therefore, usability is also important for an anonymity system. Unless the system can attract users there will be no crowds to hide in [31]. The anonymity a system provides can also be measured with entropy [29, 69]. There are a number of attacks, countermeasures and their analysis [13, 51, 56, 60, 63] documented in the literature on anonymity systems. However, in the end, it seems that long-term communication cannot be anonymous. This is shown by the results of passive logging attacks called *predecessor*, *intersection* and *disclosure*.

The predecessor attack was first introduced against Crowds [64], and later a related attack against onion routing was analyzed [73]. However, Wright et al. generalized the attack to work against all known anonymity protocols (some introduced above) and showed upper bounds for how long these protocols can maintain anonymity [77]. Later, Wright et al. improved their analysis by removing some of their previous simplifying assumptions and in simulations, they also showed that, in practice, the anonymity of real systems is worse than the theoretical upper bounds [78–80]. The predecessor attack can be established when the attacker controls some nodes in the anonymous system. The passive attacker observes path reformations where identifiable streams recur from possible initiators, and uses this information to deduce the identity of the users.

Another attack that sets limits for the anonymity provided by anonymity systems is the disclosure attack. The attacker sits on the edges of the anonymity system viewing it as a black box and tries to correlate traffic in order to identify all the peers of a particular user. The attack was first proposed by Kesdogan et al. [48] and further refined by Agrawal et al. [4, 5]. The disclosure is a NP-complete problem and therefore computationally exhausting, which led to a proposal for a statistical disclosure attack [26]. The statistical disclosure attack is further refined by Danezis and Serjantov [27], as well as by Mathewson and Dingledine [54]. The results of the analysis of the predecessor and disclosure attacks showed that long-term communications cannot be protected by known anonymity systems.

2.1.2 Name lookup

Many networking protocols as a first step try to resolve a name to a IP address. On the Internet scale, this is done with DNS, but many legacy protocols such as NetBIOS over TCP (NBT) used by the Windows operating systems can use their own name servers (e.g. WINS). These lookups can trivially reveal what kind of information you are interested in, and possibly the organization you work for [10]. For example, if the computer connects to a VPN gateway of its organization (e.g., vpn-gw.winlab.rutgers.edu), a look at a DNS log is sufficient to identify the company.

In our previous work [10] on analyzing leaks from network traces, we found leaks from e.g. trying to *resolve private network names*, *default suffixes*, *DHCP host identification and DNS registration*, *Domain Controller (Windows networking specific)* and, *file shares and printer discovery* etc.

2.1.3 Authentication and key exchange

It is prudent engineering practice [3] for cryptographic protocols that a name for the authenticating principal is given during the process of authentication:

> "If the identity of a principal is essential to the meaning of a message, it is prudent to mention the principal's name explicitly in the message." [3]

This practice is heeded in many protocols, but unfortunately, by sending the names of the principals and usually the certificates in plaintext, the protocols reveal the identity of the principals to even passive observers. This section presents protocols that are designed to protect the identity of the authentication participants, in addition to providing confidentiality of content.

The Security Architecture for Internet Protocol (IPsec) [47] working group in the IETF has inspired many authentication and key exchange protocols that provide identity protection. The earliest is the SKEME [50] protocol. SKEME provides identity protection against observers not participating in the protocol by encrypting the identity of the initiator with the responders public key. Thus, the initiator needs to reveal its identity to potential responders first. The IETF designed the IKE protocol [41], which provides identity protection in the main mode similar to SKEME, but unfortunately uses many roundtrips. The designers of Just Fast Keying (JFK) [6, 7] designed two protocols, JFKi and JFKr, that take two round-trips and provide the same level of denial of service protection. JFKi provides identity protection for the initiator against active attacks, while the JFKr provides identity protection for the responder and protects both parties' identities against passive observers. Today IKE specification has been made obsolete by IKEv2 [44], but still the initiator has to prove its identity first. Many protocols, however, do not provide identity protection at all.

As a side note, the IPsec architecture has always been surrounded with controversy. Encrypting the packet payload at the IP layer was seen as too drastic, and therefore the IETF needed to specify the IP Authentication Header (AH) [45], which does not provide confidentiality for the content, but only integrity. This could have also been established using the null encryption mode with Encapsulating Security Payload (ESP) [46], but the IETF opted to specify a header for the sole purpose of authentication. On the other hand, AH does protect also the headers of the IP packet.

Independent of IPsec related work, Abadi proposed two protocols for private authentication that protect the identities of both parties, and later formally proved the privacy properties [1, 2]. One of the important notes Abadi makes is that the used public key encryption protocol used needs to be "which-key concealing" or "key-private" [15]. This property guarantees that, if the attacker sees the ciphertext, he or she cannot tell which specific key, out of a set of known public keys, was used to create the specific ciphertext.

There are also special protocols designed to limit the exposure of identity in specific contexts. Secret handshake protocols [9, 14] provide a way to authenticate membership in a group and role in it, while a third-party observing the exchange does not learn

anything new (including whether the users doing the handshake belong to the same group, the identities of the groups or the roles of the users). However, these protocols have been asymmetric so far, in the sense that it is possible for a user belonging to a group to learn whether another user belongs to the same group and not reveal anything of herself by aborting the protocol at the appropriate time. Other schemes of interest are anonymous credentials [19, 21] for authorizing pseudonyms to access a system, short group signatures [17] for hiding who in a group signed a message, and secret sets [55] for providing sets where anybody can test for their own membership, but only the creator of the set can test another party's membership in the group she created.

2.1.4 Data transit

The payload of communication protocols or "data transit" will reveal the communication content if it is not protected. It may also help to identify the communication participants and the location of the participants. The simplest protection here is to encrypt the data traffic, with the help of authentication and key exchange protocols discussed above. However, there are also inference attacks against encrypted traffic. For example, there are timing attacks against SSH [71] and HTTP [35], statistical attacks against HTTP. [16], and variable-bit-rate encoding used in voice-over-IP and movie streaming introduces vulnerabilities even with encryption [68].

2.2 Control Plane Choke-Points

TO BE INCLUDED

Chapter 3

Conclusions

In this paper, we have aimed to identify the architectural choke-points for privacy in the current Internet architecture. We first discussed how do we approach privacy, by defining confidentially of communication, participants and location. We discussed how observers or attackers located in different locations along the communication path can violate the privacy of Internet users.

We have divided the privacy choke-points for data plane and control plane. The data plane choke-points include naming, name lookup, authentication and key exchange, and data transit. [discussion of control plane choke-points to be included]

We hope this document will help with the design and discussion of novel Internet architectures.

References

[1] Martín Abadi. Private authentication. In *Second International Workshop on Privacy Enhancing Technologies (PET 2002)*, April 2002. 11

[2] Martín Abadi and Cédric Fournet. Private authentication. *Theor. Comput. Sci.*, 322(3):427–476, September 2004. 11

[3] Martín Abadi and Roger Needham. Prudent engineering practice for cryptographic protocols. *IEEE Transactions on Software Engineering*, January 1996. 11

[4] Dakshi Agrawal and Dogan Kesdogan. Measuring anonymity: The disclosure attack. *IEEE Security & Privacy*, 1, November–December 2003. 10

[5] Dakshi Agrawal, Dogan Kesdogan, and Stefan Penz. Probabilistic Treatment of MIXes to Hamper Traffic Analysis. In *IEEE Symposium on Security and Privacy*, May 2003. 10

[6] William Aiello, Steven M. Bellovin, Matt Blaze, Ran Canetti, John Ioannidis, Angelos D. Keromytis, and Omer Reingold. Just fast keying: Key agreement in a hostile internet. *ACM Transactions on Information and System Security (TISSEC)*, 7, May 2004. 11

[7] William Aiello, Steven M. Bellovin, Matt Blaze, John Ioannidis, Omer Reingold, Ran Canetti, and Angelos D. Keromytis. Efficient, DoS-resistant, secure key exchange for internet protocols. In *9th ACM conference on Computer and Communications Security*, November 2002. 11

[8] Jari Arkko, Pekka Nikander, and Mats Näslund. Enhancing Privacy with Shared Pseudo Random Sequences. In *Proc. of Security Protocols*, April 2005. 8

[9] Giuseppe Ateniese, Marina Blanton, and Jonathan Kirsch. Secret handshakes with dynamic and fuzzy matching. In *Proc. of NDSS '07*, February 2007. 11

[10] Tuomas Aura, Janne Lindqvist, Michael Roe, and Anish Mohammed. Chattering laptops. In *8th Privacy Enhancing Technologies Symposium (PETS)*, July 2008. 10

[11] Tuomas Aura and Michael Roe. Designing the Mobile IPv6 Security Protocol. *Annales des télécommunications / Annals of telecommunications, special issue on Network and information systems security*, 61(3-4), March-April 2006. 8

[12] Tuomas Aura and Alf Zugenmaier. Privacy, Control and Internet Mobility. In *Security Protocols, 12th International Workshop*, April 2004. 8

[13] Adam Back, Ulf Möller, and Anton Stiglic. Traffic analysis attacks and trade-offs in anonymity providing systems. In *Information Hiding Workshop (IH 2001) LNCS (2137)*, April 2001. 10

[14] Dirk Balfanz, Glenn Durfee, Narendar Shankar, Diana Smetters, Jessica Staddon, and Hao-Chi Wong. Secret handshakes from pairing-based key agreements. In *Proc. of IEEE Security and Privacy*, May 2003. 11

[15] Mihir Bellare, Alexandra Boldyreva, Anand Desai, and David Pointcheval. Key-privacy in public-key encryption. In *ASIACRYPT*, December 2001. 11

[16] G. D. Bissias, M. Liberatore, D. Jensen, and B. N. Levine. Privacy vulnerabilities in encrypted http streams. In *Privacy Enhancing Technologies (PET), LNCS 3856*, May 2005. 12

[17] Dan Boneh, Xavier Boyen, and Hovav Shacham. Short group signatures. In *Proc. of Crypto*, August 2004. 12

[18] Nikita Borisov, Ian Goldberg, and Eric Brewer. Off-the-record communication, or, why not to use PGP. In *WPES '04: Proceedings of the 2004 ACM workshop on Privacy in the electronic society*, pages 77–84, New York, NY, USA, October 2004. ACM Press.

[19] Jan Camenisch and Anna Lysyanskaya. Dynamic Accumulators and Application to Efficient Revocation of Anonymous Credentials. In *CRYPTO*, August 2002. 12

[20] David Chaum. Untraceable electronic mail, return addresses, and digital pseudonyms. *Commun. ACM*, 24(2):84–90, February 1981. 8, 9

[21] David Chaum. Security without identification: transaction systems to make big brother obsolete. *Commun. ACM*, 28(10):1030–1044, 1985. 12

[22] David Chaum. The dining cryptographers problem: Unconditional sender and recipient untraceability. *Journal of Cryptology*, 1, 1988. 8, 9

[23] Ian Clarke, Scott G. Miller, Theodore W. Hong, Oskar Sandberg, and Brandon Wiley. Protecting free expression online with freenet. *IEEE Internet Computing*, 6(1):40–49, 2002. 9

[24] Ian Clarke, Oskar Sandberg, Brandon Wiley, and Theodore W. Hong. Freenet: A Distributed Anonymous Information Storage and Retrieval System. In *Designing Privacy Enhancing Technologies: International Workshop on Design Issues in Anonymity and Unobservability, LNCS 2009*, July 2001. 9

[25] R. Clayton, G. Danezis, and M. G. Kuhn. Real world patterns of failure in anonymity systems. In *Information Hiding Workshop (IH 2001), LNCS 2137*, April 2001.

[26] George Danezis. Statistical disclosure attacks: Traffic confirmation in open environments. In *Proceedings of Security and Privacy in the Age of Uncertainty, (SEC2003)*, May 2003. 10

[27] George Danezis and Andrei Serjantov. Statistical disclosure or intersection attacks on anonymity systems. In *Proceedings of 6th Information Hiding Workshop (IH 2004)*, May 2004. 10

[28] S. Deering and R. Hinden. RFC 2460: Internet Protocol, Version 6 (IPv6) Specification, December 1998. Status: Draft Standard. 7

[29] Claudia Díaz, Stefaan Seys, Joris Claessens, and Bart Preneel. Towards measuring anonymity. In *Proceedings of Privacy Enhancing Technologies Workshop (PET 2002), LNCS 2482*, April 2002. 10

[30] Whitfield Diffie and Susan Landau. *Privacy on the Line: The Politics of Wiretapping and Encryption – Updated and Expanded Edition*. The MIT Press, May 2007. ISBN 978-0-262-04240-6. 4

[31] Roger Dingledine and Nick Mathewson. Anonymity Loves Company: Usability

and the Network Effect. In *Workshop on the Economics of Information Security*, June 2006. 10

[32] Roger Dingledine, Nick Mathewson, and Paul Syverson. Tor: The Second-Generation Onion Router. In *Proceedings of the 13th USENIX Security Symposium*, August 2004. 4, 9

[33] Roger Dingledine, Nick Mathewson, and Paul Syverson. Deploying low-latency anonymity: Design challenges and social factors. *IEEE Security and Privacy*, 5(5):83–87, 2007. 9

[34] Alberto Escudero-Pascual. *Privacy in the next generation Internet: Data protection in the context of European Union policy*. PhD thesis, Royal Institute of Technology, 2002. 8

[35] E. W. Felten and M. A. Schneider. Timing attacks on web privacy. In *Proceedings of the 7th ACM Conference on Computer and Communications Security (CCS)*, November 2000. 12

[36] David Goldschlag, Michael Reed, and Paul Syverson. Onion routing. *Commun. ACM*, 42(2):39–41, 1999. 9

[37] David M. Goldschlag, Michael G. Reed, and Paul F. Syverson. Hiding Routing Information. In *Workshop on Information (LNCS 1174)*, 1996. 8, 9

[38] Philippe Golle and Kurt Partridge. On the anonymity of home/work location pairs. In *Proceedings of the 7th International Conference on Pervasive Computing*, Pervasive '09, pages 390–397, Berlin, Heidelberg, 2009. Springer-Verlag. 5

[39] Ralph Gross and Alessandro Acquisti. Information revelation and privacy in online social networks. In *Proceedings of the 2005 ACM workshop on Privacy in the electronic society*, WPES '05, pages 71–80, New York, NY, USA, 2005. ACM. 3

[40] Saikat Guha and Paul Francis. Identity Trail: Covert Surveillance Using DNS. In *Workshop on Privacy Enhancing Technologies (PET)*, June 2007. 8

[41] Dan Harkins and Dave Carrell. The Internet Key Exchange (IKE), November 1998. 11

[42] Sabine Helmers. A brief history of anon.penet.fi - the legendary anonymous remailer. *Computer-Mediated Communication Magazine (CMC)*, September 1997. 8

[43] Robert Hinden and Stephen Deering. RFC 4291: IP Version 6 Addressing Architecture, February 2006. Status: Draft Standard. 7

[44] Charlie Kaufman. RFC 4306: Internet Key Exchange (IKEv2) Protocol, December 2005. 11

[45] Stephen Kent. RFC 4302: IP Authentication Header (AH), December 2005. 11

[46] Stephen Kent. RFC 4303: IP Encapsulating Security Payload (ESP), December 2005. 8, 11

[47] Stephen Kent and Karen Seo. RFC 4301: Security Architecture for the Internet Protocol, December 2005. 11

[48] Dogan Kesdogan, Dakshi Agrawal, and Stefan Penz. Limits of anonymity in open environments. In *Information Hiding Workshop (IH 2002) LNCS (2578)*, October

2002. 10

[49] Rajeev Koodli. RFC 4882: IP Address Location Privacy and Mobile IPv6: Problem Statement, May 2007. 8

[50] Hugo Krawczyk. SKEME: a versatile secure key exchange mechanism for Internet. In *Symposium on Network and Distributed System Security (NDSS)*, February 1996. 11

[51] Brian N. Levine, Michael K. Reiter, Chenxi Wang, and Matthew K. Wright. Timing attacks in low-latency mix-based systems. In *Proceedings of Financial Cryptography (FC '04)*, February 2004. 10

[52] Janne Lindqvist. *Practical Privacy Enhancing Technologies for Mobile Systems*. Helsinki University of Technology, 2009. ISBN 978-951-22-9902-7. 3

[53] Janne Lindqvist, Tuomas Aura, George Danezis, Teemu Koponen, Annu Myllyniemi, Jussi Mäki, and Michael Roe. Privacy-preserving 802.11 access-point discovery. In *Proceedings of the second ACM conference on Wireless network security*, WiSec '09, pages 123–130, New York, NY, USA, 2009. ACM. 6

[54] Nick Mathewson and Roger Dingledine. Practical traffic analysis: Extending and resisting statistical disclosure. In *Proceedings of Privacy Enhancing Technologies workshop (PET 2004)*, May 2004. 10

[55] Refik Molva and Gene Tsudik. Secret sets and applications. *Information Processing Letters*, 65, 1998. 12

[56] Steven J. Murdoch and George Danezis. Low-Cost Traffic Analysis of Tor. In *IEEE Symposium on Security and Privacy*, May 2005. 10

[57] Steven J. Murdoch and Piotr Zielinski. Sampled traffic analysis by internet-exchange-level adversaries. In *Proc. of 7th Workshop on Privacy Enhancing Technologies*, June 2007. 6

[58] Thomas Narten and Richard Draves. RFC 3041: Privacy Extensions for Stateless Address Autoconfiguration in IPv6, January 2001. 7

[59] Thomas Narten, Richard Draves, and Suresh Krishnan. RFC 4941: Privacy Extensions for Stateless Address Autoconfiguration in IPv6, September 2007. Status: Draft Standard. 8

[60] Lasse Overlier and Paul Syverson. Locating hidden servers. In *IEEE Symposium on Security and Privacy*, May 2006. 10

[61] Andreas Pfitzmann and Marit Hansen. Anonymity, unlinkability, unobservability, pseudonymity, and identity management – a consolidated proposal for terminology, February 2008. Version v0.31 http://dud.inf.tu-dresden.de/Anon_Terminology.shtml. 5

[62] Jon Postel. RFC 791: Internet Protocol, September 1981. Status: Standard. 7

[63] Jean-François Raymond. Traffic analysis: protocols, attacks, design issues, and open problems. In *International workshop on Designing privacy enhancing technologies*, pages 10–29, New York, NY, USA, 2001. Springer-Verlag New York, Inc. 10

[64] Michael K. Reiter and Aviel D. Rubin. Crowds: Anonymity for Web Transactions.

ACM Transactions on Information and System Security, 1(1):66–92, November 1998. 9, 10

[65] Michael K. Reiter and Aviel D. Rubin. Anonymous Web Transactions with Crowds. *Communications of the ACM*, 42, February 1999. 9

[66] Alexey Reznichenko, Saikat Guha, and Paul Francis. Auctions in do-not-track compliant internet advertising. In *Proceedings of the 18th ACM conference on Computer and communications security*, CCS '11, pages 667–676, New York, NY, USA, 2011. ACM. 3

[67] Jonathan Rosenberg, Henning Schulzrinne, Gonzalo Camarillo, Alan Johnston, Jon Peterson, Robert Sparks, Mark Handley, and Eve Schooler. RFC 3261: SIP: Session Initiation Protocol, June 2002. 8

[68] T. Scott Saponas, Jonathan Lester, Carl Hartung, Sameer Agarwal, and Tadayoshi Kohno. Devices That Tell On You: Privacy Trends in Consumer Ubiquitous Computing. In *Proc. of USENIX Security*, August 2007. 12

[69] Andrei Serjantov and George Danezis. Towards an information theoretic metric for anonymity. In *Privacy Enhancing Technologies Workshop (PET 2002)*, April 2002. 10

[70] Daniel J. Solove. A taxonomy of privacy. *University of Pennsylvania Law Review*, 154(3), January 2006. 5

[71] D. X. Song, D. Wagner, and X. Tian. Timing analysis of keystrokes and timing attacks on ssh. In *In Proceedings of the 10th USENIX Security Symposium*, August 2001. 12

[72] Latanya Sweeney. k-Anonymity: a model for protecting privacy. *International Journal on Uncertainty, Fuzziness and Knowledge-based Systems*, 10, 2002. 3, 10

[73] Paul Syverson, Gene Tsudik, Michael Reed, and Carl Landwehr. Towards an analysis of onion routing security. In *International workshop on Designing privacy enhancing technologies*, pages 96–114, New York, NY, USA, 2001. Springer-Verlag New York, Inc. 9, 10

[74] Paul F. Syverson, David M. Goldschlag, , and Michael G. Reed. Anonymous connections and onion routing. In *IEEE Symposium on Security and Privacy*, May 1997. 8, 9

[75] Susan Thomson and Thomas Narten. RFC 1971: IPv6 Stateless Address Autoconfiguration, August 1996. Status: Obsolete. 7

[76] Susan Thomson, Thomas Narten, and Tatuya Jinmei. RFC 4862: IPv6 Stateless Address Autoconfiguration, September 2007. Status: Draft Standard. 8

[77] Matthew Wright, Micah Adler, Brian N. Levine, and Clay Shields. An analysis of the degradation of anonymous protocols. In *Network and Distributed Security Symposium (NDSS)*, February 2002. 10

[78] Matthew Wright, Micah Adler, Brian N. Levine, and Clay Shields. Defending anonymous communication against passive logging attacks. In *IEEE Symposium on Security and Privacy*, May 2003. 10

[79] Matthew K. Wright, Micah Adler, Brian Neil Levine, and Clay Shields. The

predecessor attack: An analysis of a threat to anonymous communications systems. *ACM Trans. Inf. Syst. Secur.*, 7(4):489–522, November 2004. 10

[80] Matthew K. Wright, Micah Adler, Brian Neil Levine, and Clay Shields. Passive-logging attacks against anonymous communications systems. *ACM Trans. Inf. Syst. Secur.*, 11(2):1–34, March 2008. 10