# Privacy-enhancing technologies for the Internet

Ian Goldberg David Wagner Eric Brewer
*University of California, Berkeley*
{iang,daw,brewer}@cs.berkeley.edu

## Abstract

*The increased use of the Internet for everyday activities is bringing new threats to personal privacy. This paper gives an overview of existing and potential privacy-enhancing technologies for the Internet, as well as motivation and challenges for future work in this field.*

## 1. Introduction

Recently the Internet has seen tremendous growth, with the ranks of new users swelling at ever-increasing rates. This expansion has catapulted it from the realm of academic research towards new-found mainstream acceptance and increased social relevance for the everyday individual. Yet this suddenly increased reliance on the Internet has the potential to erode personal privacies we once took for granted.

New users of the Internet generally do not realize that every post they make to a newsgroup, every piece of email they send, every World Wide Web page they access, and every item they purchase online could be monitored or logged by some unseen third party. The impact on personal privacy is enormous; already we are seeing databases of many different kinds, selling or giving away collections of personal data, and this practice will only become more common as the demand for this information grows.

All is not lost. While the Internet brings the danger of diminished privacy, it also ushers in the potential for expanding privacy protection to areas where privacy was previously unheard of. This is our vision: restoration and revitalization of personal privacy for online activities, and betterment of society via privacy protection for fields where that was previously impossible. We want to bring privacy to the Internet, and bring the Internet to everyday privacy practices.

The purpose of this paper is not to present new results, but rather to encourage further research in the area of Internet privacy protection, and to give an overview (necessarily brief in a short paper such as this) of privacy-enhancing technologies. Section 2 explores some motivation for studying privacy issues on the Internet, and Section 3 provides some relevant background. We then discuss Internet privacy technology chronologically, in three parts: Section 4 describes the technology of yesterday, Section 5 explains today's technology, and Section 6 explores the technology of tomorrow. Finally, we conclude in Section 7.

## 2. Motivation

The threats to one's privacy on the Internet are two-fold: your online actions could be (1) monitored by unauthorized parties and (2) logged and preserved for future access many years later. You might not realize that your personal information has been monitored, logged, and subsequently disclosed; those who would compromise your privacy have no incentive to warn you.

The threat of long-term storage and eventual disclosure of personal information is especially acute on the Internet. It is technically quite easy to collect information (such as a compendium of all posts you have made to electronic newsgroups) and store it for years or decades, indexed by your name for easy retrieval. If you are looking for a job twenty years from now, do you want your employer to browse through every Usenet posting you've ever made? If you are like most people, you have probably said something (however minor) in your past you would prefer to forget—perhaps an incautious word from your indiscreet youth, for instance. Long-term databases threaten your ability to choose what you would like to disclose about your past.

Furthermore, in recent years great advances have been made in technology to mine the Internet for interesting information. This makes it easy to find and extract personal information about you that you might not realize is available. (For instance, one of your family members might have listed information about you on their web page without your knowledge; Internet search engine technology would find this easily.) Did you know your phone number, email address, and street address are probably listed on the Web? Or that your social security number is available on any of several for-pay electronically-searchable databases? Most people probably do not want to make it easy for salesmen, telemarketers, an abusive ex, or a would-be stalker to find them.

In these ways, the Internet contributes to the "dossier effect", whereby a single query can compile a huge dossier containing extensive information about you from many diverse sources. This increasingly becomes a threat as databases containing personal information become electronically cross-linked more widely. A recent trend is to make more databases accessible from the Internet; with today's powerful search engine and information-mining technology, this is one of the ultimate forms of cross-linking. (For instance, phone directories, address information, credit reports, newspaper articles, and public-access government archives are all becoming available on the Internet.) The "dossier effect" is dangerous: when it is so easy to build a comprehensive profile of individuals, many will be tempted to take advantage of it, whether for financial gain, vicarious entertainment, illegitimate purposes, or other unauthorized use.

Government is one of the biggest consumers and producers of dossiers of personal information, and as such should be viewed as a potential threat to privacy. The problem is that today's governments have many laws, surveillance agencies, and other tools for extracting private information from the populace [6]. Furthermore, a great many government employees have access to this valuable information, so there are bound to be some workers who will abuse it. There are many examples of small-scale abuses by officials: a 1992 investigation revealed that IRS employees at just one regional office made hundreds of unauthorized queries into taxpayer databases [2]; employees of the Social Security Administration have been known to sell confidential government records for bribes as small as $10 [22]; highly confidential state records of AIDS patients have leaked [3]. Finally, there is very little control or oversight, so an corrupt leader could easily misuse this information to seize and maintain power. A number of cautionary examples are available: FBI Director Edgar Hoover had his agency spy on political dissidents, activists, and opponents; the NSA, a secret military surveillance agency, has a long history of spying on domestic targets [5]; President Clinton's Democratic administration found themselves with unauthorized secret dossiers on hundreds of Republican opponents in the "Filegate" scandal.

Anonymity is one important form of privacy protection that is often useful.

We observe that anonymity is often used not for its own sake, but primarily as a means to an end, or as a tool to achieve personal privacy goals. For example, if your unlisted telephone number is available on the web, but can't be linked to your identity because you have used anonymity tools, then this might be enough to fulfill your need for privacy just as effectively as if you had kept the phone number completely secret. Many applications of online anonymity follow the common theme of "physical security through anonymity". For instance, political dissidents living in totalitarian regimes might publish an exposé anonymously on the Internet to avoid harassment (or worse!) by the secret police.

In contexts other than the Internet, anonymous social interaction is both commonplace and culturally accepted. For example, the Federalist papers were penned under the pseudonym Publius; many other well-known literary works, such as *Tom Sawyer*, *Primary Colors*, etc. were also written anonymously or under a pseudonym. Today, home HIV tests rely on anonymous lab testing; police tip lines provide anonymity to attract informants; journalists take great care to protect the anonymity of their confidential sources; and there is special legal protection and recognition for lawyers to represent anonymous clients. The US Postal Service accepts anonymous mail without prejudice; it is well-known that anonymous voice calls can be easily made by stepping into a payphone; and ordinary cash allows everyday people to purchase merchandise (say, a copy of Playboy) anonymously. In short, most non-Internet technology today grants the ordinary person access to anonymity. Outside of the Internet, anonymity is widely accepted and recognized as valuable in today's society. Long ago we as a society reached a policy decision, which we have continually reaffirmed, that there are good reasons to protect and value anonymity off the Internet; that same reasoning applies to the Internet, and therefore we should endeavor to protect online anonymity as well.

There are many legitimate uses for anonymity on the Internet. In the long term, as people take activities they'd normally do offline to the Internet, they will expect a similar level of anonymity. In fact, in many cases, they won't even be able to imagine the extensive use this data could be put to by those with the resources and incentive to mine the information in a less-than-casual way. We should protect the ordinary user rather than requiring them to anticipate the various ways their privacy could be compromised. Moreover, the nature of the Internet may even make it possible to exceed those expectations and bring anonymity to practices where it was previously nonexistent. In the short term, there are a number of situations where we can already see (or confidently predict) legitimate use of Internet anonymity: support groups (e.g. for rape survivors or recovering alcoholics), online tip lines, whistleblowing, political dissent, refereeing for academic conferences, and merely the pursuit of everyday privacy of a less noble and grand nature. As the New Yorker magazine explained in a famous cartoon, "On the Internet, nobody knows you're a dog"[23]—and this is perhaps one of the greatest strengths of the Internet.

On the other hand, illicit use of anonymity is all too common on the Internet. Like most technologies, Internet anonymity techniques can be used for better or worse, so it should not be surprising to find some unfavorable uses of anonymity. For instance, sometimes anonymity tools are used to distribute copyrighted software without permission ("warez"). Email and Usenet spammers are learning to take advantage of anonymity techniques to distribute their marketing ploys widely without retribution. Denial of service and other malicious attacks are likely to become a greater problem when the Internet infrastructure allows wider support for anonymity. The threat of being tracked down and dealt with by social techniques currently acts as a partial deterrent to would-be intruders, but this would be eroded if they could use Internet tools to hide their identity. We have already seen one major denial of service attack [10] where the attacker obscured his IP source address to prevent tracing. Widespread availability of anonymity will mean that site administrators will have to rely more on first-line defenses and direct security measures rather than on the deterrent of tracing. Providers of anonymity services will also need to learn to prevent and manage abuse more effectively. These topics are discussed at greater length in later sections.

## 3. Background

A few definitions are in order. Privacy refers to the ability of the individual to protect information about himself. Anonymity is privacy of identity. We can divide anonymity into two cases: persistent anonymity (or pseudonymity), where the user maintains a persistent online persona ("nym") which is not connected with the user's physical identity ("true name"), and one-time anonymity, where an online persona lasts for just one use. The key concept here is that of linkability: with a nym, one may send a number of messages that are all linked together but cannot be linked to the sender's true name; by using one-time anonymity for each message, none of the messages can be linked to each other or to the user's physical identity.[1] Forward secrecy refers to the inability of an adversary to recover security-critical information (such as the true name of the sender of a controversial message) "after the fact" (e.g. after the message is sent); providers of anonymity services should take care to provide forward secrecy, which entails (for instance) keeping no logs.

Some of the more obvious uses of persistent anonymity are in "message-oriented" services, such as email and newsgroup postings. Here, the two major problems to be solved are those of sender-anonymity, where the originator of a message wishes to keep his identity private, and of recipient-anonymity, where we wish to enable replies to a persistent persona.

In contrast to "message-oriented" services, we have "online" services. In these services, which include the World-Wide Web, online chat rooms, phones, videoconferences, and most instances of electronic commerce, we wish to enable two parties to communicate in real time, while allowing one or both of them to maintain their anonymity. The added challenges for online services stem from the increased difficulty involved in sending low-latency information without revealing identity via timing coincidences; to support these online services, we want to erect a general-purpose low-level infrastructure for anonymous Internet communications. In addition, certain specific applications, such as private electronic commerce, require sophisticated application-level solutions.

## 4. Past

In past years email was the most important distributed application, so it should not be surprising that early efforts at bringing privacy to the Internet primarily concentrated on email protection. Today the lessons learned from email privacy provide a foundation of practical experience that is critically relevant to the design of new privacy-enhancing technologies.

The most primitive way to send email anonymously involves sending the message to a trusted friend, who deletes the identifying headers and resends the message body under his identity. Another old technique for anonymous email takes advantage of the lack of authentication for email headers: one connects to a mail server and forges fake headers (with falsified identity information) attached to the message body. (Both approaches could also be used for anonymous posting to newsgroups.) Of course, these techniques don't scale well, and they offer only very minimal assurance of protection.

The technology for email anonymity took a step forward with the introduction of anonymous remailers. An anonymous remailer can be thought of as a mail server which combines the previous two techniques, but using a computer to automate the header-stripping and resending process [4, 16, 17, 24]. There are basically three styles of remailers; we classify remailer technology into "types" which indicate the level of sophistication and security.

The `anon.penet.fi` ("type 0") remailer was perhaps the most famous. It supported anonymous email senders by stripping identifying headers from outbound remailed messages. It also supported recipient anonymity: the user was assigned a random pseudonym at `anon.penet.fi`, the remailer maintained a secret identity table matching up the user's real email address with his `anon.penet.fi` nym, and incoming email to the nym at `anon.penet.fi`

[1] Users of anonymity services should keep in mind that messages written by the same person tend to share certain characteristics, and that this fact has been used to identify the authors of anonymous works in the past.

was retransmitted to the user's real email address. Due to its simplicity and relatively simple user interface, the `anon.penet.fi` remailer was the most widely used remailer; sadly, it was shut down recently after being harassed by legal pressure [18].

The disadvantage of a `anon.penet.fi` style (type 0) remailer is that it provides rather weak security. Users must trust it not to reveal their identity when they send email through it. Worse still, pseudonymous users must rely on the confidentiality of the secret identity table—their anonymity would be compromised if it were disclosed, subpoenaed, or bought—and they must rely on the security of the `anon.penet.fi` site to resist intruders who would steal the identity table. Furthermore, more powerful attackers who could eavesdrop on Internet traffic traversing the `anon.penet.fi` site could match up incoming and outgoing messages to learn the identity of the nyms.

Cypherpunk-style (type I) remailers were designed to address these types of threats. First of all, support for pseudonyms is dropped; no secret identity table is maintained, and remailer operators take great care to avoid keeping mail logs that might identify their users. This diminishes the risk of "after-the-fact" tracing. Second, type I remailers will accept encrypted email, decrypt it, and remail the resulting message. (This prevents the simple eavesdropping attack where the adversary matches up incoming and outgoing messages.) Third, they take advantage of *chaining* to achieve more robust security. Chaining is simply the technique of sending a message through several anonymous remailers, so that the second remailer sees only the address of the first remailer and not the address of the originator, etc. Typically one combines chaining with encryption: the originator encrypts repeatedly, nesting once for each remailer in the chain; the advantage is that every remailer in a chain must be compromised before a chained message can be traced back to its sender. This allows us to take advantage of a distributed collection of remailers; diversity gives one a better assurance that at least some of the remailers are trustworthy, and chaining ensures that one honest remailer (even if we don't know which it is) is all we need. Type I remailers can also randomly reorder outgoing messages to prevent correlations of ciphertexts by an eavesdropper. In short, type I remailers offer greatly improved security over type 0, though they do have some limitations which we will discuss next.

## 5. Present

The newest and most sophisticated remailer technology is the Mixmaster, or type II, remailer [7, 11]. They extend the techniques used in a type I remailer to provide enhanced protection against eavesdropping attacks. First, one always uses chaining and encryption at each link of the chain. Second, type II remailers use constant-length messages, to prevent passive correlation attacks where the eavesdropper matches up incoming and outgoing messages by size. Third, type II remailers include defenses against sophisticated replay attacks. Finally, these remailers offer improved message reordering code to stop passive correlation attacks based on timing coincidences. Because their security against eavesdropping relies on "safety in numbers" (where the target message cannot be distinguished from any of the other messages in the remailer net), the architecture also calls for continuously-generated random cover traffic to hide the real messages among the random noise.

Another new technology is that of the "newnym"-style nymservers. These nymservers are essentially a melding of the recipient anonymity features of a `anon.penet.fi` style remailer with the chaining, encryption, and other security features of a cypherpunk-style remailer: a user obtains a pseudonym (e.g. `joeblow@nym.alias.net`) from a nymserver; mail to that pseudonym will be delivered to him. However, unlike `anon.penet.fi`, where the nymserver operator maintained a list matching pseudonyms to real email addresses, newnym-style nymservers only match pseudonyms to "reply blocks": the nymserver operator does not have the real email address of the user, but rather the address of some type I remailer, and an encrypted block of data which it sends to that remailer. When decrypted, that block contains the address of a second remailer, and more encrypted data, etc. Eventually, when some remailer decrypts the block it receives, it gets the real email address of the user. The effect is that all of the remailers mentioned in the reply block would have to collude or be compromised in order to determine the email address associated with a newnym-style pseudonym.

Another simple technique for recipient anonymity uses message pools. Senders encrypt their message with the recipient's public key and send the encrypted message to a mailing list or newsgroup (such as `alt.anonymous.messages`, set up specifically for this purpose) that receives a great deal of other traffic. The recipient is identified only as someone who reads the mailing list or newsgroup, but onlookers cannot narrow down the identity of the recipient any further. A "low-tech" variant might use classified advertisements in a widely-read newspaper such as The New York Times. Message pools provide strong recipient anonymity, but of course the huge disadvantage is that they waste large amounts of bandwidth and pollute mailing lists with bothersome noise.

With the increasing sophistication in remailer technology, we find that modern remailers have been burdened with a correspondingly complicated and obscure interface. To deal with this unfriendly mess, client programs have sprung up to provide a nicer interface to the remailers. Raph Levien's `premail` [21] is the archetypical example. Even so, using remailers still requires some knowledge; for even greater user-friendliness, we need this support to be integrated into popular mail handling applications.

One could reasonably argue that the problem of anonymous email is nearly solved, in the sense that we largely understand most of the principles of building systems to provide email anonymity. However, email is not the only important application on the Internet. More recently, we have begun to see privacy support for other services as well.

The "strip identifying headers and resend" approach used by remailers has recently been applied to provide anonymity protection for Web browsing as well. Community ConneXion has sponsored the Anonymizer [9], a web proxy that filters out identifying headers and source addresses from the web browser. This allowing users to surf the web anonymously without revealing their identity to web servers. However, the Anonymizer offers rather weak security—no chaining, encryption, log safeguarding, or forward secrecy—so its security properties are roughly analogous to those of type 0 remailers. Other implementations have since appeared based on the same approach [12, 15]; but technology for anonymous web browsing remains relatively unsophisticated and underdeveloped.

Finally, anonymous digital cash is another state-of-the-art technology for Internet privacy. As many observers have stressed, electronic commerce will be a driving force for the future of the Internet. Therefore, the emergence of digital commerce solutions with privacy and anonymity protection is very valuable. DigiCash's ecash [8] has the strongest privacy protection of any deployed payment system—it uses sophisticated cryptographic protocols to guarantee that the payer's privacy is not compromised by the payment protocol even against a colluding bank and payee. Thus, DigiCash's ecash has many of the privacy properties of real cash; most other deployed payment systems have only about as much privacy as checks or credit cards.

Of course, the DigiCash protocols only prevent your identity from being revealed by the protocols themselves: if you send the merchant a delivery address for physical merchandise, he will clearly be able to identify you. Similarly, if you use pay using ecash over a non-anonymized IP connection, the merchant will be able to deduce your IP address. This demonstrates the need for a general-purpose infrastructure for anonymous IP traffic, as discussed later. (The other option is to pay by email, with which you can use the existing remailer infrastructure, to preserve your privacy.) In any case, security is only as strong as the weakest link in the chain, and we need strong anonymity (such as provided by DigiCash's protocols) in our payment system as well as strong anonymity in our data transport infrastructure.

DigiCash's anonymous ecash does have a few limitations. Like the telephone or the fax machine, its success depends on seeing widespread adoption by a large number of customers and merchants; but so far it has merely a relatively small user base. Also, it currently offers only one-way anonymity—namely, anonymity for payers but not for payees—so parties who wish to sell services or information anonymously are currently not served well by DigiCash's ecash. Nonetheless, improvements are still being made, and DigiCash is a important pioneer in this crucial area.

## 6. Future

The first author has made significant progress on working around the limitations of DigiCash's ecash. His enhancements attempt to stimulate growth in the user base by making it easy to use ecash without signing up for an account at a participating bank (thus eliminating paperwork). Additionally, he developed support for currency trading and e-cashiering, where service providers may offer to buy or sell DigiCash ecash in exchange for other forms of payment. His improvements also include bi-directional anonymity to support change-making and anonymous merchants, and a Netscape plug-in to make payment more transparent. These improvements are compatible with DigiCash's system—users can take advantage of his enhancements without any changes to the bank's software.

When attempting to design anonymity support for web traffic, interactive text/voice/video chatting, remote `telnet` connections, and other similar services, we quickly see that what we need is an infrastructure to provide bi-directional anonymity protection for general-purpose low-latency interactive Internet traffic. Wei Dai has described an architecture that would provide this protection based on a distributed system of anonymizing packet forwarders, analogous to today's remailer network; he called it "Pipenet" [13]. We will use the generic term pipenet for any architecture built along these lines.

No complete pipenet design, much less implementation, is available yet. Several authors have independently attempted to build a system with similar features [26], but because they were unaware of the work of Wei Dai [14] and other cypherpunks, their design remains vulnerable to a number of attacks. Due to space limitations, we cannot give a full list of threats and attacks in this paper; we will merely confine ourself with observing that pipenet must protect against all of the attacks against remailers discussed above, as well as some others specific to low-latency long-lived connections. A future paper will discuss these threats in detail and give a number of possible countermeasures. We hope that the great applicability of a general-purpose infrastructure for anonymized Internet traffic will motivate and stimulate new research in this area.

Another great challenge that faces future researchers in Internet privacy technology is the problem of abuse. As tools and infrastructure for anonymity become available, some will abuse these resources for illicit purposes.

We have some experience with handling abuse from the deployed remailers. Abuse only accounts for a small minority of remailer usage, but it is typically much more visible. One of the most common abuses of remailers is junk email, where senders hide behind anonymity to send vast quantities of unsolicited email (usually advertising) to a large number of recipients who find it unwelcome. Remailers today include simplistic alarms when they encounter a large volume of mail in a short time; then remailer operators can delete the spammed messages and source block the spammer (i.e. blacklist the sender). Harassment of a targeted individual is another common abuse of anonymous remailers. One countermeasure is to have targeted individuals install mail filtering software. (Remailers could also provide destination blocking services, but this raises many thorny issues; the right solution is for the recipient to filter their email.)

The effect of this abuse is to place tremendous political and legal pressure on the remailer operator [18]. Of course, remailer operators receive no benefit themselves from providing anonymity services to the world, which makes it all the harder to justify spending much time, money, or effort to defend one's remailer. Each incident of abuse generates a number of complaints to the remailer operator, his ISP, and others who might be in a position to pressure them. This situation has become so acute that one of the greatest difficulties in setting up a new remailer is finding a host who will not give in to the political pressure.

Undoubtedly the magnitude and severity of abuse will increase when more infrastructure (such as pipenet) becomes available, and we will need to know how to deal with this problem. For instance, pipenet potentially allows malicious hackers to break into a remote site untraceably. We can borrow some techniques from today's remailers. For instance, intrusion detection software at the last hop in a pipenet chain may detect some attacks, but it also has some serious limitations; we can also use source blocking to shut out known trouble-makers. New techniques will probably be needed too. For example, some have suggested that requiring a small payment for the anonymity services would reduce spam, harassment, and denial of service attacks by making it too expensive to send large volumes of data; also, the resulting revenue might make it easier and more economical for providers of anonymity services to handle abuse and stand up to political pressure. In any case, abuse management and prevention is likely to remain a central challenge for future anonymity technology.

Others have proposed some special-purpose applications for Internet privacy, though implementation experience is somewhat lacking. The Eternity Service [1] is designed to provide long-term distribution of controversial anonymous documents, even when the threat model includes governments and other powerful parties, but the design has not been implemented and deployed yet. Many cryptographers have studied the problem of electronic voting, and cryptographic protocols abound [25]—but more practical experience with building and deploying large voting systems is needed. The need for more application-specific privacy-respecting systems will no doubt arise as the Internet continues to grow.

Perhaps the most important challenge facing Internet privacy advocates is to ensure that it sees widespread deployment. The issues include educating users about the need for special privacy protection to restore the privacy lost in an online world, building privacy software that is integrated with popular applications, winning over those who fear anonymity, and building systems that meet the needs of real users. It is important that this technology reaches the users who most need it.

## 7. Conclusion

We have surveyed a number of privacy technologies currently available to the Internet user. We have also listed a number of challenges and directions for future research.

We wish to see a variety of means by which users can protect their privacy, preferably by putting privacy-enhancing technology directly into their own hands. Where the cooperation of others is necessary to ensure personal privacy, the system should not be easily subverted by the mere collusion or compromise of a few participants.

We conclude with an important piece of wisdom from the cypherpunks [19, 20]. The cypherpunks credo can be roughly paraphrased as "privacy through technology, not through legislation." If we can guarantee privacy protection through the laws of mathematics rather than the laws of men and whims of bureaucrats, then we will have made an important contribution to society. It is this vision which guides and motivates our approach to Internet privacy.

## References


[1] Ross Anderson, "The Eternity Service," *PRAGOCRYPT 96*. `ftp://ftp.cl.cam.ac.uk/users/rja14/eternity.ps.Z`

[2] Gary Anthes, "IRS uncovers bogus access to tax records (Internal Revenue Service's Atlanta office investigation)," *Computerworld*, vol. 27 no. 32, 9 Aug 1993, p. 15.

[3] Associated Press, 19 Sept 1996.

[4] Andre Bacard, "Anonymous Remailer FAQ," 1996. `http://www.well.com/user/abacard/remail.html`

[5] James Bamford, *The Puzzle Palace*, Penguin Books, New York, 1983.

[6] Douglas Barnes, "The Coming Jurisdictional Swamp of Global Internetworking (Or, How I Learned to Stop Worrying and Love Anonymity)," unpublished manuscript, 16 Nov 1994.
http://www.communities.com/paper/swamp.html

[7] David Chaum, "Untraceable Electronic Mail, Return addresses, and Digital Pseudonyms," *Communications of the ACM*, February 1981, vol. 24 no. 2.
http://www.eskimo.com/˜weidai/mix-net.txt

[8] David Chaum, "Blind Signatures for Untraceable Payments," *CRYPTO 82*, Plenum, pp. 199-203.

[9] Community ConneXion, "Anonymous Surfing," 1996.
http://www.anonymizer.com/

[10] Elizabeth Corcoran, "Hackers Strike at NY Internet Access Company," *The Washington Post*, 12 Sept 1996, p. D09.

[11] Lance Cotrell, "Mixmaster & Remailer Attacks," 1995.
http://www.obscura.com/˜loki/remailer/remailer-essay.html

[12] Ray Cromwell, "Welcome to the Decense Project," 1996.
http://www.clark.net/pub/rjc/decense.html

[13] Wei Dai, "PipeNet," Feb 1995, post to the `cypherpunks` mailing list.

[14] Wei Dai, personal communication.

[15] Laurent Demailly, "Announce: Anonymous Http Proxy (preliminary release)," Usenet post.
http://www.lyot.obspm.fr/˜dl/anonproxy.txt

[16] Arnoud Engelfriet, "Anonymity and Privacy on the Internet," 19 Dec 1996.
http://www.stack.nl/˜galactus/remailers/index.html

[17] C. Gulcu and G. Tsudik, "Mixing E-mail with Babel," *Proc. Symp. Network and Distributed System Security*, 1996, pp. 2–16.

[18] Johan Helsingius, press release, 30 August 1996.
http://www.cyberpass.net/security/penet.press-release.html

[19] Eric Hughes, "A Cypherpunk's Manifesto," 9 March 1993.
ftp://ftp.csua.berkeley.edu/pub/cypherpunks/rants/.manifesto.html

[20] Steven Levy, "Crypto Rebels," *Wired*, May/June 1993, vol. 1 no. 2, pp. 54–61.
http://www.hotwired.com/wired/1.2/features/crypto.rebels.html

[21] Raph Levien, "premail".
http://www.c2.net/˜raph/premail.html

[22] The Nando Times, 20 Nov 1996, New York, staff and wire reports.

[23] The New Yorker, 5 July 1993, p. 61.

[24] Andreas Pfitzmann and Michael Waidner, "Networks without user observability—design options," *EUROCRYPT 85*, LNCS 219, Springer-Verlag, pp. 245–253.
http://www.informatik.uni-hildesheim.de/˜sirene/publ/PfWa86anonyNetze.html

[25] Bruce Schneier, *Applied Cryptography*, second edition, John Wiley & Sons, 1996.

[26] Paul Syverson, David Goldschlag, Michael Reed, "Anonymous Connections and Onion Routing," draft manuscript.
http://www.itd.nrl.navy.mil/ITD/5540/projects/onion-routing/overview.html