

Using caching for browsing anonymity

ANNA M. SHUBINA

and

SEAN W. SMITH

Dartmouth College

Privacy-providing tools, including tools that provide anonymity, are gaining popularity in the modern world. Among the goals of their users is avoiding tracking and profiling. While some businesses are unhappy with the growth of privacy-enhancing technologies, others can use lack of information about their users to avoid unnecessary liability and even possible harassment by parties with contrary business interests, and to gain a competitive market edge.

Currently, users interested in anonymous browsing have the choice only between single-hop proxies and the few more complex systems that are available. These still leave the user vulnerable to long-term intersection attacks.

In this paper, we propose a caching proxy system for allowing users to retrieve data from the World-Wide Web in a way that would provide recipient unobservability by a third party and sender unobservability by the recipient and thus dispose with intersection attacks, and report on the prototype we built using Google.

Categories and Subject Descriptors: K.4.1 [**Computers and Society**]: Public Policy Issues—*Privacy*

General Terms: Human factors, Security, Legal aspects

Additional Key Words and Phrases: Anonymity, privacy, caching proxies

1. INTRODUCTION.

Following the most recent political and economic developments, from increased consumer profiling for many ends and purposes (see, for example, [Electronic Privacy Information Center 2003]) to plans of the US Government for Terrorism Information Awareness (formerly Total Information Awareness, see [DARPA 2003]), a number of previously unconcerned people started to contemplate being more careful with exposing or giving out any information about themselves. A number of businesses that have been collecting data about their customers have also taken damage from adversary parties. It is no longer a purely hypothetical possibility that a vendor of electronic equipment will find his web logs subpoenaed on murky legal grounds. An example is the recent DirecTV case, where customer lists of a number of different companies that sold smart cards were seized, and people on these lists accused of

This research has been supported in part by the Mellon Foundation, NSF (CCR-0209144), AT&T and the Office for Domestic Preparedness, Department of Homeland Security (2000-DT-CX-K001). This paper does not necessarily reflect the views of the sponsors.

Permission to make digital/hard copy of all or part of this material without fee for personal or classroom use provided that the copies are not made or distributed for profit or commercial advantage, the ACM copyright/server notice, the title of the publication, and its date appear, and notice is given that copying is by permission of the ACM, Inc. To copy otherwise, to republish, to post on servers, or to redistribute to lists requires prior specific permission and/or a fee.

© 2003 ACM 0000-0000/2003/0000-0011 \$5.00

piracy [Poulsen 2003].

The lack of privacy of transactions on the web being well-known (as documented, for example, at [Electronic Privacy Information Center 1997]), more people are now starting to try services that claim to provide anonymous browsing on the web, such as `anonymizer.com` or `the-cloak.com`.

There currently exist a number of very different services - free or paid, with more or fewer features and paranoia, with better or worse performance, with bigger or smaller quotas - which claim that they would anonymize one's web browsing. Looking closer, however, one cannot help noticing that most of them are based on the same underlying assumptions, the same attacker model and the same design.

The few more complicated systems that consider a more powerful attacker have yet to go into wide use and require extra software download and setup from the user (and some of them may also come with extra obligations, such as automated request-forwarding). Of these, we are aware of only two that are currently publicly available: Crowds [Reiter and Rubin 1998] and JAP [Berthold et al. 2001]. The first commercial mix-net, Freedom Network [Goldberg and Shostack 2001] is now offline permanently [Shostack and Goldberg 2000]; Onion Routing [Syverson et al. 1997] is offline temporarily (we hope), waiting for the test of its second generation system.


This does not leave many options to a user concerned about his privacy, raising the question: what are the simplest anonymizing systems really giving a user in search of protection, and is there a way to get better protection than that?

2. ANONYMITY AND UNOBSERVABILITY.

In the terms of [Pfitzmann and Köhntopp 2000], *anonymity is the state of not being identifiable within a set of subjects*. *Sender anonymity* means that the sender of the message cannot be identified, while *recipient anonymity* means that the recipient of the message cannot be identified. *Unlinkability of sender and recipient* means that it cannot be identified that a given sender and recipient are communicating with each other. Even stronger than anonymity, *unobservability* means that the existence of the message itself cannot be detected. *Sender (or recipient) unobservability* means that it is impossible to detect whether any sender (recipient) from the set of subjects is sending (receiving), while *relationship unobservability* means that it is impossible to detect the existence of a message sent from a possible sender to a possible recipient.

For this definition to work, one needs to consider also two concepts: that of *the attacker* against whom anonymity is achieved and of *the degree of anonymity* (as described in [Reiter and Rubin 1998]), the certainty with which the attacker can pinpoint the sender, the recipient, or link them to each other. The attackers can be *passive* - merely watching the traffic, or *active* - capable of inserting their own traffic. They can be *local*, watching only one node or wire, or *global*, watching multiple nodes or wires. They can watch traffic over short or long periods of time. Theoretical treatments also need to consider the computational power of the attacker.

Fig. 1. Single-hop proxy



3. ANONYMIZING SERVICES ON THE WORLD-WIDE WEB.

Single-hop proxy functionality.

Most anonymizing services currently available to the general public at their base are merely single-hop proxies that make the http request for the user. Figure 1 illustrates the operation of such a proxy. While using a single-hop proxy, the user submits the destination URL to the service and the service immediately issues an http request to this URL. The http request appears to have originated not at the user's computer but at the proxy. As the target computer replies, sending back an html document, the proxy sends back to the user this document with all links rewritten so that they point back to the proxy, not to the sites they originally pointed to. This means that when using, for example, the free part of `anonymizer.com`, a link to `http://slashdot.org` gets rewritten, becoming `http://anon.free.anonymizer.com/http://slashdot.org`.

The connection of the user to the proxy may or may not be encrypted. Some services (see, for example, `the-cloak.com`) provide both; some (see, for example, `anonymizer.com`) provide only unencrypted surfing for free and charge for encrypted surfing.

Besides concealing the user's IP and possibly encryption, extra functionality of anonymizing single-hop proxies may include: filtering out or specially handling cookies; filtering out or rewriting JavaScript, Java, or other active content; filtering out advertisements and banners; proxying or blocking https; faking the `http_user_agent` field in the http header (that is, not revealing information about the user's OS and browser); faking the `http_referer` field (that is, not disclosing the previously visited site).

Single-hop proxy examples. Table I lists a few available services. Note that most such proxies aren't run by businesses that have obligation to keep them up, and therefore may appear and disappear unpredictably. Some of these services (such as, for example, `anonymizer.com`) have been subject to public review for a long time and have a reputation for not violating their users' privacy. For others, caveat emptor may apply.

Single-hop proxy attacker model. A single-hop proxy attempts to protect the identity of the sender of a request from the attacker who can monitor the traffic of the destination site.

What this really means is, that leaving aside bugs and faults in design or imple-

Table I. Some available single-hop proxies as of June 2003

Name	URL	Encryption
Anonymizer	http://www.anonymizer.com	paid version only
the-Cloak	http://www.the-cloak.com	yes
ProxyWeb.net	http://www.proxyweb.net	paid version only
SnoopBlocker.com	http://www.snoopblocker.com	paid version only
Proxify.com	http://proxify.com	no
Anonymouse	http://anonymouse.ws	no
Web Warper	http://webwarper.net	no
Anonymization	http://www.anonymization.net	no
PurePrivacy	http://www.pureprivacy.com	no

mentation that may make such services able to disclose the information they claim to conceal (an example is in [Martin and Schulman 2002]) and leaving aside their additional content filtering capabilities, the main function of a single-hop proxy is the concealment of the user’s IP from the site he or she accesses.

The destination site can do traffic analysis (see, for example, [Song and Korba 2002] or [Raymond 2001]) to learn about and from the user’s browsing patterns. The site can see what paths the user takes inside it (*communication pattern attack*), the intervals of time between requests (*timing attack*), the amount of transmitted data (*packet volume and counting attack*), as discussed in [Song and Korba 2002]. The site may also be able to correlate different accesses by the same user that occur at different times (*intersection attack*).

Single-hop proxies also do not address the model of the attacker who may be watching user’s traffic as it goes from the user to the proxy and from the proxy to the user. Sometimes they do not encrypt such traffic, allowing any observer on its path the full view of what is going on.

Even if the traffic from the user to the proxy and back is encrypted, the global attacker that is able to watch both the user-proxy and the proxy-destination sides of the proxy might be able to see what data the user requests and gets by doing traffic analysis on the ingoing and outgoing traffic.

Going beyond a single-hop proxy. In 1981, Chaum introduced the concept of “Mix-nets” (see [Chaum 1981]). “Mix-nets” are groups of servers that provide anonymity by passing user’s traffic through nodes called *Mixes* which may delay, reorder, reencrypt, pad and forward traffic passing through them. In addition to providing sender anonymity, Mix-nets attempt also to provide sender and recipient unlinkability against a global attacker. Among the examples of Mix-nets are Onion Routing [Syverson et al. 1997], Zero Knowledge Systems’ Freedom Network [Goldberg and Shostack 2001], Web MIXes [Berthold et al. 2001], Tarzan [Freedman and Morris 2002].

Among the other approaches to the problem of anonymous web transactions is the Crowds system [Reiter and Rubin 1998], which allows participating nodes to forward requests within their crowd with a certain probability. Crowds attempt to provide sender and recipient anonymity, but unlike the Mix-nets, they do not attempt to provide sender and recipient unlinkability against a global attacker.

A problem: long-term intersection attacks. A subset of intersection attacks and a major unsolved problem in anonymity systems are so-called “long-term intersection attacks” (see [Berthold and Langos 2002]), where the attacker is able to watch all or many nodes, including a number of entrance and exit points of the users traffic. Users tend to have a certain behaviour while on-line. They tend to send messages to same or similar destinations. If user *A* in a mix daily reads a

Table II. Anonymizing systems compared

	Sender anonymity from		Recipient anonymity from		Sender, recipient unlinkability	Single point of failure
	recipient	3rd party	sender	3rd party		
Single-hop proxies	yes	no	no	no	no	yes
Mix-nets	yes	yes	no	yes	short-term	no
Crowds	yes	yes	no	yes	no	no
Unencrypted caching ¹	yes	no	no	no	no	yes
Encrypted caching ¹	yes	yes	no	yes	yes	yes

certain site, after very many observations it may be possible to match A with the outgoing request to that site, using the knowledge of A 's presence or A 's absence together with the information whether or not the site was visited during that time.

One of the directions for dealing with this is dummy traffic (as discussed in [Berthold and Langos 2002], [Jiang et al. 2001]). Dummy traffic consists of messages sent even when the user has nothing to send. A totally different approach is proposed in PipeNet [Dai 1998], where the whole mix-net is supposed to shut down when one node stops to communicate with the network.

4. USING CACHING FOR ANONYMITY.

We propose an anonymizing web-browsing system that would deal with the global attacker who is able to do long-term observation. The proposal is to make user requests to proxy independent in time from the proxy's requests to target sites. The proxy will request and cache information on its own and send it to the users when requested. To do this, the proxy server would collect and store documents which its users are likely to request; provide an encrypted channel to access its contents; and rewrite links in the documents provided to the user to point back at the proxy, similarly to the workings of an ordinary single-hop proxy. This would mean that a global observer looking at the traffic to and from the server will see only regularly scheduled cache updates and encrypted user traffic, thus removing correlation between the users requests and the server's accesses to remote sites.

This caching proxy will act like a single-hop proxy, in that the destination server does not find out who communicated with it except for the proxy. But it does more than the single-hop proxy in many other respects: the destination server does not find out that someone except the caching proxy communicated with it at all, and neither does a third-party observer; the global attacker cannot correlate requests from the sender and to the recipient; and long-term intersection attacks are meaningless since there is no connection between the proxy's fetches and the user's requests.

In short, the attacker model includes not only the attacker at the destination site (as that of existing single-hop proxies) but also the global attacker who is able to do not only short-term, but also long-term attacks.

Similarly to the existing single-hop proxies, the user will not be protected from the attacker at the proxy itself.

Such a proxy would compare with the existing anonymity services as described in Table II.

¹Proposed for providing anonymity in this paper.

5. CANDIDATES FOR A PROTOTYPE.

Caching proxies that periodically update their content. The best known example of a caching proxy that collects data on its own for future use is Google's cache [Google]. Every few weeks it crawls all the web, collecting, storing and indexing text information from every webpage [Brin and Page 1998]. Attempting to scale to the size of the whole World Wide Web, it has been optimized for memory usage and speed.

Google's primary objective is indexing text information. It does make the attempt to index images but only caches thumbnail images, relying on the original image being still present on the server. Such usage may be due also to data size considerations, but there are also legal reasons for it. In a recent case, the 9th U.S. Circuit Court of Appeals ruled that it is legal for search engines to cache thumbnail images [Olsen 2003a], but it is yet to be determined whether it is legal to cache a full-size image copy. According to the Court, the use of thumbnail images is fair use, because they 'do not supplant the need for originals' and 'benefit the public by enhancing information-gathering techniques on the Internet'.

Other search engines that allow access to their caches are Comet Web Search [Comet]² and Gigablast [Gigablast]. An interesting web cache is Wayback Machine [WayBack]. This is an internet archive that contains multiple copies of webpages from 1997 to 2002. Similarly to the caches of search engines, it cannot be used for anonymous browsing. Using JavaScript, Wayback Machine handles links to point back to the cache where archived versions are available, but if that fails, it sends the request to the original site.

“Using Google cache as anonymous surf”. The idea of using Google cache for anonymity has occurred to Google's users before, but Google's design is not targetted towards this goal. Written apparently as a result of such attempts, an FAQ entitled 'Don't Use the Google Cache as Anonymous Surf', [The Virtual Chase 2003] explains that clicking on a link in a cached page would lead you to a page outside Google cache, an embedded image would attempt to load directly from the original site, and redirection might make the whole page load from the original site.

6. PROTOTYPE.

We prototyped³ an anonymizing caching proxy that outsources all caching to Google by forwarding all requests to Google cache and rewriting links in the resulting documents. As shown in Figure 2, when a URL is submitted to the script, the script requests a cached copy from the Google cache. Any link in the document gets rewritten to submit the link to this script, instead of going directly to the requested page. Thus, if we deal with plain text data, every request would be redirected to the Google cache.

The implementation has to detect all outward going links, detect their data type and to rewrite them according to whether their content may be available from

²Comet Web Search stopped allowing access to their cache sometime after the submission of this article. The reasons for this are unknown to the authors.

³The prototype is at <http://www.cs.dartmouth.edu/~ashubina/google.html>. It is a CGI script that can be located at the user's machine or elsewhere.

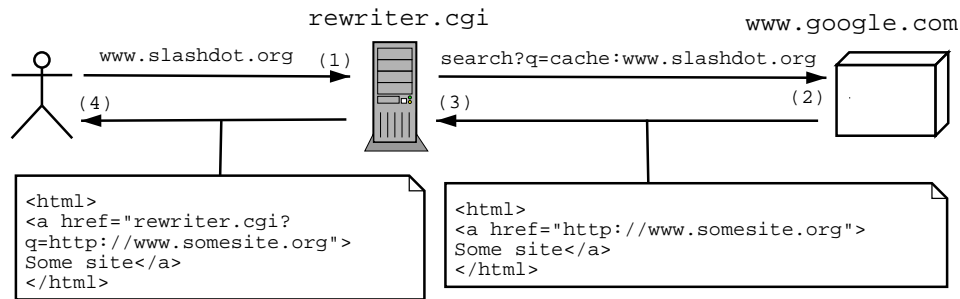


Fig. 2. Anonymizing caching proxy

Google cache.

Since Google cache does not store images, images are not provided unless the user explicitly specifies that he or she wants the script to act like a single-hop proxy that fetches them for him from the destination server. Such a transaction would disclose the user's browsing to the global observer. If the script is located at the user's machine, it would also disclose his IP to the destination site, while if the script is not located at the user's machine, this behaviour would be no different from that of a single-hop proxy.

7. THE CACHING PROXY CONTROVERSY.

It appears to be an unfortunate law of the modern civilization that whenever anyone wants to give out something for free, someone else is going to have a problem with it. Caching proxies are no exception to this rule.

Most Web sites strive for high search engine rankings (on how to play this game, see, for example, [Calishain and Dornfest 2003]). However, as described in a recent *news.com* article [Olsen 2003b], many of them are unhappy with caching of their contents.

[Olsen 2003b] gives the following list of complaints about search engine caching:

- **Caching of removed data.** Caching lets users access pages that are temporarily unavailable. Thus it sometimes provides also access to content that has been deliberately taken off-line.
- **Access to registration-only sites.** Due to faults in the design of a number of sites that require registration, caching may end up indexing and caching their content.
- **Detouring traffic from the original sites.** The very feature that we propose to use for anonymity is a subject of such complaints. The claim is that by detouring traffic from the sites where the original information is stored, caching may make Web publishers lose their income.

Interestingly, for the current implementations of all the web caches we reviewed ([Google], [WayBack], [Comet], [Gigablast]), the last claim is not quite valid. Among the many ways in which the user's anonymity from the original site provided by the use of such caches is violated are images, which usually include also ads. Since ads are not cached, a cached copy of a webpage with an ad usually displays

the ad downloaded from its original provider. Such ads may, however, not be able to track to which site they really owe this request very well, and thus may mistake the web cache for the referrer. (However, the proxy we propose has a different behaviour.)

Google's response to the last complaint is that web sites can easily prevent Google from caching their pages without preventing it from indexing them. All it takes is adding to your page either `<META NAME="ROBOTS" CONTENT="NOARCHIVE">` to exclude all robots, or `<META NAME="GOOGLEBOT" CONTENT="NOARCHIVE">` to exclude just Google. However, some sites are reluctant to do this for fear that this may affect their search rankings.

The Digital Millennium Copyright Act has a narrow exception for Web caching, allowing internet service providers to keep local copies of Web pages. It is not clear, however, whether this would protect also search engine caches or archives similar to [WayBack] if tested.

The legality of a caching proxy such as we propose under the modern copyright laws is yet to be determined, as is the legality of search engine caches. Due to the very features that make our scheme provide anonymity that are not there for other web caches, our scheme may not be ruled legal even if other web caches are.

8. EXTENSIONS.

Obviously, a proxy such as we propose would not allow the users to use any dynamic content. A user could not provide any information to be passed to the destination site without informing the destination site and the global observer that someone is using the proxy to look at the destination site. It might be possible to design a proxy that submits such a request for the user together with other people's requests or a number of randomly generated requests, but such an extension is bound to be subject to attacks.

However, a natural extension would be allowing to search Google cache, using Google itself and rewriting the results to point back into the cache. It might also be possible to design a caching proxy that provides thumbnail images, as does Google.

Google (as well as the other caching proxies described above) does not provide encryption, hence the information going to and from Google is not encrypted. An encrypted proxy would allow protection from an attacker that can watch traffic going to it. Against such an attacker, it would also help to bundle packets together or break them up, normalizing their size. Such a proxy could also attempt to conceal its operations from an attacker at the very proxy by means of private information retrieval.

9. CONCLUSION.

In this paper we presented an approach that would allow retrieving text information from the web with better anonymity than a single-hop proxy allows to achieve and would also eliminate long-term correlation attacks by the global attacker. The approach uses a caching proxy that contains its own copy of the World Wide Web (or of its large subset). One example of such a proxy is Google cache, which we used for a prototype.

Our prototype is already useful both for achieving limited anonymity, and for

accessing cached copies of sites that for some reason are unaccessible at the moment.

As privacy of web transactions becomes more and more important and harder and harder to achieve, it may be worthwhile to use caching proxies to achieve anonymous access to what may be the most important achievement of the human civilization - written text information.

REFERENCES

- BERTHOLD, O., FEDERRATH, H., AND KÖPSELL, S. 2001. Web MIXes: A system for anonymous and unobservable Internet access. *Lecture Notes in Computer Science 2009*, 115–??
- BERTHOLD, O. AND LANGOS, H. 2002. Dummy traffic against long term intersection attacks. In *Privacy Enhancing Technologies 2002*. Springer-Verlag.
- BRIN, S. AND PAGE, L. 1998. The anatomy of a large-scale hypertextual search engine. In *Seventh International World Wide Web Conference*.
- CALISHAIN, T. AND DORNFEST, R. 2003. *Google Hacks*. O'Reilly.
- CHAUM, D. 1981. Untraceable electronic mail, return addresses and digital pseudonyms. *Communications of the ACM* 24, 2 (February), 84–88.
- COMET. Comet web search. <http://search.cometsystems.com>.
- DAI, W. 1998. PipeNet 1.1. <http://www.eskimo.com/~weidai/pipenet.txt>.
- DARPA. 2003. Report to Congress regarding the Terrorism Information Awareness program. http://www.darpa.mil/body/tia/tia_report_page.htm.
- ELECTRONIC PRIVACY INFORMATION CENTER. 1997. Surfer beware: Personal privacy and the internet. <http://www.epic.org/reports/surfer-beware.html>.
- ELECTRONIC PRIVACY INFORMATION CENTER. 2003. Privacy and consumer profiling. <http://www.epic.org/privacy/profiling>.
- FREEDMAN, M. J. AND MORRIS, R. 2002. Tarzan: A peer-to-peer anonymizing network layer. In *Proceedings of the 9th ACM Conference on Computer and Communications Security (CCS 2002)*. Washington, D.C.
- GIGABLAST. Gigablast search. <http://gigablast.com>.
- GOLDBERG, I. AND SHOSTACK, A. 2001. Freedom network 1.0 architecture and protocols.
- GOOGLE. Google. <http://www.google.com>.
- JIANG, S., VAIDYA, N. H., AND ZHAO, W. 2001. Power-aware traffic cover mode to prevent traffic analysis in wireless ad hoc networks. In *IEEE Infocom*.
- MARTIN, D. AND SCHULMAN, A. 2002. Deanonymizing users of the SafeWeb anonymizing service. Tech. Rep. 2002-003. 11.
- OLSEN, S. 2003a. Court backs thumbnail image linking. http://news.com.com/2100-1025_3-1023629.html?tag=bplst.
- OLSEN, S. 2003b. Google cache raises copyright concerns. http://news.com.com/2100-1032_3-1024234.html?tag=fd_lede2_hed.
- PFITZMANN, A. AND KÖHNTOPP, M. 2000. Anonymity, unobservability, and pseudonymity — a proposal for terminology. In *Designing Privacy Enhancing Technologies: Proceedings of the International Workshop on the Design Issues in Anonymity and Observability*. Vol. 2009. Springer-Verlag, 1–9.
- POULSEN, K. 2003. DirecTV dragnet snares innocent techies. <http://www.securityfocus.net/news/6402>.
- RAYMOND, J.-F. 2001. Traffic analysis: Protocols, attacks, design issues and open problems. In *Designing Privacy Enhancing Technologies: Proceedings of International Workshop on Design Issues in Anonymity and Unobservability*, H. Federrath, Ed. LNCS, vol. 2009. Springer-Verlag, 10–29.
- REITER, M. K. AND RUBIN, A. D. 1998. Crowds: anonymity for Web transactions. *ACM Transactions on Information and System Security* 1, 1, 66–92.
- SHOSTACK, A. AND GOLDBERG, I. 2000. How not to design a privacy system: reflections on the process behind the Freedom product. In *Proceedings of the tenth conference on computers*, ACM SIGecom Exchanges, Vol. 4, No. 2, September 2003.

freedom and privacy: challenging the assumptions. ACM Press, Toronto, Ontario, Canada, 85–87.

SONG, R. AND KORBA, L. 2002. Review of network-based approaches for privacy. In *Proceedings of the 14th Annual Canadian Information Technology Security Symposium*. Ottawa, Ontario.

SYVERSON, P. F., GOLDSCHLAG, D. M., AND REED, M. G. 1997. Anonymous connections and Onion Routing. In *IEEE Symposium on Security and Privacy*. Oakland, California, 44–54.

THE VIRTUAL CHASE. 2003. Don't use the Google cache as anonymous surf. http://www.virtualchase.com/ask_answer/google_cache.html.

WAYBACK. WayBack Machine. <http://www.archive.org>.

Received July 2003; revised September 2003; accepted September 2003