

# Anonymizing Censorship Resistant Systems

Andrei Serjantov  
Andrei.Serjantov@cl.cam.ac.uk

University of Cambridge Computer Laboratory,  
William Gates Building,  
JJ Thomson Avenue,  
Cambridge CB3 0FD, UK

March 1, 2002

## Abstract

In this paper we propose a new Peer-to-Peer architecture for a censorship resistant system with user, server and active-server document anonymity as well as efficient document retrieval. The retrieval service is layered on top of an existing Peer-to-Peer infrastructure, which should facilitate its implementation. The key idea is to separate the role of document storers from the machines visible by the users, which makes each individual part of the system less prone to censorship.

## 1 Introduction

Many censorship resistant systems have been proposed recently, yet most still lack one crucial feature – protection of the servers hosting the content.

In the past this was not considered an issue. For instance, in Anderson’s eternity service [And96], it was deemed sufficient to guarantee that a document was always available through the system. However, many examples indicate that servers hosting content are vulnerable to censorship, due to “Rubber Hose Cryptanalysis” – various kinds of pressure applied by attackers to shut down servers or remove files. Examples of documents subjected to censorship include DeCSS [DeC], the paper detailing the attack on SDMI [CMW<sup>+</sup>01], and documents (or quotes

from documents) which the Church of Scientology described as their secrets. In cases like this, the server administrators receive “cease and desist” letters when the censor finds the offending document on their server.

Most modern censorship resistant systems, for example Publius [WRC00] and Freenet [CSWH01], have not addressed this problem in a satisfactory way. It is possible for a malicious reader to find out which servers content is stored on, and subsequently try to pressure the server administrators to remove it. With Publius in particular, the situation is slightly more complicated as each document is encrypted and the key is split into  $n$  shares, any  $k$  of which are re-combinable to form the document back. In this case, the attacker needs to remove content from  $n - k + 1$  servers, all of which he can easily locate.<sup>1</sup> With the number of servers reasonably small and static, the job of censoring documents becomes easier than one might expect.

An alternative approach to dealing with censorship resistance is taken by systems like Dagster [SW01] and Tangler [WM01] which prevent removal of any single document from the system by entangling documents together. However, in our view, these are not effective enough at dealing with the problem either: if the offending document was entangled with

---

<sup>1</sup>Indeed, if one server has been pressured into removal, the other server administrators may simply follow the precedent and remove the offending content themselves.

the Declaration of Independence, Das Capital and the Little Red Book, censoring it (thereby removing all of the above from the system) would not be a major problem as all the other documents are readily available from other sources. Given the other documents are *not* available, the censor is unlikely to be discerning enough to want to keep some of them. Furthermore, if the system is run on a single server (eg. Dagster), then the censor may simply try to shut down the entire server.

In our system, we consider the storers of the files valuable entities, and protect them against “Rubber Hose Cryptanalysis”. Furthermore, we protect the documents they are storing by providing *active-server document anonymity* (as first introduced in [DFM01]). This is a property which states that the storer should not be able to determine (parts of) which document it is storing, not even by retrieving the document from the system. We now proceed with a description of the system and then give an analysis and critique of it.

## 2 System Description

Our system consists of many identical peers, each of which can fulfil four different roles:

- Publisher  $P$ . The node which has a document and wishes to make it available and censorship resistant.
- Forwarder  $a$ . A node which has an anonymous pointer to a node storing part of a document.
- Storer  $s$ . A node which stores part of a document.
- Client  $c$ . A node which retrieves a document.

The system is built on top of an existing Peer-to-Peer file sharing service like PAST [DR01]. PAST can be viewed as a network of machines (peers), each with a unique identifier. Neighbouring machines (machines within a certain distance of each other within a logical name space) share state. The only thing required to send a message to a machine is its *id*, furthermore, PAST guarantees that the message takes no more than  $\log N$  hops, where  $N$  is the number of nodes in the system. We also assume a




Figure 1: Publishing

public key infrastructure, so any peer is able to learn any other peer’s public key. This is further discussed in the next section. Using an existing Peer-to-Peer architecture allows us to abstract from routing, clients leaving and joining the network, and other low level issues. An architecture like PAST also provides robustness, which is further discussed in the next section. We also make use of an anonymous connection system such as Onion Routing [GRS99] which is capable of handling reply blocks.

As usual in censorship resistant systems, the operations available to a node are publishing and retrieval. There is no search facility, therefore we rely on a broadcast mechanism like an anonymous newsgroup to transmit retrieval information to potential readers. We do not support content deletion or modification.

The overall publishing process is illustrated in Figure 1. To publish a document (see Figure 2), the publisher  $P$  splits it into  $n + 1$  shares  $h_i$ , any  $k + 1$  of which can be combined to form the whole document again. This can be done using one of the standard algorithms like Shamir’s secret sharing [Sha79]. He then generates  $n + 1$  keys  $k_i$  and encrypts each share with the corresponding key. He now picks  $n + 1$  peers  $a_0 \dots a_n$  at random to act as forwarders and constructs onions<sup>2</sup> to send (via the anonymous connections layer) each of them the encrypted share  $\{h\}_{k_i}$ ,

<sup>2</sup>This is a technique first described in [Cha81]. A (standard) onion for destination  $d$  with message  $M$  and peer sequence  $a_0 \dots a_n$  is  $\{d, \{M\}_{k_d}\}_{k_{a_n}} \dots \}_{k_{a_0}}$  which is sent to  $a_0$ .  $a_i$  is the address of the server and  $k_{a_i}$  is its public key.




Figure 2: Inserting share  $a_0$ . All communication is done via the anonymous connection system using randomly constructed onions. If the message is sent using a return address, it is displayed at the base of the arrow. Anonymous return addresses are denoted by  $r$ , eg.  $r_{a_0}$

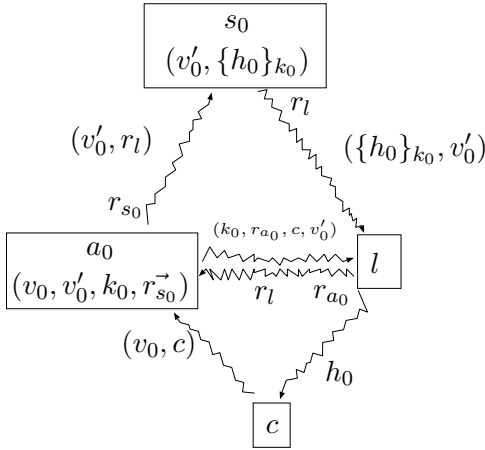


Figure 3: Retrieval

the corresponding key  $k_i$  and a random integer  $v_i$  together with a return address (reply onion)<sup>3</sup> The publisher can now wait for a confirmation to come back from each of the  $a_i$ s (via the reply onion) saying whether the publishing has been successful or not. If the operation failed, the publisher should try different  $a_i$ s.

Each of the forwarders (take  $a_0$  as an example) receives the message, finding an encrypted share  $\{h_0\}_{k_0}$ , a key  $k_0$  and a random number  $v_0$ . He then picks a storer  $s_0$  to store the share and a small number  $v'_0$  which the storer would associate the share with. He constructs an onion for delivering these to the storer. Thus, he puts the encrypted share,  $v'_0$ , as well as its own anonymous return address  $r_{a_0}$  into the onion as the message and sends it off. If the onion is received by  $s_0$ , it stores the share and issues a number of different return addresses  $r_{s_0}^{\rightarrow}$  (to be used for retrieval), sending them back to  $a_0$  via the return address  $r_{a_0}$ . Now  $a_0$  associates  $v_0$ ,  $v'_0$  and  $k_0$  with the return addresses  $r_{s_0}^{\rightarrow}$ , forgets  $s_0$ , and replies “ok” to the publisher. Once all the shares have been stored, the publisher destroys them and announces the name of the file, together with the  $n + 1$  pairs  $(a_i, v_i)$  to potential users.

To retrieve a document (see Figure 3b), the client  $c$  asks the forwarder  $a_0$  (and each  $a_i$  in the same way) to retrieve the share  $h_0$  by sending them an anonymous message with  $v_0$  and their address  $c$ . The forwarder  $a_0$  then picks a random server  $l$  to act as a decrypter and sends it  $k_0$ , the key it is storing which decrypts the stored share,  $v'_0$ ,  $c$ , and a return address  $r_{a_0}$ , getting back a return address for  $l$ . Now  $a_0$  forwards  $r_l$  and  $v'_0$ , which identifies the share, to  $s_0$  via one of the  $r_{s_0}^{\rightarrow}$  ( $r_{s_0}$ ). Now  $s_0$  looks up the encrypted share corresponding to  $v'_0$  and forwards it and  $v'_0$  to  $l$ , which decrypts the share and sends it to the client. The process continues until  $c$  has accumulated enough shares to reconstruct the document.

<sup>3</sup>A return address is a kind of onion which, if included in an anonymous message, can be used to reply to that message without revealing the original sender (see [Cha81] for details).

### 3 Discussion

In this section we discuss the limitations of our system. First, one should question the validity of assuming a public key infrastructure on a P2P network. We need each peer to be able to retrieve the public key of any other peer and verify that the key belongs to that particular peer. The simple solution is to use a global repository. However, such a scheme would limit the scalability of the system. We believe that better solutions exist, and are actively working on this problem.

Secondly, we should consider the forwarders – they are certainly visible to the attacker. However, we argue that they are much less likely to be subjected to “Rubber Hose Cryptanalysis” than, for example, Publius servers for the following reasons:

- They are not storing the offending document, not even in an encrypted form, so their connection with it is somewhat indirect.
- They do not store the identities of the  $s_i$ s, so an attack to try to get it out of them will not succeed.
- The share does not actually go through the peers  $a_i$  after publication is completed.
- The forwarders can deny that the request they received from  $c$  had anything to do with a share arriving at  $c$  some time later.

A slight modification to the protocol (which is beyond the scope of this paper) can be introduced to further reduce the role the  $a_i$  play in the protocol, and therefore reduce the potential for them to be attacked.

Thirdly, we must consider the number of compromised peers it takes to remove a document from the system. A possible attack is as follows: each peer  $a_i$  remembers (rather than forgetting) the corresponding  $s_i$  at the time of publication in the hope of exposing  $s_i$  later, if the content turns out to be offending. Once the document has been successfully published,  $a_i$  notes the correspondence between the random integer ( $v_i$ ) published with the document and one in its lookup table, and works out the fact that  $s_i$  is storing a share of a particular document. It can now pressure  $s_i$  into removing the share. However, the chances of

the  $n - k + 1$  peers picked as  $a_i$  being compromised are small and the peers have to be compromised at the time of publication, otherwise the attack fails.

Furthermore, we need to consider what happens when a particular node goes down. Indeed, this may cause problems as a forwarder failing makes the share behind not accessible. This is handled by the fact that PAST shares state among neighbouring nodes. Therefore, requests will automatically be routed to nearby nodes which store the same information. In particular, the forwarders share the data they have  $(v_0, v'_0, k_0, r_{s_0}^{\rightarrow})$  with neighbouring nodes which can therefore also answer requests. Similarly, the storer shares  $(v'_0, \{h_0\}_{k_0})$ . The decrypter does not need to share anything as he will only get one request to decrypt the share and will then give up this role. It is worth noting that, we are relying on PAST to replicate state on a relatively short time scale to survive denial of service attacks.

Having stated that we provide active server anonymity, we must pay attention to the amount of information the storer can gain by repeatedly requesting the document and noticing the requests coming in for the share it is storing. However, a suitable number of random requests generated by the forwarders should weaken this attack. Again, the precise details are beyond the scope of this paper.

We also note that we are presenting just one part of a design of a system. Many questions are left unanswered and a few attacks are not addressed. For example, we do not deal with accountability in any systematic way. Consider a scenario where the attacker is powerful enough to insert many nodes into the system. Each of these, when asked to act as a forwarder, replies “ok”, but drops the share and fails to answer subsequent requests. In this design, we are relying on the inability of the attacker to insert enough malicious nodes to censor documents in this way.

Another attack is simply trying to flood the censorship resistant system with random data so that “real” documents cannot be inserted. Again, this design does not include protection against this. We felt that although standard methods which are summarised in [DFM00] can be used in this system, they will be inappropriate in this context or will not provide ade-

quate protection. Therefore, we leave this for future work.

## 4 Related Work

A variety of censorship resistant systems have been designed, some of which have also been implemented. We have already discussed Publius, Dagster and Tangler but, perhaps the system closest to ours in terms of the aims it tries to achieve is Free Haven [DFM01].

It is built on top of an underlying network of anonymous remailers and deals with reader anonymity, server anonymity and censorship resistance. However, it uses a Gnutella-like search for retrieval of shares of the document. More specifically, a user request to retrieve a particular document gets broadcast from the user node to all the neighbouring nodes, and so on. When a request arrives at a peer which has a matching share, it gets sent off to the requester via a chain of remailers. This scheme for locating files is rather inefficient, and, as the Gnutella experience has shown, does not scale for large numbers of peers. Furthermore, we note that peers frequently exchange shares with each other. This is costly in terms of network bandwidth and makes it hard to provide guarantees that a document will be located. Our system aims to be more efficient both in terms of bandwidth and share retrieval. Finally, Free Haven has not been implemented, perhaps because it contains complex notions of share trading, reputation, etc. We note, however, that moving the individual shares around in the system is an interesting technique for increasing censorship resistance which may be incorporated into our system. For example, the shares may be moved periodically, and the state in the forwarders updated.

## 5 Conclusion

We have presented a design of a system which deals with censorship resistance and satisfies strong anonymity requirements. Although (like many other similar systems) it has not yet been implemented, we

hope that the fact that it is based on top of an existing infrastructure will make the job easier.

We have argued for building an anonymous censorship resistant system on top of a peer to peer architecture and demonstrated the feasibility of doing so to provide strong anonymity and robustness guarantees.

Having described an anonymity system and claimed that it satisfies some properties, we consider it worthwhile to formalise those and prove them rigorously in the future.

## 6 Acknowledgements

I acknowledge support from EPSRC grant GRN24872 *Wide-area Programming* and EU grant PEPITO. A variety of ideas have resulted from conversations with Richard Clayton, George Danezis, Peter Pietzuch and Peter Sewell and from comments by various members of the Cambridge Security group.

## References

- [And96] R. J. Anderson. The eternity service. In *Pragocrypt*. 1996. <http://www.cl.cam.ac.uk/users/rja14/eternity/eternity.html>.
- [Cha81] D. Chaum. Untraceable electronic mail, return addresses and digital pseudonyms. *Communications of the A.C.M.*, 24(2):84–88, 1981.
- [CMW<sup>+</sup>01] S. A. Craver, J. P. McGregor, M. Wu, B. Liu, A. Stubblefield, B. Swartzlander, D. S. Wallach, D. Dean, , and E. W. Felten. Reading between the lines: lessons from the SDMI challenge. In *Information Hiding Workshop*. 2001.
- [CSWH01] I. Clarke, O. Sandberg, B. Wiley, and T. W. Hong. Freenet: A distributed anonymous information storage and retrieval system. In Federrath [Fed01], pages 46–66. <http://freenet.sourceforge.net>.
- [DeC] Gallery of CSS descramblers. <http://www-2.cs.cmu.edu/~dst/DeCSS/Gallery/>.
- [DFM00] R. Dingledine, M. J. Freedman, and D. Molnar. *Peer-to-Peer: Harnessing the Power*

- of Disruptive Technologies*, chapter 16. O'Reilly, 2000.
- [DFM01] R. Dingledine, M. J. Freedman, and D. Molnar. The Free Haven Project: Distributed anonymous storage service. In Federath [Fed01], pages 67–95. <http://freehaven.net>.
- [DR01] P. Druschel and A. Rowstron. Past: A large-scale, persistent peer-to-peer storage utility. In *The 8th Workshop on Hot Topics in Operating Systems*. 2001.
- [Fed01] H. Federrath, editor. *Designing Privacy Enhancing Technologies: International Workshop on Design Issues in Anonymity and Unobservability*, volume 2009 of *Lecture Notes in Computer Science*. Springer-Verlag, 2001. ISBN 3-540-41724-9.
- [GRS99] D. Goldschlag, M. Reed, and P. Syverson. Onion routing for anonymous and private internet connections. *Communications of the ACM (USA)*, 42(2):39–41, 1999.
- [Sha79] A. Shamir. How to share a secret. *Communications of the ACM*, 22:612–613, 1979.
- [SW01] A. Stubblefield and D. Wallach. Dagster: Censorship-resistant publishing without replication. Technical report, Rice University, 2001.
- [WM01] M. Waldman and D. Mazieres. Tangler: A censorship resistant publishing system based on document entanglements. In *8th ACM Conference on Computer and Communication Security (CCS-8)*. 2001.
- [WRC00] M. Waldman, A. D. Rubin, and L. F. Cranor. Publius: A robust, tamper-evident, censorship-resistant, web publishing system. In *Proc. 9th USENIX Security Symposium*. 2000.