The Economics of Censorship Resistance

George Danezis and Ross Anderson

University of Cambridge, Computer Laboratory,
15 JJ Thomson Avenue, Cambridge CB3 0FD, United Kingdom.
(George.Danezis, Ross.Anderson)@cl.cam.ac.uk

Abstract. We propose the first economic model of censorship resistance. Early peer-to-peer systems, such as the Eternity Service, sought to achieve censorshop resistance by distributing content randomly over the whole Internet. An alternative approach is to encourage nodes to serve resources they are interested in. Both architectures have been implemented but so far there has been no quantitative analysis of the protection they provide. We develop a model inspired by economics and conflict theory to analyse these systems. Under our assumptions, resource distribution according to nodes' individual preferences provides better stability and resistance to censorship. Our results may have wider application too.

1 Introduction

Peer-to-peer designs have evolved in part as a response to technical censorship of early remailer systems such as `penet` [13], and early file distribution systems such as Napster [5]. By distributing functionality across all network nodes, we can avoid an obvious single point of failure where a simple attack could incapacitate the network. Although maintaining the availability of the files and censorship resistance is the *raîson d'étre* of such systems, only heuristic security arguments have been presented so far to assess how well they fulfil their role. Other network design issues have often overshadowed this basic security goal.

Two main paradigms have emerged in the last few years in peer-to-peer systems. The first is to scatter resources randomly across all nodes, hoping that this will increase the opponent's censorship costs (we will call this *the random model*). The theory is that censorship would inconvenience everyone in the network – including nodes that are not interested in the censored material – and more support would be gathered against censorship. Anderson's Eternity Service [2–4] follows this strategy, followed by freenet [6] and Mojo Nation [25]. Structured peer-to-peer systems, including distributed hash table-based systems [21, 18] and others [24], scatter files around in a deterministic way on random nodes, which achieves a similar effect.

The second paradigm allows peer nodes to serve content they have downloaded for their personal use, without burdening them with random files (we will call this *the discretionary model*). The popular Guntella [17] and Kazaa [15] are real-world examples of such systems. Users choose to serve files they are interested in and have downloaded. Other theoretical designs have assessed how such systems scale, by performing distributed information retrieval [7, 22].

The comparison between these two paradigms has so far concentrated on their efficiency in systems and network engineering terms: the cost of search, retrieval, communications and storage. In this paper we are going to compare the two paradigms' ability to resist censorship, as per the original intention of peer-to-peer systems. We present a model that incorporates the heterogeneous interests of the peer nodes, inspired by conflict theory [14] and economic analysis. We will use this to estimate the cost of defending networks against censorship using the two different paradigms.

2 The red-blue utility model

We consider a network of N peer nodes. Each node n_i has a different set of interests from the other nodes: it may prefer news articles to political philosophy essays, or prefer nuclear physics to cryptology. Furthermore each node might even have different political views from other nodes. We model this by considering two types of resources: red and blue[1]. We assign to each node n_i a preference for red $r_i \in [0, 1]$ and a preference for blue $b_i = 1 - r_i$ (note that $r_i + b_i = 1$).

Each node likes having and serving resources, but it prefers to have or serve a balance of resources according to its preference r_i and b_i. For this reason we consider that the utility function of a node holding T resources out of which R are red resources and B are blue resources is (with $T = R + B$):

$$U_i(R, B) = -T(R/T - r_i - 1)(R/T - r_i + 1) \tag{1}$$

This is a quadratic function (as shown in figure 1) with its maximum at $R = r_i T$, scaled by the overall number of resources T that the node n_i holds. This utility function increases as the total number of resources does, but is also maximal when the balance between red and blue resources matches the preferences of the node. Other unimodal functions with their maxima at $r_i T$, such as a normal distribution, give broadly similar results.

This model diverges from the traditional economic analysis of peer-to-peer networks, under which peers have ex-ante no incentives to share [11]. This assumption might be true for copied music, but does not hold for other resources such as news, political opinions, or scientific papers. For example, a node with left-wing views might prefer to read, and redistribute articles, that come 80% from *The Guardian* and 20% from *The Telegraph* ($r_i = 0.8, b_i = 0.2$) while a node in the middle of the political spectrum might prefer to read and redistribute them equally ($r_i = 0.5, b_i = 0.5$) and a node on the right might prefer 80% from *The Telegraph* and 20% from *The Guardian* ($r_i = 0.2, b_i = 0.8$). Furthermore their respective utility increases the more they are able to distribute this material in volumes that are in line with their political preferences.

[1] We follow the economics tradition of only considering two goods. Real life preferences have finer granularity, and our results also apply to n goods.

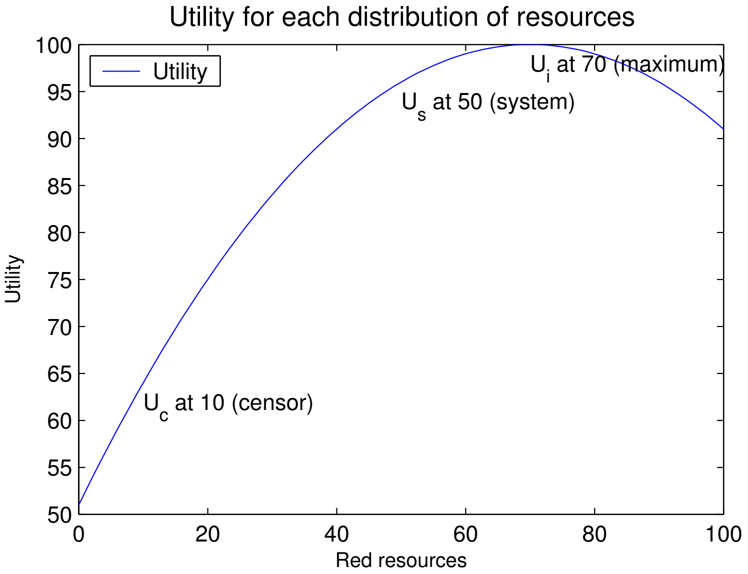

Fig. 1. Example utility model for discretionary, random and censored distribution.

3 The utility of discretionary and random distribution

We will first examine the utility of the network nodes when they can choose which files they store and help to serve (the discretionary model). This corresponds to the Gnutella or Kazaa philosophy. Assuming that a node has the ability to serve T files in total, its utility is maximised for a distribution of red and blue resources that perfectly matches its preferences: $R = r_i T$ and $B = b_i T$. Our proposed utility function U_i is indeed maximal for $U_i(r_i T, b_i T)$ at each node n_i.

On the other hand, architectures such as Eternity and DHTs scatter the red and blue resources randomly across all nodes n_i. What is the average or expected utility of each node n_i? If there exists in the system in total $\mathcal{R}$ red resources and $\mathcal{B}$ blue resources we can define a system-wide distribution of resources (r_s, b_s) that each node in the system will on average hold, with:

$$r_s = \frac{\mathcal{R}}{\mathcal{R}+\mathcal{B}} \qquad b_s = \frac{\mathcal{B}}{\mathcal{R}+\mathcal{B}} \tag{2}$$

Each node n_i will have on average a utility equal to $U(r_s T, b_s T)$. It is worth noting early on that the utility each node will attain in the random case is always lower or equal to the utility a node node has under the discretionary model:

$$U_i(r_i T, b_i T) \geq U_i(r_s T, b_s T) \tag{3}$$

For this reason the discretionary peer-to-peer paradigm will be preferred, given the choice and in the absence of other mechanisms.

It is worth exploring in more detail the implications of the lower utility provided by the random distribution model. The equality $U_i(r_i T, b_i T) = U_i(r_s T, b_s T)$ is true when $r_s = r_i$ and $b_s = b_i$, in other words when the distribution of resources in the whole system aligns with the preferences of the particular node. The first observation is that this cannot hold true for all nodes, unless they share the same preferences. Secondly it is in the self-interest of all nodes to try to tip the balance of $\mathcal{R}$ and $\mathcal{B}$ towards their preferences. With a utility function slightly more biased towards serving – which we might call an 'evangelism utility' – this could mean flooding the network with red or blue files according to their preferences.

An alternative is subversion. Red-loving nodes who consider the network overly biased toward Blue could just as easily try to deny service of Blue files, and in extremis they might try to deny service generally. Systems such as freenet and distributed hash tables can be quite prone to flooding, whether of the evangelism or service-denial variety.

A number of systems, including FreeHaven [8], Mojonation [25] and Eternity [2–4], recognised that where the utility function places more value on consumption than on service, nodes have incentives to free-ride. Eternity proposed, and Mojonation tried to implement, a payment mechanism to align the incentives that nodes have to store and serve files. By storing and serving file nodes acquire *mojo*, a notional currency, that allows them to get service from other nodes. Due to implementation failures, poor modelling and inflation [25] Mojonation provided sub-standard service and did not take off. FreeHaven took the route of a reputation system, whereby peers rate the quality of service that they provide each other and prioritise service to good providers. This is an active area of research. Perhaps such a system could be used to rate nodes using some collaborative mechanism that established r_s and b_s through voting, and then rated peers in accordance with their closeness to this social norm. This is not trivial; voting theory, also known as social choice theory, tells us that it is hard to create a voting system that is both efficient and equitable [20, 19]. The additional constraints of peer-to-peer networks – nodes frequently joining and leaving, transient identities, and decentralisation – make a 'democratic' system a non-trivial problem.

Some systems attempt to hide from the nodes which resources they are storing or serving, by encrypting them or dispersing them. This is thought to protect the nodes by providing plausible deniability against censors, but also preventing nodes deleting resources they do not like. In our framework these techniques amount to hiding from the nodes the actual distribution of red and blue resources they hold, and can even go as far as hiding the overall distribution r_s, b_s of the system. Effectively hiding this information makes these systems very expensive. The effects of the participating nodes' state of uncertainty, on their incentives to honesty participate in such a network, should be the subject of further study.

In what follows, we will for simplicity start off by considering the attackers to be exogenous, that is, external to our system.

4 Censorship

So far we have compared the utility of nodes in the random model versus the discretionary model, and have found the latter to always provide as good or higher utility for all nodes in the absence of censorship. We will now examine how nodes react to censorship.

We model censorship as an external entity's attempt to impose on a set of nodes a particular distribution of files r_c and b_c. The effect of the censor is not fixed, but depends on the amount of resistance offered by the affected nodes.

Assume that a node that is not receiving attention from the censor can store up to T resources. A node under censorship can chose to store fewer resources $(T-t)$ and invest an effort level of t to resist censorship. We define the probability that a node will successfully fight censorship (and re-establish its previous distribution of resources) as $P(t)$. With probability $1-P(t)$ the censor will prevail and impose the distribution r_c, b_c.

We first consider the discretionary case, where nodes select the content they serve. Knowing the nodes' preferences r_i, b_i, the censor's distribution r_c, b_c, the total resource bound T and the probability $P(t)$ that it defeats the censor, we can calculate the optimal amount of resources a node will invest in resisting censorship. The expected utility of a node under censorship which invests t in resisting the censor is the probability of success, times the utility in that case, plus the probability of failure times the utility in that case:

$$U = P(t)U_i(r_i(T-t), b_i(T-t)) + (1-P(t))U_i(r_c(T-t), b_c(T-t)) \tag{4}$$

Our utility functions U_i are unimodal and smooth, so assuming that the functions $P(t)$ are sufficiently well-behaved, there will be an optimal investment in resistance t in $[0,T]$ which we can find by setting $\frac{dU}{dt}=0$.

We will start with the simplest example, namely where the probability $P(t)$ of resisting censorship is linear in the defense budget t. Assume that if a node invests all its resources in defence it will prevail for sure, but will have nothing left to actually serve any files. At the other extreme, if it spends nothing on lawyers (or whatever the relevant mode of combat) then the censor will prevail for sure. Therefore we define $P(t)$ as:

$$P(t) = \frac{t}{T} \tag{5}$$

By maximising (4) with $P(t)$ being defined as in (5) we find that the optimal defence budget t_d will be:

$$t_d = \frac{T}{2}\frac{2U_i(r_c,b_c) - U_i(r_i,b_i)}{U_i(r_c,b_c) - U_i(r_i,b_i)} \tag{6}$$

The node will divert t_d resources from file service to resisting censorship. We will also assume, for now, that the cost of the attack for the censor is equal to the node's defence budget t.

So much for the case of a discretionary peer-to-peer network like Kazaa or gnutella where nodes choose the resources they distribute. We now turn to the case of Eternity or DHTs where resources are scattered randomly around the network and each node expects to hold a mixture r_s, b_s of files.

As in the previous example, the utility of a node under censorship will depend on its defence budget t, and the censor's choice of r_c, b_c, but also the system's distribution of files r_s, b_s:

$$U = P(t)U_i(r_s(T-t), b_s(T-t)) + (1-P(t))U_i(r_c(T-t), b_c(T-t)) \tag{7}$$

A similar approach is followed as above to derive the optimal defence budget t for each node:

$$t_s = \frac{T}{2}\frac{2U_i(r_c, b_c) - U_i(r_s, b_s)}{U_i(r_c, b_c) - U_i(r_s, b_s)} \tag{8}$$

However, not all nodes will be motivated to resist the censor! Some of them will find that $U_i(r_sT, b_sT) \leq U_i(r_cT, b_cT)$ i.e. their utility under censorship increases. This is not an improbable situation: in a network where half the resources are red and half are blue ($r_s = 0.5, b_s = 0.5$) a censor that shifts the balance to $r_c = 0$ will benefit the blue-loving nodes, and if they are free to set their own defence budgets then they will select $t = 0$.

5 Who fights censorship harder?

We have derived the defence budget t_d of a node in a discretionary network, and that t_s of a node in a random network. These also equal the censor's costs in the two types of network. We will now show that the aggregate defence budget, and thus the cost of consorship, is greater in the discretionary model than in the random one, except in the case where all the nodes have the same preferences (in which case equality holds).

The example in figure 2 illustrates the defence budgets of the random model versus the discretionary model.

Note that for the maximum value of the defence budget t to be positive in the interval $[0, T]$ the following condition must be true:

$$0 < \frac{T}{2}\frac{2U_i(r_c, b_c) - U_i(r_s, b_s)}{U_i(r_c, b_c) - U_i(r_s, b_s)} \tag{9}$$

in other words,

$$2U_i(r_c, b_c) < U_i(r_s, b_s) \tag{10}$$

When this is not true, a node maximises its utility by not fighting at all and choosing $t = 0$ (as illustrated in figure 3).

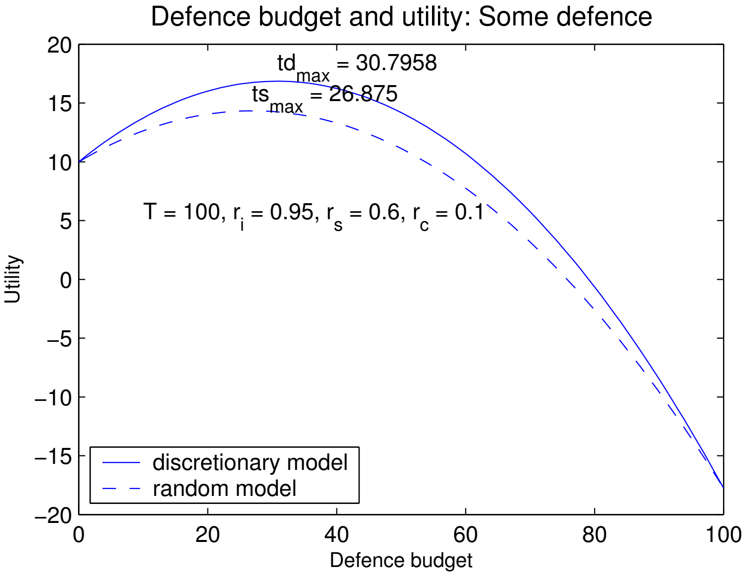

Fig. 2. Example defence budget of the two models.

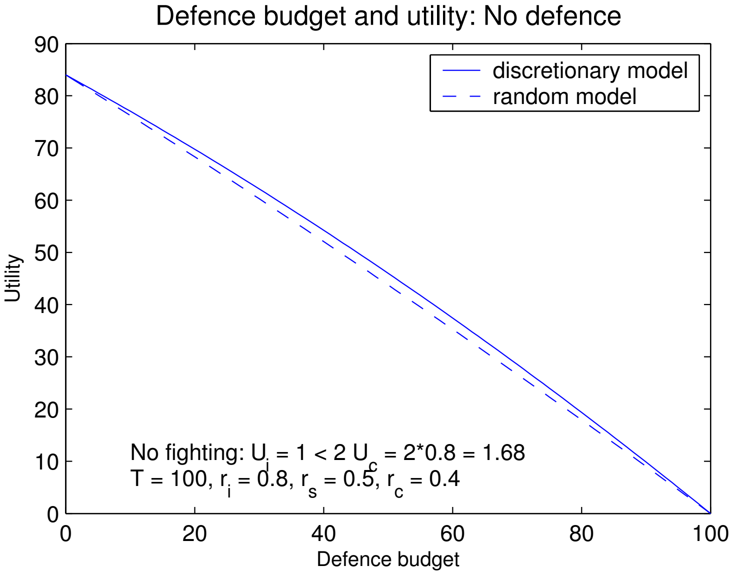

Fig. 3. Example of zero defence budget in both models.

Given the observations above it follows that:

$$\forall i \in \mathcal{S}, t_d \geq t_s \Rightarrow \sum_{i \in \mathcal{S}} t_d \geq \sum_{i \in \mathcal{S}} t_s \quad (11)$$

Whatever the strategy of the attacker, it is at least as costly or more costly to attack a network architecture based on the discretionary model, compared with the random model. Equality holds when for each node, $t_d = t_s$, which in turn means that $r_i = r_s$. This is the case of homogeneous preferences. In all other cases, the cost to censor a set of nodes will be maximised when resources are distributed according to their preferences rather than randomly.

6 Discussion

The above model is very simple but still gives some important initial insights into the economics of censorship resistance. Censorship is an economic activity; whether a particular kind of material is repressed using the criminal or civil law, or using military force, there are costs for the censor. Defence expenditure by the target (whether on lawyers, lobbying, technical protection measures or even on armed conflict) can diminish the censor's prospect of success.

Until now, much work on censorship resistance has seen censorship as a binary matter; a document is either proscribed by a court or it is not, while a technical system is either vulnerable to attack or it is not. We believe such models are as unrealistic as the global adversary sometimes posited in cryptography – an opponent that can record or modify all messages on network links. We suspect that all-powerful opponents would make censorship uninteresting as a technical issue; resistance would be impossible. Similarly, assuming that no-one can censor any nodes would provide little intuition into real systems.

Technology changes can greatly affect the parameters. For example, the introduction of moveable type printing made it much harder to suppress books thought to be heretical or seditious, and this change in the underlying economics helped usher in the modern age. It is also possible that developments such as online publishing and trusted computing may make censorship easier once more, with effects we can only guess at. Therefore trying to analyse the cost of both censorship and resistance to censorship is important. Our work presents a first model and a framework for doing this.

6.1 Preferences and utility

Modelling the node's preferences also provides important insights. Simply assuming that all nodes will fight censorship for an abstract notion of "freedom of speech" restricts the model to a fraction of potential real-world users. To take real-life examples, there have been online tussles between Scientologists and people critical of their organisation, and about sado-masochistic material that is legal under Californian law but illegal in Tennessee. The average critic of Scientology may not care hugely about sexual freedon, while a collector of spanking pictures may be indifferent to religious disputes. While there are some individuals who would take a stand on freedom of speech on a broad range of issues, there may be many more who are prepared to defend it on a specific issue.

On the other hand, assuming that nodes will not put up any resistance at all and meekly surrender any disputed documents or photographs, is also unfaithful to real-world experience. Allowing nodes to express heterogeneous interests, and preferences, when it comes to material they want to promote and protect, enables us to enhance the system's stability and security. It also enables us to defend against certain types of service-denial attack. For example, when an Eternity Service was first implemented and opened for public use [3], one of the first documents placed in it (by an anonymous poster) advocated sex between men and under-age boys. While some people would defend such speech, many would feel reluctant to use a system that expressed it. A discretionary peer-to-peer system can deal with such issues, much as ISPs currently decide which usenet newsgroups to support depending on local laws and their clients' preferencs. Thus objectionable content need not provide a universal attack tool.

Our model provides a first framework for thinking about such issues. In particular we have found that, in the presence of heterogeneous preferences, systems that distribute material randomly across all nodes are less efficient at resisting censorship than those which allow storage according to node preferences. As this inefficiency increases with heterogeneity, we expect random distribution to be more successful within groups with roughly homogeneous interests. When interests diverge, systems should either allow users to choose their resources or they will tend to be unstable. Nodes will prefer to form alternative networks that match better their preferences. Using our model we explain why while most initial peer-to-peer systems advocated random distribution of resources, the most successful ones implemented a discretionary approach.

Our model might be extended in a number of ways. Most obviously we use red and blue resources as a simple example. It is imperative that nodes can express arbitrary and much finer grained preferences, and the results we present can be generalised to unimodal multi-dimensional utility functions. We chose to model node utility locally, and did not take into account how available a resource is "globally" through the network. Forming a global view of availability is hard in many peer-to-peer systems. We also ignore the costs associated with search and retrieval. Some systems, such as distributed hash tables, allow very efficient retrieval but at a high search cost; other systems are more balanced or less efficient overall. Then again, there are various specific attacks on peer-to-peer systems, whether in the literature or in the field; the model might be extended to deal with some of them.

Random distribution may also introduce social choice problems that discretionary distribution avoids. There may be a need for explicit mechanisms to determine r_s and b_s, the relative number of red versus blue files that a typical node will be on average asked to serve. Nodes have incentives to shift these towards their preferences, they may be tempted to manipulate whatever voting or reputation systems are implemented. Making these robust is a separate topic of research.

6.2 Censorship

Our model of censorship has been carefully chosen not to introduce additional social choice issues. The censor is targeting a set of individual nodes, and the success or not of his attack on a particular node depends only on the defence budget of that node. Of course, in the case where nodes are subject to legal action, a victory against one node may create a precedent that makes enforcement against other nodes cheaper in the future. Defence may thus take on some of the aspects of a public good, with the problems associated with free-riding and externalities. How it works out in detail will depend on whether the level of defence depends on the least effort, the greatest effort or the sum-of-efforts of all the nodes. For how to model these cases, see Varian [23].

Our model also assumes a censor that wishes to impose a certain selection of resources on nodes. This is appropriate to model censorship of news and the press, but might not be so applicable to the distribution of music online. There the strategy of the music industry is to increase the cost of censorship resistance to match the retail price of the music. In that context, our model suggests that it would be much harder for the industry to take on a diffuse constellation of autonomous fan clubs than it would be to take on a monolithic file-sharing system. In that case, there might be attack-budget constraints; some performers might be unwilling to alienate their fans by too-aggressive enforcement.

Our particular censorship model provides some further insights. For both random distribution and discretionary distribution, the censor will meet resistance from a node once his activities halve its utility. This is because, for the particular utility function we have chosen, a node will react to mild censorship by investing in other resources rather than engaging in combat. So mild censorship may attract little reaction, but at some point there will be nodes that start to fight back, starting with those nodes whose preferences are most different from the censor's. This is consistent with intuition, and real-life experience.

We have not modelled the incentives of the censor, or tried to find his optimal strategy in attacking the network. Better attack models probably require more detail about the network architecture and operation.

Finally, our model may have wider implications. It is well known that rather than fighting against government regulation and for market freedom in the abstract, firms are more likely to invest effort, through trade associations, in fighting for the freedoms most relevant to their own particular trade. There is also the current debate about whether increasing social diversity will necessarily undermine social solidarity [1, 12, 16]. The relevance of our model to such issues of political economy is a matter for discussion elsewhere.

7 Conclusions

We have presented a model of node preferences in peer-to-peer systems, and assessed how two different design philosophies (random and discretionary distribution of resources) resist censorship. Our main finding is that, under the assumptions of our model, discretionary distribution is better. The more heterogeneous the preferences are, the more it outperforms random distribution. Nodes will on average invest more in fighting censorship of resources they value. Our model of censorship resistance is simple, but can be extended to explore other situations.

In the discretionary model nodes do not have to collectively manage the overall content of the network, which gives them fewer incentives to subvert the control mechanisms. This in turn allows for simpler network designs that do not require election schemes, reputation systems or electronic cash, which can be cumbersome and difficult to implement. Finally, the discretionary model also leads to a more stable network. Each node can better maximise its utility and is less likely to leave the network to seek a better deal somewhere else.

Acknowledgements The authors would like to thank Rupert Gatti and Thierry Rayna for their valuable input on the appropriateness of different utility functions to the problem studied, and the anonymous reviewers for valuable comments. This work was done with the financial assistance of the Cambridge MIT Institute (CMI) as part of the project project on 'The design and implementation of third generation peer-to-peer systems'.

References

1. David Aaronovitch. The joy of diversity. The Observer, February 29 2004.
2. Ross Anderson. The eternity service. In *1st International Conference on the Theory and Applications of Cryptology (Pragocrypt '96)*, pages 242–252, Prague, Czech Republic, September/October 1996. Czech Technical University Publishing House.
3. Adam Back. The eternity service. *Phrack Magazine*, 7:51, 1997. `http://www.cypherspace.org/~adam/eternity/phrack.html`.
4. Tonda Benes. The strong eternity service. In Ira S. Moskowitz, editor, *Proceedings of Information Hiding Workshop (IH 2001)*. Springer-Verlag, LNCS 2137, April 2001.
5. Bengt Carlsson and Rune Gustavsson. The rise and fall of napster - an evolutionary approach. In Jiming Liu, Pong Chi Yuen, Chun Hung Li, Joseph Kee-Yin Ng, and Toru Ishida, editors, *AMT*, volume 2252 of *Lecture Notes in Computer Science*, pages 347–354. Springer, 2001.
6. Ian Clarke, Oskar Sandberg, Brandon Wiley, and Theodore W. Hong. Freenet: A distributed anonymous information storage and retrieval system. In Federrath [10], pages 46–66.
7. Edith Cohen, Amos Fiat, and Haim Kaplan. Associative search in peer to peer networks: Harnessing latent semantics. In *IEEE INFOCOM 2003, The 22nd Annual Joint Conference of the IEEE Computer and Communications Societies*, San Franciso, CA, USA, March 30 - April 3 2003. IEEE, 2003.

8. Roger Dingledine, Michael J. Freedman, and David Molnar. The free haven project: Distributed anonymous storage service. In Federrath [10], pages 67–95.
9. Peter Druschel, M. Frans Kaashoek, and Antony I. T. Rowstron, editors. *Peer-to-Peer Systems, First International Workshop, IPTPS 2002, Cambridge, MA, USA, March 7-8, 2002, Revised Papers*, volume 2429 of *Lecture Notes in Computer Science*. Springer, 2002.
10. Hannes Federrath, editor. *Designing Privacy Enhancing Technologies, International Workshop on Design Issues in Anonymity and Unobservability, Berkeley, CA, USA, July 25-26, 2000, Proceedings*, volume 2009 of *Lecture Notes in Computer Science*. Springer, 2001.
11. Philippe Golle, Kevin Leyton-Brown, and Ilya Mironov. Incentives for sharing in peer-to-peer networks. In *ACM Conference on Electronic Commerce*, pages 264–267. ACM, 2001.
12. D Goodhart. Discomfort of strangers. The Guardian, Feruary 2 2004. pp 24–25.
13. Johan Helsingius. Johan helsingius closes his internet remailer. `http://www.penet.fi/press-english.html`, August 1996.
14. Jack Hirshleifer. *The dark side of the force*. Cambridge University Press, 2001.
15. Kazaa media desktop. Web site, 2004. `http://www.kazaa.com/`.
16. The kindness of strangers? The Economist, February 26 2004.
17. Qin Lv, Sylvia Ratnasamy, and Scott Shenker. Can heterogeneity make gnutella scalable? In Druschel et al. [9], pages 94–103.
18. Antony I. T. Rowstron and Peter Druschel. Pastry: Scalable, decentralized object location, and routing for large-scale peer-to-peer systems. In Rachid Guerraoui, editor, *Middleware*, volume 2218 of *Lecture Notes in Computer Science*, pages 329–350. Springer, 2001.
19. A Sen. Social choice theory. In K Arrow and MD Intriligator, editors, *Handbook of Mathematical Economics*, volume 3, pages 1073–1181. Elsevier, 1986.
20. Andrei Serjantov and Ross Anderson. On dealing with adversaries fairly. In *The Third Annual Workshop on Economics and Information Security (WEIS04)*, Minnesota, US, May 13–14 2004.
21. Ion Stoica, Robert Morris, David Kargerand M. Frans Kaashoek, and Hari Balakrishnan. Chord: A scalable peer-to-peer lookup service for internet applications. In *ACM SIGCOMM 2001 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communication*, pages 149–160, San Diego, CA, USA, August 27-31 2001. ACM.
22. Chunqiang Tand, Zhichen Xu, and Sandhya Dwarkadas. Peer-to-peer information retrieval using self-organising semantic overlay networks. In *SIGCOMM'0*, Karlsruhe, Germany, August 25-29 2003. ACM.
23. H Varian. System reliability and free riding. In *Workshop on Economics and Information Security*, University of California, Berkeley, May 2002. `http://www.sims.berkeley.edu/resources/affiliates/workshops/econsecurity/econws/49.pdf`.
24. Weatherspoon, Westley Weimer, Chris Wells, and Ben Y. Zhao. Oceanstore: An architecture for global-scale persistent storage. In *ASPLOS-IX Proceedings of the 9th International Conference on Architectural Support for Programming Languages and Operating Systems*, volume 28 of *SIGARCH Computer Architecture News*, pages 190–201, Cambridge, MA, USA, November 12-15 2000.
25. Bryce Wilcox-O'Hearn. Experiences deploying a large-scale emergent network. In Druschel et al. [9], pages 104–110.