Enforcing Behaviour with Anonymity

Joss Wright, Susan Stepney
University of York
Heslington, York. YO10 5DD
Email: {joss,susan}@cs.york.ac.uk

***Abstract*—We discuss applications of an underlying anonymous infrastructure to enforce fair behaviour on participants in a distributed resource-sharing system. This approach aims to prevent users from forming self-rewarding cliques in order to gain unfair advantages in the use of shared resources.**

We deliberately avoid considering the more traditional applications of anonymous systems in an attempt to show the potential for the use of restricted access to identifying user information in applications where privacy is not the main motivation.

We also briefly explore the problem of enforcing anonymity on users who may not wish to be anonymous, and consider the effect that low-level identification may have on the overall behaviour that we seek to enforce.

I. Introduction

Research into anonymity systems typically focuses on the provision of anonymity as a goal in itself, and relies on a small number of well-known justifications for the creation of such systems. Among the most common aims in anonymity research are the development of censorship resistant systems, and the general desire for privacy in a world which is increasingly under surveillance. While these are certainly sufficient to justify the existence of systems that provide anonymity, it is useful to consider more general applications for these systems.

Recent developments in anonymity research have produced comparatively robust and efficient anonymous communication systems. This raises the question of what uses they may have beyond the basic goal of hiding the identity of their users. The potentially useful side effects of anonymity systems are the focus of the work presented here.

II. Contribution

In this discussion paper we propose anonymity as a mechanism to enforce fair behaviour in distributed computing systems. This approach relies on an underlying anonymous system's ability to prevent users from distinguishing other users, either for the purpose of creating beneficial partnerships or for exploiting users outside a group.

We consider the possibility of taking this approach, and the potential advantages and disadvantages of schemes for enforcing fairness based on anonymity in the network. We discuss some difficulties that must be overcome in order for these schemes to be practical. We also briefly examine the problem of enforcing a level of anonymity on users that actively desire to identify themselves.

This paper is intended as an exploration of achieving fairness through anonymity as the basis for future research. As such we do not present detailed results or propose specific solutions to all of the problems we identify, but instead show a number of open issues and practical considerations that are worthy of investigation.

III. Background

Serious work into computer anonymity systems began in 1981 with Chaum's development of the mix system for untraceable message sending [1]. This was followed in 1988 by the dining cryptographer protocol [2] that provides provable anonymity for participants under certain stringent assumptions. Since then, there has been a great deal of work into systems for providing unlinkability between senders and recipients in communicating networks. Much work has been conducted into variations on the basic mix design, along with a wide variety of attacks and quantifications of the effectiveness of different approaches. Recent work has tended to focus on low-latency communications suitable for interactive services such as web browsing, the most well-known example of which is Tor [3].

In this paper we are concerned with applications that overlay these anonymity services, and therefore do not consider here the entire history of research into anonymity systems. We instead refer the reader to an appropriate survey such as [4].

In recent years anonymity systems have become increasingly mature, and have spread to encompass a user-base sufficient to demonstrate their feasibility for real-world usage; anonymous traffic networks are now a serious proposition. It may be considered that the basic issues with anonymising communications have been addressed sufficiently that their use for more mainstream applications are worthy of consideration. It is one such application that we propose in this paper.

The desire for anonymity is not a new phenomenon and has historically been employed for a variety of purposes, both beneficial and otherwise. With the advent of the Internet, a number of proposed systems have arisen that directly exploit anonymity for one purpose or another. Here, we briefly examine some relevant work.

A. Voting systems

Voting is one of the most important and widespread applications of anonymity, both in new approaches to electronic voting and in its more traditional incarnation. Importantly, voting is also recognised as a "legitimate" application of anonymity and does not suffer from many of the objections levelled against more general anonymity services.

The importance of secret ballots in voting has been recognised for hundreds of years, and their use is now generally

considered a requirement for fair elections. Anonymous voting protects voters from intimidation or coercion, as well as preventing the possibility of vote-buying. This functions to protect both voters *and* candidates by preventing both coercion employed against voters and of accusations of coercion levelled at candidates. Given the huge importance that the outcome of elections can have at the national level, voting is almost certainly the most influential application of anonymity systems at present. As such, electronic voting is an active area of research.

B. The Eternity Service

In this paper we make an effort not to base our proposals on those systems in which privacy or anonymity is a goal in itself. However one notable case where the need for anonymity naturally combines with other practical considerations is that of the Eternity Service, as presented in [5].

The Eternity Service is based on a very simple concept: that distributed storage of information is resistant to denial of service attacks, and that this can be used to ensure that published material remains published.

In the basic Eternity Service, a user wishing to publish data distributes their content to a large number of remote data storage servers that make certain guarantees regarding reliability and availability. This transfer of data is accompanied by the creation of an annuity payment in order to provide an economic incentive for servers to retain data for a given time period. Once the data has been transferred to the required number of servers, the identities of a small subset of these servers are retained by the user and the rest are deleted.

The retained identifiers can be periodically checked by the user to ensure that their data is reliably stored. As the set of stored server identifiers is randomised, no server can silently drop data without risking discovery.

Deleting the large majority of identifiers prevents a user from being forced to remove a document that they have published, as well as preventing a third party from being able easily to censor the data by learning the location of all copies. By deleting the identifiers of the majority of servers storing the content, the publisher no longer has knowledge of the location of all copies and cannot be forced to reveal them.

The idea of the Eternity Service has served as an inspiration for several censorship-resistant systems, notably Publius [6].

C. The cocaine auction protocol

The cocaine auction protocol [7] deals with the hypothetical problem of ensuring a fair, and safe, auction between rival drug dealers for a shipment of narcotics. This is achieved by use of anonymous broadcast of secrets, and deals largely with the issue of protecting users from each other whilst allowing the seller and winning buyer to complete their transaction in secret.

The design of the protocol aims to ensure the anonymity of each bidder from their peers for the purposes of safety. Of note, however, is that the protocol also prevents a seller from learning the identity of the winning buyer before a committing to a sale, and allows a winning bidder to prove that they won a bid if the seller unfairly tries to sell to another, lower-bidding, buyer.

The basic form of the cocaine auction protocol is implemented over a short-range wireless broadcast network. It is assumed that this medium achieves a sufficient level of untraceability to ensure that each user's broadcast cannot be easily linked to them.

The auction takes place over a number of rounds of bidding. At the start of each round the seller broadcasts a desired price; if a bidder is willing to pay that price, they broadcast the result of a one-way function on a freshly-generated nonce. The first observed bid for a round is taken as the winner of that round, whereupon the price is incremented and a new round begins. If a bid is not received for any given round, the winner of the previous round wins the auction.

The winning bidder can prove their identity to the seller by producing the nonce used to generate their bid. The use of a one-way function ensures that only the bidder can know this value. In order that this verification need not take place immediately, the seller can use the information in the bid to initiate a Diffie-Hellman key exchange with the bidder and broadcast details of a secret rendezvous under the newly negotiated shared key.

This approach ensures the anonymity of buyers until the end of an auction and thus prevents discrimination against, or favouritism towards, particular buyers.

There are a variety of other features imparted by anonymising users, and there are many alternative schemes and systems that make use of anonymity. However, it is this ability of anonymous systems to prevent collaboration or discrimination that we examine in detail.

IV. Enforcing Fairness

Anonymity, in its traditional sense, is already a widely-used method for enforcing behaviour in situations where knowledge of individual users could lead to favouritism or discrimination. The marking of exam papers, review of funding applications and the increasing use of double-blind academic peer-review are all well known examples where anonymity is employed to prevent undesirable bias from entering a judgement process.

In these situations anonymity prevents individuals in the system from making, consciously or unconsciously, prior judgements. Examiners familiar with the authors of exam scripts may be more lenient towards favoured students or, equally harmfully, less lenient towards others; funding applications from well-known institutions may be treated favourably based on past connections with the awarding body. Anonymising these requests, even if imperfectly, reduces the potential for favouritism.

Despite the involvement of more complex technology, the cocaine auction example presented in Section III-C is still largely concerned with affecting the behaviour of the humans in the protocol; it is a concern if the seller exhibits biased behaviour due to the identity of the buyer, or if buyers exhibit

violent behaviour based on the identifiable bids of other users. The protocol overcomes this element.

We may extend this concern, however, to situations in which clients behaving in a network are deliberately programmed to take advantage of identifying information to cheat a protocol that is intended to be fair. Further, we may be concerned less with dramatic behaviour that swings the outcome of a single decision from one side to another. In more complex scenarios, a user may exploit subtle variations in behaviour in order to gain a statistical advantage over other users without being easily detectable.

In such systems, as in the case of the cocaine auction protocol, we are concerned with the distributing of some resource amongst a number of participants. A useful example is that of certain peer-to-peer systems in which the combined bandwidth of all users accessing some data is shared out for more effective distribution. A similar concept is that of distributed computing grids, where large amounts of processor time are shared between users.

The behaviour that we focus on preventing is that of groups of users collaborating in order to exploit the system. In sharing resources it is possible for users to form preferential cliques that offer greater bandwidth to other members while dropping or assigning low priorities to requests from members outside the clique. This allows the colluding parties to monopolise a resource whilst minimising their own contribution.

The potential for such an approach is mentioned, to some extent, in [8] and has been partially implemented in a number of clients for the BitTorrent network that recognise other users of the same client and behave preferentially towards them.

The purpose of applying anonymity, therefore, to a system in the fashion described in this paper is to enforce a level of *fairness*. This term has a variety of interpretations under different assumptions, however we are largely concerned with the inability of any subset of users to gain a greater access to resources than that intended by the system. This does not necessarily ensure that any one user receives an equal share of some resource, or that all users contribute equally, although this may be the case. Many systems simply attempt to maintain a ratio between resources contributed and resources consumed by users; in these systems any user behaviour that unbalances this ratio may be considered unfair.

Another consideration is whether preferential behaviour has detrimental effects for other members of the network. One claim of the work described in [8] is that the efficiency of the network as a whole may be improved if all clients adopt the same behaviour.

Whilst anonymity provides a mechanism to overcome collaboration between users, it at the same time allows for users to behave in an undesirable fashion with a negligible chance of being caught. This results in many of the well-known problems associated with anonymity systems, such as denial of service, resource poisoning and the potential for an individual to make multiple registrations in order to control a larger portion of the network [11].

As implied above, a system's method of defining "fairness" may leave the potential for exploitation. Many simple peer-to-peer content distribution systems enforce fairness based on the volume of data contributed. This ensures, to some extent, that a user does not "leech" from the system by simply downloading content while refusing to reciprocate with uploads.

This approach leaves itself open to the collaboration of individuals who maintain high upload bandwidth between themselves to more rapidly disseminate data obtained from the rest of the network. By doing so, collaborating users can maintain an "honest" ratio of contribution to consumption whilst still benefiting from dishonest behaviour. As we have already noted, this need not necessarily harm the network as a whole.

A. Optimal Behaviour

A critical element of the approach presented in this paper is that of drastically reducing the information available to an actor for making decisions. By removing or restricting information, the system alters the ability of the user to make decisions; the perceived optimal course of action for a user may change drastically given more or less information. In the context of anonymous systems, the lack of identifying information prevents an actor from incorporating identity into their decisions regarding behaviour towards peers.

Collaboration occurs as the result of a user perceiving a benefit from behaving in a particular way towards a particular user. By removing the ability to distinguish between peers, the optimal behaviour for a user is to treat all peers equally. This approach extends to other sources of information beyond identity, however we choose to focus on this feature.

V. Applicable Systems

We consider applications that overlay an anonymous communication system. As such we assume that an appropriate system for anonymising communications is already in place and do not discuss specifics of the underlying network. One important assumption that we make, however, is that participating nodes do not choose their communication partners directly but are instead randomly provided with them by the underlying anonymous network.

This is a large assumption, and hides a number of serious issues regarding routing, reliability and security weaknesses; however we do not address these here. One approach to satisfying this constraint could be for actors to select their communication partners randomly from a list of short-term pseudonymous identifiers in the manner of private information retrieval systems [9]. The crucial aspect is that it must be infeasible for two participants in a transaction to rediscover each other once a communication has ended.

In what types of system could we apply such a scheme? Any system in which some set of users have access to a resource that is to be shared has the potential for discrimination and cheating. Often, this can be overcome by centralised resource management or tracking, or by reputation-based systems that maintain persistent statistics for each node. Each of these

follows the traditional model of strong, and largely global, identification.

We propose that in a more chaotic network, in which centralised servers are infeasible or undesirable, an anonymity mechanism can provide an alternative approach. Of course, in a system where anonymity is desirable for other reasons we may take advantage of the potential for fairness inherent in the system.

A. The human element

Anonymity provides an approach to fairness in two forms of system: those that rely on human judgements, in which anonymity can be used to override irrational human decisions; and those that rely on automated judgements, in which case anonymity prevents users from making use of identifying information to manipulate the system. In later sections we focus on the issues surrounding the second of these, where human psychology is not a major issue, but here mention issues related to both types.

Many of the traditional applications of anonymity have aimed to prevent humans from the potential for making biased decisions. The anonymisation of data in exam marking and double-blind peer review have already been mentioned as systems in which identifying information may prevent fair decisions from being made.

For these examples, almost uniquely for anonymity when considered as a security property, the parties from whom data is being hidden are typically willing participants in the anonymisation. As in the case of elections, it is generally to the benefit of both voters and candidates for anonymity to be preserved. Consequently, there is often no "attacker" in the traditional sense and the breaking of anonymity is likely to occur by accident if at all[1].

In systems that rely on human judgement, bias may come about unintentionally. An examiner, on recognising a known good student's script, may provide the benefit of the doubt when an answer is questionable. Alternatively, the examiner may be unusually strict towards that student in an attempt not to show favouritism.

Human judgements are particularly vulnerable to faulty or irrelevant data, especially in giving exaggerated weighting to personal experience. A prospective customer consulting a system that ranks customer satisfaction with companies may be unwilling to trade with a company with whom they have had a single poor experience in the past. This bias may often override a large number of positive reviews from other users. Anonymisation of company names in such a system would, in a rather sinister fashion, force a user to base their choice solely on the information that the system deems important. By drastically restricting the information available to a user for decision-making, a system can prevent decisions from being made according to "unauthorised" criteria.

Human psychology causes extremely complex behaviour that is hard to model, and is often open to exploitation by actors that behave rationally. If human judgements are critical to an application this exploitation can have significant effects, but these are outside the scope of this paper. We therefore focus on situations in which anonymity may be used to enforce behaviour in a stricter sense.

[1]This ignores the possibility of students seeking to identify themselves to their examiner in order to exploit their good reputation.

B. Automated systems

As we have implied, anonymity is of use beyond preventing the effects that identifying knowledge may have on human judgement. In computer systems, the identification of other members of a network allows for a variety of behaviours in which nodes collude, or discriminate, in order to exploit more effectively the available resources. By restricting the information available to a peer, it is possible to force their optimal behaviour to be in line with that that is most beneficial for the system as a whole.

To return to the example of a peer-to-peer content distribution system, identification allows users to select peers with whom they preferentially distribute data. As has been noted, it may be possible that such an "unfair" approach allows for greater efficiency in the network as a whole.

The potential effects of favourably selecting peers have been demonstrated in an adapted client for the BitTorrent network, known as BitTyrant [8], in which the authors demonstrate that preferentially allocating bandwidth to those peers that have uploaded more to the current node can improve efficiency of that node's connection. The paper does not, however, aim to maximise bandwidth at the cost of other nodes on the network and stops short of considering actively malicious, or even entirely selfish, nodes.

Apart from distribution of data, the distribution of processing power is becoming increasingly popular. Grid computing is an active area of research, with distributed computing clusters such as BOINC [10] being well established. While there are not currently examples of these systems that take a fully distributed peer-to-peer approach, it is not unreasonable to assume that such systems may appear in the future. The application of anonymity as a method for fairness in these networks are similar to those of the peer-to-peer content distribution networks discussed earlier. We may also apply these ideas to distributed storage, as well as numerous other systems.

VI. Considerations with anonymous systems

There are many considerations in applying a fully anonymous systems. The ability to distinguish users, to maintain repeated connections with them, and to make judgements based on past interactions are all fundamental tools in forming networks that become useless with the removal of identifying information.

Of course, the advantage most often associated with the use of anonymity systems is anonymity itself, and this is not surprising. Even the most usable current anonymity systems impose a penalty for their use in terms of speed or computation, and notably in causing well-behaved users to be indistinguishable from misbehaving elements. Many Internet

services habitually block traffic from networks such as Tor. The cost of using anonymity systems is often sufficiently high that the only advantage valuable enough to make them worthwhile is anonymity itself.

We examine here some of the other advantages that arise from the use of an anonymous network, as well as the disadvantages introduced by this approach.

A. Advantages

The advantage of anonymity explored on in this paper is the inability of dishonest users to collaborate effectively when their identities are unknown to each other. Anonymity prevents the formation of cliques in a network, or the exploitation of such cliques if formed outside the network. In systems where resources are shared there is great potential for collaboration between dishonest users that is detrimental to the network as a whole; anonymity provides one solution to maintaining the balance.

Another side to this indistinguishability is that it becomes difficult to harm individuals directly. This allows for the well-known censorship-resistant properties of anonymity systems that are mentioned above. In a purely anonymous network, the only way for a node to "punish" another node is to refuse to deal with it on a per-connection basis. This allows for a self-regulation of connections that does not rely on a centralised policy.

B. Limitations

Networks that seek to enforce fairness by anonymising connections inherit many of the issues that hamper other anonymous networks. Denial of service is arguably the most severe problem, as it is almost impossible to block or drop users from the network. Each new connection must be considered individually, making it impossible to blacklist misbehaving users.

Similarly, systems built on anonymous networks suffer from so-called *Sybil attacks*, as detailed in [11]. These attacks occur when an individual creates a large number of accounts in some anonymous system, and can therefore come to control a significant portion of the network. The anonymity of users makes these attacks hard to prevent, or even detect.

A fundamental issue in applying anonymity to enforce any form of user behaviour is the potential identification of users, either by their own efforts or through the more traditional traffic-analysis of attackers. The loss of anonymity could have serious effects for a network that relies on anonymity for its overall behaviour. If fairness is enforced by requiring users to remain anonymous, the ability of a small number of users to communicate reliably with each other could leave a system much more susceptible to the very behaviour that it is designed to prevent.

Alternatively, the ability of a small set of users to identify themselves may not be sufficient to ruin the overall fairness of the network. This is particularly the case if identification requires a significant expenditure of time or resources.

This last consideration is one of the more unusual considerations in the application of anonymity in enforcing fairness: it is difficult to anonymise users against their will, as we discuss in Section VIII.

We now examine some specific approaches to the use of an underlying anonymity systems as a method for enforcing fair user behaviour.

VII. Incentives for honesty

Anonymity in the connection between users ensures that distinctions cannot be made between individuals. The positive effect on which we have focused so far is the inability of actors to treat certain actors preferentially.

For a dishonest user, the simplest response when seeking to gain from a system is to consume equally from other users, but to refuse to contribute. The anonymity of the system makes this behaviour difficult to penalise in the traditional sense, as the non-contributing user cannot easily be distinguished.

In order to enforce "honest" behaviour in anonymous nodes there must be some incentive for each host to contribute to the network. In systems with strong identification it is possible to associate an actor with a reputation based on their contribution to the system, and to grant access to resources based on this reputation. Distributed processing systems can store the amount of CPU hours contributed by an actor; storage systems can judge the amount of data stored for other actors. In an anonymous system these associations become largely impossible and it is necessary to find alternative approaches to ensuring contribution.

One approach is to "embed" incentives into the system in such a way that it is in the users' best interest to contribute, but without the requirement for a global reputation system. We now discuss this possibility.

A. Per-connection ratios

The simplest method to approach this problem is for each node to judge the contribution of each peer based only on information observed during the current connection with that peer. This approach is effective in those systems where peers consistently trade some resource, such as in peer-to-peer content distribution systems in which the ratio of upload to download bandwidth can be maintained.

The method is restricted, however, to systems where the connection itself provides sufficient evidence of work performed. For peer-to-peer content systems, each node can easily observe the ratio of data exchanged with their peer partners. For other systems, such as the long-term storage required by the Eternity Service, or the processing workload in grid systems, this simple approach is not sufficient. The transfer of information to a node does not guarantee that the data will be stored for the appropriate length of time; it may simply be deleted by the receiving node when the transfer is complete. Similarly, a distributed processing node may accept a particular unit of work to process, then silently drop it without appropriate processing. Such systems require some

method for verifying that both parties have fulfilled their obligations without relying on a persistent identifier.

This leads us to consider more complex schemes for maintaining a balance between resources contributed and resources consumed.

B. Artificial symmetry

An ideally anonymous network, as viewed by a participant, is partitioned into "self" and "not-self". A user will clearly not act in such a way that they are themselves harmed in the long run[2]. A dishonest user participating in the network wishes to gain some benefit from its interaction with the system while contributing as little as possible. In many cases, ignoring incoming requests for resource usage, such as storage or processing, is unlikely to cause sufficient harm to the user as the likelihood of receiving their own requests is sufficiently low that they easily gain a net benefit.

It would be ineffective for a distributed storage network such as the Eternity Service to ensure that users contribute to the overall data store by allowing portions of their own data to be randomly routed back to them. A user that drops all incoming storage requests could rely on their data being distributed across the rest of the network; indeed, the design of the Eternity Service specifically defends against data being dropped in such a fashion. This prevents censorship caused by attacking individual nodes or users, but allows anonymous users to exploit the system with impunity.

A more complex approach to ensuring honest behaviour in anonymous networks is to address this asymmetry. The work performed on behalf of a particular user by the system must be made dependent on work performed for others. This is simple in systems where a user can be associated with the work that they have performed. In anonymous systems the problem is much more complex and we do not propose a solution here.

Some work into this area has been presented in [14], relying on a relatively complete e-cash implementation. An alternative and possibly simpler approach is to make the benefit that a user gains from the system dependent on certain "critical" portions of work that are anonymously redirected to the originating user when requesting access to resources. In order to guarantee that their own workload is processed, the user must process all incoming resource requests or risk disproportionately harming themselves by deleting critical data. This approach has the potential for introducing security or reliability risks, however has the advantage of effectively embedding the incentive mechanism into the network.

A simple example of this would be for data stored in an eternity-service style system to be encrypted with a key that is routed back to the originator without their knowledge. That user, wishing their own data to remain available, cannot silently drop incoming data storage requests without risking invalidating their own stored data. This creates a simple point of attack for an Eternity Service-style system, however the approach may be applicable in other forms of distributed resource sharing.

[2]It has been proposed in [12] that allowing participants to revoke the reputation of another user at the cost of their own reputation allows for an effective reputation management system in dynamic networks. This "suicide for the common good" approach is more commonly seen, in a less drastic sense, in human interactions where participants will willingly harm themselves to prevent perceived unfair behaviour by their partner. This behaviour, reciprocity, is a well-known effect in social science and economics [13].

VIII. Enforcing Anonymity

The history of research into anonymity since Chaum's mix was presented in 1981 demonstrates that providing anonymity to users is a difficult problem. Slight variations in user behaviour allow attackers to draw identifying information from traffic patterns, and there are an endless number of side channels and sources of information leakage that serve to link users to their actions.

Research into anonymity systems typically assumes that the anonymised users wish to be anonymous. This assumption is partially due to the natural view of anonymity as a method of somehow protecting users, and partially due to the sheer number of sources of information leakage in all but the most restrictive systems. Any side channel provides a means for users wishing to identify themselves to succeed almost trivially.

One area that considers the enforcement of anonymity is electronic voting, where there is a requirement for *receipt-freeness* [15]. As has been noted, in order to ensure resistance against coercion in such systems it must be impossible for a voter to obtain proof of their own vote that can be presented to a third party. The requirement for receipt-freeness has many parallels with attempts to enforce anonymity in systems such as those we discuss here.

The attempt to enforce anonymity in spite of malicious users is also seen in some censorship-resistant systems. In [16] the author describes a scheme that prevents clients from learning the identities of the servers that store their data, even if the client attempts to discover this information during publication. This breaks the implicit assumption that the client wishes the location of their data to remain hidden, however does not consider clients that actively seek to identify themselves as the publishers of data.

The use of anonymity as a method to enforce fair behaviour in a network, however, creates a situation in which actors may seek to identify themselves to desired partners and maintain these connections despite the underlying anonymity of the communication medium. This radically shifts, and to some extent reverses, the attacker model that is traditionally considered in anonymity research.

In attempting to mutually identify with a desired partner in the systems we consider, there are two features of an anonymous network that must be overcome: authentication and discovery.

A. Authentication

The simplest, indeed almost trivial, problem for the self-identifying user to overcome is how to authenticate themselves to a connected peer. In a network that does not restrict communications, the user may employ any one of a wide

variety of authentication protocols. Shared secrets or public keys can be negotiated and distributed out-of-band across the general Internet.

Even in a system designed to prevent or detect explicit "unauthorised" handshakes, it is possible to hide authentication in normal requests, particularly in systems where general data or processing tasks are communicated. In a system in which data packets are somehow restricted or subject to scrutiny, pre-agreed delays in the timing of handshake packets or variations in the flow of traffic provide ample hidden channels for dishonest users that are all but undetectable to honest partners. It is highly unlikely that anonymity can be enforced at this level in any useful system.

B. Discovery

The more difficult problem for self-identifying users to overcome is that of discovering members of their clique with whom to collude. If we assume a network that prevents routing to specified communication partners then it becomes difficult for users who wish to collaborate to make initial contact.

One solution to this problem is to employ the "secret handshake" approach mentioned in the previous section, but to refuse to communicate with those partners who do not reciprocate: an actor seeking to form a partnership with a member of their clique, on receiving a connection from some new partner, initiates a secret authentication process. If the partner is a member of the clique then the secret handshake is completed and communication is established. If the handshake is not completed, the requester simply drops the connection and requests another partner. With sufficient repetition, this will eventually result in pairing with a member of the desired group of users.

This approach quickly becomes infeasible for all but the smallest networks (or the largest cliques). Simple countermeasures, such as the underlying network to impose a small delay on requests for new communication partners, are possible but unlikely to be necessary. With these constraints, the most feasible method for partners to communicate is to make use of out-of-band communications in an attempt to overcome the routing imposed by the network. These attacks, however, will rely on implementation details of individual systems and so are outside the scope of this paper.

IX. Conclusions

In this paper we have discussed the concept of using network-wide anonymity as a method to control user behaviour. We have described the nature of the systems where this approach may be used, and shown some of its associated advantages and disadvantages. We have also presented some simple attacks and considerations that must be resolved.

The potential for using anonymity as method to control user behaviour relies on the ability of the system to enforce some level of honest behaviour in nodes that cannot be identified. This lack of identity prevents reliance on centralised servers to regulate network behaviour, and "punishment" of misbehaving participants according to their long term behaviour. A significant requirement for maintaining this behaviour is the prevention of nodes from discovering and identifying each other in spite of the underlying anonymous network.

We have focused on systems for which anonymity and privacy are not a goal in themselves, however the approaches that we propose here also demonstrate a useful intrinsic feature of existing anonymous systems. The exploitation of these features may have benefits for systems that rely on anonymity for other purposes.

X. Future Work

The purpose of this paper is to present a number of ideas of interest based on the concept of applying anonymity as a method to enforce fairness. As such there are many open questions arising from this work.

The most critical issue is clearly that of maintaining anonymity in the face of unwilling participants. The issues surrounding this enforcement are complex. It is clear that forcibly anonymising users who are determined to identify themselves is ultimately doomed to failure in most non-trivial systems, however there is potential for approaches that associate a high cost with mutual *discovery* by specific users. This raises the question of how robust these approaches are against against low levels of identification.

The potential for users in distributed systems to form self-rewarding cliques has not yet been fully explored, although the concept is treated to some extent in [8], among other places. The effectiveness of small groups of users in maximising their benefit at the expense of the network is worthy of exploration, as is the potential for more malicious behaviours such as selective denial of service and poisoning of resources. These other success criteria are important in systems aimed at goals such as censorship-resistance, and anonymity is likely to prove more of a benefit than a hindrance to attackers in this case.

Even for nodes that seek to use the system to their advantage, and so have an interest in the functioning of the network, forcing honest behaviour for individual nodes that may anonymously misbehave is not trivial. The attempt to create an artificial symmetry between resources consumed and resources produced, as mentioned in Section VII-B, poses some interesting problems.

More widely, the approach of deliberately designing network-level behaviour through simple rules based on highly restricted information, rather than enforcing such behaviour directly through centralised servers, is of great interest. This approach to network behaviour could have application in more chaotic networking technologies such as peer-to-peer and mesh networking in which networks are dynamic and lacking in network-wide connectivity and routing information.

Beyond this, there are still many well-known issues to be solved in anonymous networks when applied for general purposes. Reliability and dependability are difficult to ensure against users whose only interest is to harm the system. Denial of service, poisoning of shared resources and similar attacks are, and are likely to remain, difficult problems for any network in which users may act anonymously.

REFERENCES

[1] D. Chaum, "Untraceable electronic mail, return addresses, and digital pseudonyms," *Communications of the ACM*, vol. 4, no. 2, February 1981.

[2] ——, "The dining cryptographers problem: Unconditional sender and recipient untraceability," *Journal of Cryptology*, vol. 1, pp. 65–75, 1988.

[3] R. Dingledine, N. Mathewson, and P. Syverson, "Tor: The second-generation onion router," in *Proceedings of the 13th USENIX Security Symposium*, August 2004.

[4] G. Danezis and C. Diaz, "A survey of anonymous communication channels," Microsoft Research, Tech. Rep. MSR-TR-2008-35, January 2008.

[5] R. Anderson, "The eternity service," in *Proceedings of Pragocrypt '96*, 1996.

[6] 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*, August 2000, pp. 59–72.

[7] F. Stajano and R. J. Anderson, "The cocaine auction protocol: On the power of anonymous broadcast," in *Information Hiding*, 1999, pp. 434–447.

[8] M. Piatek, T. Isdal, T. Anderson, A. Krishnamurthy, and A. Venkataramani, "Do incentives build robustness in bittorrent?" in *NSDI'07*, Cambridge, MA, April 2007.

[9] L. Sassaman, B. Cohen, and N. Mathewson, "The pynchon gate: A secure method of pseudonymous mail retrieval," in *Proceedings of the Workshop on Privacy in the Electronic Society (WPES 2005)*, Arlington, VA, USA, November 2005.

[10] D. P. Anderson, "Boinc: A system for public-resource computing and storage," in *GRID '04: Proceedings of the Fifth IEEE/ACM International Workshop on Grid Computing*. Washington, DC, USA: IEEE Computer Society, 2004, pp. 4–10.

[11] J. R. Douceur, "The sybil attack," in *IPTPS '01: Revised Papers from the First International Workshop on Peer-to-Peer Systems*. London, UK: Springer-Verlag, 2002, pp. 251–260. [Online]. Available: http://portal.acm.org/citation.cfm?id=687813

[12] J. Clulow and T. Moore, "Suicide for the common good: a new strategy for credential revocation in self-organizing systems," *SIGOPS Oper. Syst. Rev.*, vol. 40, no. 3, pp. 18–21, 2006.

[13] E. Fehr and S. Gächter, "Fairness and Retaliation: The Economics of Reciprocity," *The Journal of Economic Perspectives*, vol. 14, no. 3, pp. 159–181, 2000.

[14] D. Figueiredo, J. Shapiro, and D. Towsley, "Incentives to promote availability in peer-to-peer anonymity systems," in *ICNP '05: Proceedings of the 13TH IEEE International Conference on Network Protocols*. Washington, DC, USA: IEEE Computer Society, 2005, pp. 110–121.

[15] S. Delaune, S. Kremer, and M. Ryan, "Coercion-resistance and receipt-freeness in electronic voting," *csfw*, vol. 0, pp. 28–42, 2006.

[16] A. Serjantov, "Anonymizing censorship resistant systems," in *IPTPS*, 2002, pp. 111–120.