Incentives to Promote Availability in Peer-to-Peer Anonymity Systems

Daniel Figueiredo
Dept. of Computer Science
University of Massachusetts
Amherst, MA 01002
ratton@cs.umass.edu

Jonathan Shapiro
Computer Science & Eng.
Michigan State University
East Lansing, MI 48824-1226
jshapiro@cse.msu.edu

Don Towsley
Dept. of Computer Science
University of Massachusetts
Amherst, MA 01002
towsley@cs.umass.edu

Abstract

Peer-to-peer (P2P) anonymous communication systems are vulnerable to free-riders, peers that use the system while providing little or no service to others and whose presence limits the strength of anonymity as well as the efficiency of the system. Free-riding can be addressed by building explicit incentive mechanisms into system protocols to promote two distinct aspects of cooperation among peers—compliance with the protocol specification and the availability of peers to serve others. In this paper we study the use of payments to implement an incentive mechanism that attaches a real monetary cost to low availability. Through a game theoretic analysis, we evaluate the effectiveness of such an incentive, finding that peer availability can be significantly increased through the introduction of payments under many conditions. We also demonstrate how a payment-based incentive that preserves anonymity can be implemented and integrated with a popular class of P2P anonymity systems.

1 Introduction

Early anonymity protocols [25] and some recent ones [13] rely on a core architecture in which users send messages through a relatively small set of highly available forwarding nodes. Due to scalability concerns as well as doubt about the commercial viability of an anonymous core network[1], there has been considerable recent attention paid to peer-to-peer (P2P) architectures, in which users cooperatively provide anonymity by forwarding messages for each other [19, 26, 27].

P2P anonymity systems have the potential to offer anonymous communication to very large user populations. However, as with other P2P systems, the ability of such protocols to function correctly and efficiently is threatened by free-riders—users who consume the service while providing little or no service to others. Within P2P anonymity systems, two aspects of user cooperation are of major concern. First, it is essential that peers be *compliant* with the protocol specification. For example, peers must reliably forward each others' traffic while joined to the system. Second, it is important that peers be *available* to the system. For example, peers should remain joined for extended periods of time.

Both non-compliance and low availability have the potential to undermine P2P anonymity systems. A significant presence of non-compliant peers will result in a high incidence of failed routes and path reconstructions. Low availability will induce a high rate of churn associated with peers frequently joining and leaving the system. Besides degrading system performance, both consequences facilitate certain types of attacks on anonymity [3, 32]. In addition, low availability due to free-riding implies that the expected number of peers joined to the system at any point in time will be smaller. Since many measures of anonymity improve monotonically with the average number of peers in the system [22, 23, 31], free-riding directly impacts the fundamental quality of the anonymity service.

The free-riding problem in P2P systems can generally be addressed by introducing explicit incentive mechanisms for aligning users' self interest with global system objectives. However, applying general techniques to a particular system typically requires a careful analysis of application-specific characteristics, which, in the case of anonymity systems, are particularly challenging. While incentive solutions to mitigate non-compliant behavior in anonymity systems have been studied previously [12, 14], we are not aware of any prior work addressing the problem of low peer availability in such systems. Yet, low availability is likely to be pervasive, as it can easily be achieved by fully compliant peers simply leaving the system once their immediate needs have been fulfilled. Because of the potential seriousness of this

[1] Recent efforts to deploy a high quality commercial anonymity system based on the mix network architecture in [9] have failed, mainly due to high operational costs coupled with low user subscription rates [20].

form of free-riding and because compliance-enforcing incentives are insufficient to prevent it, we focus on incentive mechanisms that can promote high availability.

We propose the use of a payment-based incentive mechanism, where peers requesting service must financially compensate the peers that contribute to servicing the request. This mechanism attaches a financial cost to using the system. The intuition is that in order to minimize costs, users will change their behavior when using the system. For example, the financial cost of using the system can be partially or totally recouped by remaining joined to the system and providing service to others. However, moving from a free system to a paid system can have several consequences on users' behavior – not all desirable from the perspective of the system designer. It is therefore important to evaluate the impact that a payment mechanism will have on the system. Another orthogonal issue is the design of efficient payment mechanisms that can be implemented and coupled to P2P anonymity system. In this paper, we make two contributions in these directions:

- We develop and evaluate a game-theoretic model to provide insights on peer availability when the P2P anonymity system is augmented with a payment mechanism. Our analytical results show that peer availability can significantly increase under certain conditions. Our findings further suggest that under a paid system, all but the most demanding users can fully recover the upfront payments when using the service.

- We present the design of an anonymity-preserving payment mechanism that can be used to promote high peer availability. Our scheme borrows from work in anonymous digital cash and micropayments to embed small payments for service in the messages exchanged in path-based P2P anonymity systems. Our scheme has several desirable security and anonymity properties, and can easily be coupled with existing P2P anonymity systems.

2 Background on P2P Anonymity Systems

In this paper, we will consider interactive, path-based, P2P anonymity systems. We assume users can freely join and leave the system at any point in time. Users can send anonymous messages as long as they are joined to the system. The term initiator will be used to denote the peer originating anonymous messages. In order to send anonymous messages, the initiator must first construct a path through a sequence of peers participating in the P2P system that terminates at the intended destination. This is known as the *anonymous path*. The peers that comprise any given anonymous path are randomly chosen from the set of peers joined to the system. Note that a peer on the anonymous path and the destination does not know if their immediate predecessor is the initiator or simply another peer on the path. An anonymous path constructed by the initiator can be used to send several messages to the same destination, but must be eventually destroyed[2]. Figure 1 illustrates the P2P anonymity system under consideration.

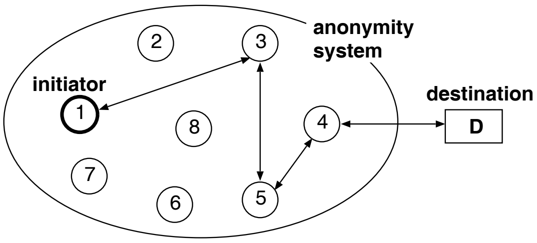

Figure 1. Anonymity system with node 1 communicating with destination D

We assume that the identity of the peers along an anonymous path is fully known to the initiator at the time a message is sent, which is a necessary condition for the applicability of the payment mechanism we propose. Anonymity systems based on Chaumian mixes [9], such as Tarzan [19], MorphMix [27] and GAP [5], satisfy this requirement.[3] However, our incentive mechanism would not trivially apply to a system like Crowds [26], where the path is not known to the initiator.

3 Motivation for Payment-Based Incentives

In order to mitigate the impact of free-riders, researchers have recently suggested building explicit incentive mechanisms into various types of P2P applications. The most widely studied mechanisms fall into two categories:

- Reputation systems in which peers punish or reward each other based on observed behavior [7, 14, 28].

- Payment-based mechanisms that require tokens or money to be exchanged in return for service [2, 4, 8, 33].

Reputation mechanisms require the capability (a) to evaluate a peer's cooperativeness on the basis of observable behavior and (b) to provide a punishment or reward with the desired incentive. In highly scalable anonymity systems (e.g. MorpMix [27]) low availability is largely unobservable due to decentralized group membership and path

[2] There are several reasons why anonymous paths must be periodically reconstructed through the system, such as coping with changes in group membership and defending against attacks on anonymity [13, 27, 32].

[3] Note that the path construction mechanism of MorphMix, while distributed in nature, ultimately delivers to the initiator the identities of the peers along the path prior to the sending of a message.

construction functions. Adding observability would require substantial redesign, impose significant communication overhead, and weaken the anonymity system's robustness against so-called intersection attacks, which work by tracking changes in group membership. Moreover, even if availability was completely observable, the most obvious punishment mechanism—refusing or granting low priority to service requests—is not feasible in systems where service is requested anonymously.[4] Other punishments are feasible, but possibly less effective. For example, refusing to include peers with poor reputations in anonymous paths has the effect of denying cover traffic to those peers. However, such punishment may have limited influence on peers already inclined toward low availability, who are likely to have weak anonymity requirements to begin with.

Because of the aforementioned concerns about reputation mechanisms, we have chosen to explore payment-based incentives as an alternative. In a payment-based mechanism for anonymous systems, the initiator would compensate each peer on the path for forwarding its message to the next hop. Peers would accumulate payments at a rate proportional to their availability. Such payments would attach a cost to using the system, but this cost could be partially or totally recouped by remaining joined to the system and providing service to others.

4 Modeling the Impact of Payment-based Incentives on Availability

In this section we conduct a qualitative analysis of payment-based incentives and investigate its impact on peer availability. We will assume that all peers participating in the P2P system are self-interested, meaning they all have consistent objectives (e.g., utility function), and will respond rationally to well-defined incentives, such as financial incentives. Users are solely interested in minimizing their costs when using the system. Since the costs imposed on a user depend on decisions made by other users, this model naturally falls into the game-theoretic framework. We adopt the concept of Nash equilibria points (NEPs) to characterize the equilibrium of the system[5]. In what follows, we demonstrate the relationship between system parameters and NEPs, which provide several insights into the use of payment-based incentive mechanisms.

Consider a general path-based anonymity system. In the presence of payments, the system requires the initiator to pay all peers along its anonymous path for each message it generates. Let us assume that the price for forwarding a message, q, is the same for all users and determined by the system designer, and that all paths through the system are of a fixed length L. Thus, the cost to send a message is $L\,q$. Let N denote the total number of potential users of the system. In order to send a message anonymously, a peer must join the system for a minimum amount of time. Let s denote the time required for the system to deliver a single message anonymously. This includes the time required to transmit the request and its reply, and any other protocol overheads. We assume that s is much smaller than 1 ($s \ll 1$).

Each individual user i in our model is characterized by three variables:

- **Demand (l_i):** The demand of a user for the anonymity system represents the number of messages generated by the user per unit of time. The demand cannot be arbitrarily high. In particular, if anonymous messages are serialized then the following must hold: $s\,l_i \leq 1$.

- **Level of availability (c_i):** The level of availability represents the fraction of time the user is joined to the system. The level of availability is bounded from below by the demand of the user, as it must be joined to the system in order for the system to deliver its own anonymous traffic. In particular, we have $s\,l_i \leq c_i \leq 1$.

- **Payment from external funds (p_i):** This represents the average unrecovered cost per message. We interpret this value as an amount of money payed from external funds and injected into the system for each message sent by user i. The price for sending a message anonymously is $L\,q$, and thus, the amount paid from external funds is bounded from above, such that $0 \leq p_i \leq L\,q$.

The strategy space of a user corresponds to the space defined by the variables above and a strategy corresponds to a choice of values for these variables.

Recall that peers pay to send their messages anonymously through the system. The rate at which user j injects money into the system is $L\,q\,l_j$. This monetary influx comes either entirely from external funds (p_j), from money accumulated by forwarding other peers' messages, or from a combination of the two. Since other peers receive payments when forwarding messages, the aggregate influx of money generated by all initiators is dispersed among the peer population. Let λ_i be the rate at which user i receives revenue from all initiators in the system. In order to determine λ_i, we introduce the event $E_{i,j}$ which indicates that peer i has been selected to be on the path constructed by peer j. Of course, peer i can only be selected to be on j's path if it is joined to the system when the path is constructed. If we assume that peers are equally likely to be chosen to occupy a

[4] An equivalent argument holds for mechanisms that reward good reputations.

[5] A Nash equilibrium point (NEP) is defined by a set of decisions, one for each peer, where no individual peer can reduce its costs by unilaterally changing their decision.

position in the path constructed (e.g., Tarzan [19]), we have

$$\Pr[E_{i,j}] = 1 - \left(1 - \frac{1}{1 + \sum_{k=1,k\neq i}^{N} c_k}\right)^L \quad , \quad i \neq j \quad (1)$$

Note that this probability does not depend on peer j, the node constructing the path, as we have assumed that the path construction mechanism is identical for all peers. However, (1) depends on the number of peers joined to the system, which on average, assuming that peer i is joined to the system, is given by $1 + \sum_{k=1,k\neq i}^{N} c_k$.

We can now determine λ_i by considering all initiators in the system. In particular, we have

$$\begin{aligned}\lambda_i &= \sum_{j=1,j\neq i}^{N} l_j\, q \Pr[E_{i,j}] \\ &= q\left(1 - \left(1 - \frac{1}{1 + \sum_{k=1,k\neq i}^{N} c_k}\right)^L\right) \sum_{j=1,j\neq i}^{N} l_j \quad (2)\end{aligned}$$

The above equation assumes peer i is joined to the system. Since peer i only receives revenue while joined to the system, the long term rate at which it accumulates revenue is $\lambda_i\, c_i$.

There are several factors that determine the costs and the rewards of a user that participates in an anonymity system. A precisely defined utility function for such a user is usually subjective and can be rather complex (for example, see [1]). For our present purposes, we introduce a relatively simple utility function that considers users to be subject to the following costs and rewards:

1. A cost per unit of time for being joined to the system. There are many reasons why users may suffer such costs, ranging from the commitment of local resources for forwarding other users' traffic, to the increased risk of scrutiny incurred by participating in an anonymity preserving system.
2. The cost of using external funds to send messages anonymously. Since users naturally value money, using external funds represents a clear cost.
3. The reward of sending messages anonymously. Clearly, users of an anonymity system see value in sending their traffic anonymously.

Assuming that costs and benefits scale linearly in the user's decision variables, a user's utility function, which in this case represents the total cost, is given by

$$u_i(l_i, c_i, p_i) = (l_i\, p_i - \lambda_i\, c_i) + \alpha_i\, c_i - \beta_i\, l_i \qquad (3)$$

The first term in (3) represents the financial cost of sending messages through the system. Note that this cost is the difference between the rate of money used from external funds and the rate of money accumulated for providing forwarding service. The second term in (3) represents the cost per unit of time associated with being joined to the system. The third term represents the reward obtained from sending anonymous messages at a particular rate. In a diverse population, users are likely to have different sensitivities to the different costs and rewards. To capture this heterogeneity among users, we assign user-dependent weights α_i to the costs of being joined and β_i to the rewards of sending messages anonymously.

Given the above utility function and the constraints on the user's decision variable, we can now introduce the optimization problem each user i will solve. In particular, we have:

$$\begin{aligned} \min_{l_i, c_i, p_i} \quad & (l_i\, p_i - \lambda_i\, c_i) + \alpha_i\, c_i - \beta_i\, l_i \qquad (4) \\ \text{subject to} \quad & 0 \leq l_i \leq 1/s \\ & s\, l_i \leq c_i \leq 1 \\ & 0 \leq p_i \leq L\, q \\ & l_i\, p_i + \lambda_i\, c_i \geq L\, q\, l_i \end{aligned}$$

The last constraint in the problem above comes from the fact that each user must have sufficient funds to cover the cost of its demand.

It is important to note that both the objective function and constraints in the optimization problem above depend on the decision variables of all other users. In particular, λ_i depends on the level of availability of all users in the system. Moreover, both the objective function and the fourth constraint are non-linear in the decision variables of user i.

To make our model analytically tractable, we group users into a small number of classes, M, where users within a class are identical. Note that the subscript i will now denote a class and not a user. For convenience, we also assume that each class i contains N users. Since users within a class have identical costs and rewards, we are interested only in equilibria points where the decisions of users within a class are also identical. Therefore, (l_i, c_i, p_i) will denote the choice of variables for all users in class i, with $i = 1, \ldots, M$. We will abuse terminology and continue refer to user i, when actually meaning a representative user of class i.

Note that λ_i now denotes the rate at which a single user in class i receives revenue from all initiators while joined to the system. Since all other users in each class have the same value for their decision variables, we can simplify equation (2). Moreover, we assume that the number of users in each class is very large. This simplifies the problem since λ_i will no longer be class dependent. Hence, all users when joined to the system, receive revenue from initiators at the same rate. The limit for λ_i is given by

$$\lim_{N \to \infty} \lambda_i \quad = \quad \frac{L\, q \sum_{j=1}^{M} l_j}{\sum_{j=1}^{M} c_j} \tag{5}$$

Note that λ_i still depends on the aggregate decisions of users from each class in the system. However, since each class has a very large number of users, the impact of the choices of a single user on λ_i is negligible.

4.1 Fixed Demand, Variable Payment Model

We start by evaluating a model where the user's demand for the anonymity system is fixed. That is, each user in class $i, 1 \leq i \leq M$, generates messages at a fixed rate l_i that must be delivered anonymously through the system. Note that l_i is now a parameter of the model instead of a decision variable. A user in class i is left with two decision variables, namely, c_i and p_i.

Under this model, each user $i, 1 \leq i \leq M$ solves the following optimization problem:

$$\begin{array}{rl} \min\limits_{c_i, p_i} & (l_i\, p_i - \lambda_i\, c_i) + \alpha_i\, c_i - \beta_i\, l_i \\ \text{subject to} & 0 \leq l_i \leq 1/s \\ & s\, l_i \leq c_i \leq 1 \\ & 0 \leq p_i \leq L\, q \\ & l_i\, p_i + \lambda_i\, c_i \geq L\, q\, l_i \end{array} \tag{6}$$

Note that the above problem is similar to the original formulation presented in problem (4), but with i referring to a representative user from class i and not all individual users, and with l_i now being an input parameter. This two assumptions greatly simplify the original problem, as the utility function and the constraints are now linear in the representative user's decision variables.[6]

It can be shown that the system defined by problem (6) always has at least one NEP (the problem satisfies the conditions required by Debreu's NEP existence theorem [11]). In what follows we characterize these equilibria points analytically. Note that the last term in the cost function in problem (6) is a constant term (as a consequence of fixing l_i) and will be ignored in the analysis that follows.

Consider the system when $M = 1$, in which case the system has a completely homogeneous user population (we will drop the indices as we now refer to a single user class). Note that in a NEP of system with a single class, no user should benefit from deviating from the choice of decision variables made by all users in the class. Our system is fully characterized by the parameters L, q, l and α. We can establish the NEPs for this system as follows.

Theorem 1 *In the system defined by L, q, l and α with a single class of users ($M = 1$), the Nash equilibria points are determined as follows (proof available in [17]):*

$$\begin{aligned} p^* &= 0 \\ c^* &= \begin{cases} 1 & \textit{if } l\, Lq/\alpha \geq 1 \\ s\, l & \textit{if } l\, Lq/\alpha \leq s\, l \\ \left[l\, \frac{Lq}{\alpha},\ \max(1,\ 2\, l\, \frac{Lq}{\alpha}) \right] & \textit{otherwise} \end{cases} \end{aligned} \tag{7}$$

For certain system parameters, the theorem establishes a range of possible Nash equilibria points (third condition of (7)). However, there is a strict ordering of these equilibria from the perspective of the users' cost function. In particular, among the range of NEPs, the point $c^* = l\, Lq/\alpha$ yields the lowest cost to all users. In game theory, this equilibrium point is known as a Pareto optimal Nash equilibrium. Since this NEP is preferred by all users, it is reasonable to consider it as the equilibrium point of the system. Note, however, that any NEP within this range yields higher level of availability than the minimum (with the preferred NEP yielding the lowest availability within the range of NEPs).

There are several interesting observations that can be drawn from Theorem 1. First, we note that if the system does not require users to pay to send messages anonymously ($q = 0$), the availability at equilibrium is given by $c^* = s\ l$, which is the minimum fraction of time that users must join the system to deliver their messages. In a paid system ($q \neq 0$), it is possible to have an equilibrium point where peer availability is much higher than the required minimum. Second, in all possible equilibria, no user is required to consistently introduce money from external funds ($p^* = 0$). Thus, the anonymity service is actually free of charge over sufficiently long time scales, as up-front payments to send messages are fully recouped[7]. Third, the level of user availability scales linearly with the users' demand (l). In particular if $l > \alpha/(Lq)$, then users are always joined to the system ($c^* = 1$). Fourth, availability scales inversely with the users' sensitivity to being joined to the system (α). If users have a low cost for being joined, the availability can be very high, while if users are very sensitive to being joined, availability is minimal. As a last observation, note that availability scales linearly with the price for forwarding a message (q). Thus, the system designer could arbitrarily increase the price to forward a message to force users to a NEP with high availability. Although the system would still remain free of charge ($p^* = 0$), users might need to make a much larger up-front payment. This characteristic is mainly due to our assumption that users must send all messages that they generate. In the next section we will explore a different model that allows users to moderate their demand.

[6] Recall that the original formulation, as stated in (4), is non-linear and consequently, much harder to solve analytically.

[7] The model presented ignores possible transaction fees imposed by the bank, which may introduce some small cost for the service.

The above result can be extended to the case of a heterogeneous user population ($M = 2$), yielding analytical results that reflect the scaling of peer availability with different system parameters. In the heterogeneous case, however, we find that the anonymity service is no longer free of charge to all users. In particular, users with high demand will pay a positive amount from external funds ($p^* > 0$), whereas those with low demand will still pay nothing. Space considerations prevent us from presenting the details for the heterogeneous case here; more information can be found in [17].

4.2 Fixed Payment, Variable Demand Model

We now consider a model where the amount of money each user introduces from external funds is fixed. That is, from the perspective of the general model introduced earlier, the quantities p_i are now parameters and not decision variables. However, the demand of each user for the anonymity system, l_i, will now be treated as a decision variable. Thus, each peer can decide the rate at which it will generate messages that have to be sent anonymous through the P2P system.

Under this variation, each user of class $i, 1 \leq i \leq M$ solves the following optimization problem:

$$\begin{aligned} \min_{l_i, c_i} \quad & (l_i\, p_i - \lambda_i\, c_i) + \alpha_i\, c_i - \beta_i\, l_i \qquad (8) \\ \text{subject to} \quad & 0 \leq l_i \leq 1/s \\ & s\, l_i \leq c_i \leq 1 \\ & 0 \leq p_i \leq L\, q \\ & l_i\, p_i + \lambda_i\, c_i \geq L\, q\, l_i \end{aligned}$$

The above problem is similar to problem (6), but with l_i being a decision variable and p_i a fixed parameter. Again, it can be shown that this problem always has at least one NEP (the problem satisfies the conditions required by Debreu's NEP existence theorem [11]).

Consider again the case of a homogeneous user population, $M = 1$. As before, all users within this class have identical preferences and, as before, we are interested in characterizing only symmetric NEP. Problem (8) is fully described by the parameters p, L, q, α and β (we again drop the indices as we are considering a single user class). The NEP of this problem is characterized as follows.

Theorem 2 *In the system defined by p, L, q, α and β, with a single class of users ($M = 1$), the Nash equilibria points are determined as follows (proof available in [17]):*

$$\begin{aligned} l^* &= \begin{cases} 0 & \text{if } \beta - Lq < \alpha\, s \\ [0,\, 1/s] & \text{if } \beta - Lq = \alpha\, s \\ (0,\, 1/s] & \text{if } \beta - Lq > \alpha\, s \end{cases} \qquad (9) \\ c^* &= s\, l^* \end{aligned}$$

The result in the above theorem establishes a range of possible NEPs under the condition $\beta - Lq \geq \alpha\, s$. However, once again there is a strict ordering of these NEPs from the perspective of the users' cost function. The Pareto optimal Nash equilibrium in this case is given by $l^* = 1/s$, $c^* = 1$, that is, users set their demands to the capacity of the system and remain joined all the time. Since we have not imposed any user-specific upper limit on the demand, users will choose to use the system as much as possible. These results are readily generalized to the case where users have a maximum demand for the system.

We can obtain several insights from the Theorem 2. First, if users place a low value on sending messages anonymously (low β) or a high cost to being joined to the system (high α), then not using the system ($l^* = c^* = 0$) is the only equilibrium point. Despite this negative result, we expect that privacy-concerned users will place a high value on sending their traffic anonymously, such that, in general, $\beta > \alpha$. Second, under some conditions, there are no Nash equilibria where users have zero demand for the system ($l^* = 0$). In this case, users always prefer to generate some demand for the system, as opposed to not using it. Third, the results show the dependence of the price to forward a message, q, which is stipulated by the system designer, on the user's decisions. Note that by arbitrarily increasing q the system designer can force users to an equilibrium where they have zero demand for the system.

The essential point in the analysis of the two models described above is to establish conditions under which a payment mechanism provides users a clear incentive to increase their level of availability. It's important to note that under some conditions, a payment mechanism may not improve peer availability. Fortunately, this is not always the case. In fact, our results show that it's possible to boost availability and at the same time have an anonymity system that is essentially free of charge, despite requiring peers to pay to send messages anonymously. This is possible since users can recoup all expenses by providing service to others.

The above analysis focuses on characterizing the different equilibria of the system and understanding the impact of the proposed incentive mechanism at the equilibria points. Another important consideration is understanding the dynamics by which users may reach an equilibrium point. Convergence to a NEP is in general a hard problem and while it is an important consideration, is beyond the scope of this paper. We leave such investigation for future work.

5 Micropayment Scheme

We propose an off-line, debit-based anonymous micropayment scheme that allows peers to make payments to each other in exchange for message forwarding service. Our scheme is a composition and adaptation of two schemes published in the literature. Since anonymity is an application requirement, we must adopt a payment scheme, where payments cannot be linked to the identity of any specific payer. Such features are readily available in the anonymous digital cash scheme proposed in [10]. At the same time, since we envision a system where peers pay per message forwarded, an efficient payment scheme is also required. Such efficiency requirements can be met by using micropayment schemes, such as Payword [30], which typically achieve simplicity at the expense of anonymity. Our scheme will combine features such as blind signatures and split identity information, together with hash chains and verifiable certificates to deliver an efficient anonymous payment scheme.

As with any other payment mechanism, our scheme uses a *bank*, a trusted entity[8] that will keep account information (e.g., cash balance) of all peers in the system. Our payment scheme interacts with the bank in an off-line fashion, meaning that the bank is not contacted at the time of a transaction.

In short, the payment scheme works as follows: Suppose a peer, P, wants to use the anonymity service provided by the P2P anonymity system. Before P can request any peer to forward its messages, it must purchase a *certificate* for a given amount from the bank. This certificate is signed blindly by the bank and has no explicit binding information to the identity of P. Peer P then selects a peer, say Q, that will serve as an intermediary peer along the anonymous path. Peer P then binds the certificate to Q and sends the certificate anonymously to Q. The certificate is equivalent to a payment in advance, however, peer Q can only cash in this certificate together with valid *tokens* that will be provided by P. Peer Q checks the validity of the certificate and its binding. To pay Q for its forwarding service, P includes a token in each anonymous message it sends through Q. Peer Q verifies the validity of the token received against the certificate and stores it. At a later point in time, peer Q presents the certificate as well as the last token received to the bank to redeem its payments. The bank verifies the validity of the certificate and the tokens and credits the proper amount into Q's account. Note that peer P must give a certificate to each peer along its anonymous path.

From the brief description above, we observe that the payment scheme can be decomposed into three phases: *(a)* peer P–bank interaction (purchasing certificates); *(b)* peer P–peer Q interaction (making payments); *(c)* peer Q–bank interaction (redeeming payments). We now provide the details of each of these interactions:

5.1 Purchasing certificates

Suppose initiator P wants to use the anonymity service. P must first purchase *certificates* for specific amounts from bank. To do so, it prepares a certificate that contains a globally unique identification number (chosen randomly by P from a very large set), its monetary value v, its committed split identity information[9] and the final value of a hash chain of length v[10]. The bank receives and verifies the legitimacy of the certificate. If P has enough funds in its account, the bank withdraws the corresponding amount from P's account, signs the certificate, and sends the signed certificate back to P. The bank must sign the certificate *blindly*, such that it cannot link the certificate with the initiator that purchased the certificate (recall that we do not trust the bank to safeguard the identity of peers). This means that the bank will not see the certificate it signs in plain text, although it almost surely knows what are the terms (i.e., the certificate value) being signed. An existing mechanism for obtaining blind signatures, such as cut-and-choose [10], can be readily used. This message exchange for certificate purchase is illustrated in Figure 2(a).

5.2 Making payments

The first-time peer P selects a peer, say peer Q, to be part of its anonymous path, it binds and sends Q one of its certificates. To bind a certificate to Q, peer P uses some unique information about Q (e.g., its IP address) to determine the split identity information that will be revealed to Q. A message containing the certificate and the appropriate half of the split identity information is sent to Q through the anonymous path itself, and not directly, so that Q does not learn the identity of P. Peer Q then verifies the validity of the certificate through the bank's signature and the correctness of the split identity information revealed by P. Peer Q now agrees to forward packets on behalf of P. Each message that P sends through Q will include a token (the next value in the hash chain) together with the certificate identification number. The token is the payment for forwarding this message. Q verifies that the payment is valid by applying the well-known hash function to this token and comparing it with the previously received token (or the token

[8] The bank is trusted as a financial institution but not necessarily trusted to safeguard payers' identities.

[9] Identity splitting is a technique based on secret sharing used in offline cryptographic payment protocols to discourage *double spending*. With each transaction, the payer verifiably reveals part of its split identity in response to a challenge. With high probability, the payer's true identity can be reconstructed if the same unit of currency is used in more than one transaction [6, 16].

[10] The final value of a hash chain of length v is the result of applying a unique hash function h, recursively, v times to some random value r, which is denoted the root of the chain. Thus, the i-th value of the hash chain is given by $h^{(i)}(r) = h(h(\cdots h(r)\cdots))$, where function h is applied recursively i times.

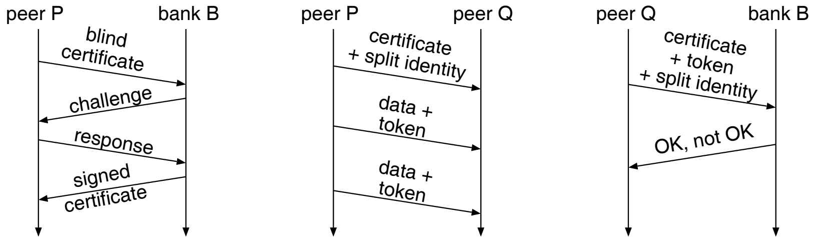

(a) (b) (c)

Figure 2. Messages exchanged during each interaction of the payment scheme: (a) purchasing a certificate; (b) making payments; (c) redeeming payments.

in the certificate)[11]. Since the certificate ensures the payment of at most v tokens, peer Q should forward at most v messages for P. Of course, P can send Q a new certificate when its balance reaches zero if it intends to continue to use Q to forward its messages. Figure 2(b) shows the messages exchanged between P and Q. For clarity, this figure suppresses the details of the anonymous channel connecting P and Q.

The micropayments described above can be readily integrated with path-based anonymous protocol in which the initiator has complete knowledge of the peers along the path and uses recursive encryption when sending messages to communicate secretly with each peer. A recursively encrypted message P_1, traversing a path with L hops, has the following form:

$$\begin{aligned} P_L &= \{D, M, C_L\}_{K_L^+} && (10) \\ P_i &= \{S_{i+1}, P_{i+1}, C_i\}_{K_i^+} \quad \text{for } L-1 \geq i \geq 1 && (11) \end{aligned}$$

where S_i is the address of the i-th peer in the path, M is the anonymous message destined to D, C_i is payment data sent by the initiator which is destined to peer i (e.g., a certificate or token), and $\{X\}_{K_i^+}$ denotes message X encrypted with public key K_i^+. This message, in virtue of its layered structure, is commonly referred to as an *onion* [25].

The initiator will forward the onion to the first peer of the path, S_1. Each intermediary peer i in the path will have access to payload P_i after decrypting the message with its private key. The payload contains the address of the next hop S_{i+1}, an encrypted payload to be sent to that hop, and payment information destined to peer i, C_i. When the initiator first decides to form a path using peer i, the contents of C_i are simply a certificate that has been bound to i and the corresponding split identity information. Peer i can verify the legitimacy of the certificate and split identity information before it forwards any messages on behalf of the initiator. In subsequent messages sent by the initiator, C_i will contain payments to peer i in the form of tokens. However, to prevent intermediate peers from receiving payments without forwarding messages, an acknowledgment mechanism is used, such that a payment can only be effectively obtained *after* the payload has been properly forwarded to the next hop. To accomplish this, the token for peer i is encrypted with a random symmetric key generated by the initiator. This random key is only visible by the successor of peer i on the path. Upon receiving a message, each peer returns the symmetric key to its predecessor in an acknowledgment message, which enables the previous hop to obtain its payment. On receiving an acknowledgment from its successor, peer i decrypts its token. Thus, the contents of C_i in this case are

$$C_i = \{\{t_i^m\}_{K_i^m}, K_{i-1}^m\} \tag{12}$$

where t_i^m is the m-th token sent to peer i, K_{i-1}^m is the symmetric key associated with the m-th token that should be returned to peer $i-1$, and the notation $\{X\}_K$ denotes a message X encrypted with symmetric key K. Figure 3 illustrates the operation of the protocol with payments.

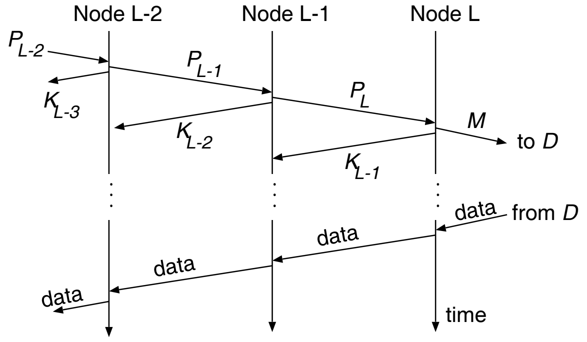

Figure 3. Message exchange for the last three hops of an anonymous path.

Eventually, the message reaches the last peer in the path

[11] The first token sent is $h^{(v-1)}(r)$, which can be verified as legitimate by applying the hash function to it once and comparing the result with the value imprinted in the certificate.

(L), which then forwards message M to destination D. Since D is assumed to be outside the anonymous system, we cannot expect D to perform any additional functionality that is pertinent to the anonymous protocol. Thus, the token sent to the last hop is not encrypted by the initiator. Finally, intermediate peers must forward any response from D along the reverse path towards the initiator.

There are two important considerations in the above protocol that are not necessarily solved by the current incentive structure. First, peers along the intermediate path must forward acknowledgment messages to their predecessors. Second, the last peer in the path must forward the message to the final destination. It is possible that the financial incentives given to peers (i.e., token values) will be sufficient to motivate them to comply with the protocol. If this is not the case, then failing to forward such messages can be treated as non-compliant behavior and handled by some other mechanism. In Section 6 we briefly discuss how non-compliant behavior can be addressed in P2P anonymity systems.

5.3 Redeeming payments

Peer Q will generally contact the bank to redeem its payments once the total number of tokens guaranteed by a certificate has been received. Peer Q sends the bank the certificate, the split identity information revealed by peer P and the last token received for that certificate. The bank verifies if the certificate is valid by checking its signature and also confirms that the certificate is bound to Q by checking the split identity information provided by Q. It then verifies if the token presented is legitimate by recursively applying the known hash function to the token until the final hash value printed in the certificate is obtained. The bank then credits an amount of v into Q's account.

5.4 Discussion

The payment scheme described above has several desirable security and anonymity properties. For example, it provides protection against *double-payment* and *double-spending*. The former can occur when a peer attempts to redeem the same certificate multiple times at the bank. The latter can occur when an initiator exploits the off-line nature of the payment mechanism to spend the same certificate with more than one peer. Among other security properties, the scheme also prevents a peer from reusing a legitimate certificate that it has received. Among desirable anonymity properties, the scheme preserves the identity of the initiator even when the bank and some peers collude. Many of these properties follow directly from the security of previous payment schemes [10, 30] and the anonymity provided by path-based anonymous protocols [19, 27].

As specified above, the payment scheme allows the initiator to pay intermediate peers along the forward path — that is, the path taken by the message as it travels to the destination. However, we have not provided payments for response messages that originate at the destination and follow the reverse path to the initiator. Although we could assume that payments in the forward path provide sufficient incentives for intermediary peers to relay back response messages, it is also possible that explicit payments will be required for forwarding these responses. Recall that we target our payment scheme to anonymous protocols that use symmetric paths, such that forward and reverse paths are formed by the same set of peers. There are several ways in which payment for response messages can be incorporated. For example, the initiator could provide payments for reverse messages in a separate recursively encrypted message after it has received the response, which can perhaps be piggybacked in a subsequent requests. This method has the advantage of allowing the initiator to pay in proportion to the size of the response.

Finally, the properties required by the payment scheme presented above (initiator knows that path and recursive encryption) form the basis for Chaumian mixes [9] and are present in many P2P anonymity systems that have been proposed in the literature (e.g., [5, 19, 27]). These systems are particularly well suited for integration with a payment mechanism, as the initiator can embed certificates and payments for each peer on the path without the risk of being stolen by an eavesdropper.

6 Attacks and Unintended Consequences

The previous sections have shown that payments can provide an incentive for high availability and can be implemented anonymously. We now consider whether this additional mechanism might affect the overall security of the system, possibly in subtle and unintended ways. A number of threat models must be considered, including attacks designed against both anonymity and the payment mechanism itself. In this section, we briefly enumerate and discuss some possible attacks and undesired consequences.

Attackers can deny service to others by failing to comply with the forwarding protocol by, for example, accepting certificates but refusing to forward packets.

This type of denial of service attack is possible in any anonymity system; the addition of an incentive does not eliminate it. However, our payment-based incentive introduces the possibility of financial loss under such an attack. Because our payment scheme does not allow the initiator to reclaim or reuse certificates that have been given to intermediary peers, it is possible for the initiator to lose money if peers refuse to forward messages in exchange for tokens. In this case, the initiator has no means of recovering its certificate. We emphasize that the peer holding a certificate cannot redeem it at the bank without the appropriate tokens, so there is no financial incentive for denying service. Nevertheless, malicious peers can mount such attack on the system. We discuss how this attack can be mitigated in the next consideration.

Peers can fail to comply with the protocol if the offered financial incentives are not sufficient or if they are unresponsive to such incentives.

It is possible that the financial incentives given to peers (i.e., token values) will be insufficient to motivate them to comply with the protocol. Failing to comply generates broken paths, which ultimately forces the initiator to create a different anonymous path. Over time, such behavior eventually reduces the flow of tokens that a non-compliant peer receives. Thus, there is an implicit incentive to comply with the protocol and then leave the system once the costs of using the system have been recouped.

The incentive to comply can be rendered explicit through the use of an additional reputation mechanism that can identify and isolate non-compliant peers, thereby reducing their impact on the system. Initiators can use this mechanism to minimize the risk of losing tokens to broken paths. We have designed and evaluated an anonymity-preserving reputation mechanism and show that under certain conditions, it can effectively isolate non-compliant peers. Space limitations prevent us from presenting the mechanism here, but readers can find more details in a related technical report [17].

The Bank can use its privileged position to conduct timing attacks, possibly aided by colluding peers.

The bank can attempt to mount a straightforward timing analysis attack on the system by correlating the signing of certificates requested by peer P to the time when a certificate is later redeemed. The bank can also collude with other participating peers to correlate purchase and usage of certificates. For example, if P purchases certificates for a given face value v and then immediately uses them, the bank and colluding peers can correlate the purchase of the certificate with its immediate usage (or redemption at the bank) in order to infer P's identity.

To counteract such traffic analysis, peers should avoid having the bank issue certificates on a per-session timescale, but instead should purchase many certificates at once divided into a few well-defined face values. Ideally, the bank should be a publicly accessible authority providing digital cash services for various other businesses other than the anonymity system. In this case, correlating payments with purchases is potentially more difficult, as transactions from peers in the anonymous system will be interleaved with a large number of unrelated transactions.

Attackers wishing to analyze traffic now have an additional monetary incentive to operate additional peers under pseudonyms, lending additional power to passive traffic analysis attacks.

While passive traffic analysis attacks are a threat in all anonymous systems, the addition of payments do not necessarily strengthen their effect. If one assumes that an attacker serious about breaking anonymity will deploy all resources available at hand (e.g., multiple machines, multiple identities, etc.) regardless of any additional incentive, then adding payments does not expose the system to a new type of threat. The real question is whether the income received from providing service can be converted into additional resources for an even more powerful attack. Note that, honest users can also operate multiple peers, not with the intent of attacking the system but to obtain financial benefits. If the net effect is an increase in the overall number of peers in the system, there is no obvious advantage to the attacker.

New attacks may emerge with the goal of stealing money from the system.

While it might be possible to exploit weaknesses in, say, path construction protocols to extract undeserved currency, by far the easiest way to obtain money is to become a compliant, highly available peer. An interesting question, however, is whether users who perform work for payment but never use the service should be regarded as a threat. Although such users do not provide cover traffic to the system, their presence does increase the overall size of the group and system stability. We are therefore inclined to view this profiteering behavior as benign at moderate levels. If, however, profiteering were to become a dominant behavior, the ability of ordinary users to recover costs would be diminished and the relative lack of cover traffic in the system could reduce the quality of anonymity, thereby driving out paying users. One might expect an equilibrium level of profiteering to emerge in this case, an intriguing question that we leave to future research.

Payments for service, if not recoverable, will represent a net financial cost for using the system. This cost could depress demand for anonymity, which might negatively affect the quality of anonymity.

Although our model indicates that the cost of payments can be recouped through increased availability, one must interpret this result conservatively. In practice, overheads such as bank fees, lost tokens, and profiteering might make full recovery of costs difficult, can induce users to abandon using the service. Under a shrinking total user population, the average group size could become small even as remaining users increased their availability.

The likelihood of this scenario depends on the elasticity of demand for anonymous communication, which, unfortunately, is not well understood. While it is widely claimed that most people will not pay for anonymity, it is also the case that most people do not use even free anonymous services. It may be that non-financial costs such as the reduced performance and inconvenience of anonymous communication create a self-selection effect that keeps out peers for whom anonymity is of marginal value. In this case, demand should be relatively inelastic and payments could help to increase the number of willing users actively joined at any point in time. Further empirical study is needed to decide this question.

7 Related Work

Payment-based incentive mechanisms have been proposed for various specific P2P systems with the intent of promoting compliance with the system protocol. MojoNation was a deployed P2P system for robust file storage and retrieval in which peers traded a form of private currency called *mojo* in exchange for both the storage and retrieval of data.[12] A number of payment-based schemes have been proposed for ad hoc networks and wireless multi-hop networks. In [8] the authors use payments to drive the system to an equilibrium point where nodes comply with the protocol. Their approach requires all participating peers to have access to tamper-proof hardware to enforce honest exchange of payments. In [33], a payment scheme called *Sprite* is proposed. Sprite uses a centralized record keeping authority along with a cryptographic scheme for deferred payments, and does not require specialized hardware. Another approach for safely establishing payments in multi-hop cellular networks is described in [4]. All of these payment mechanisms do require the knowledge of the identity of the payer and thus cannot be used in an anonymity system.

All the proposed incentive mechanisms for P2P systems have so far focused solely on the problem of non-compliant peers. Although the problem of low availability has been raised in the literature [24, 29], we are not aware of any incentive mechanism that specifically encourages peer availability. Note that [15] presents an incentive mechanism that is robust in the face of high turnover rate, but their mechanism does not explicitly promote availability.

Incentive mechanisms focused on promoting compliance have also been proposed for anonymous communication systems based on an anonymous core architecture. The reputation schemes of [14] and [21] assume a relatively small and static population of forwarding nodes, and are thus poorly suited to P2P systems. In [18], the authors propose a micropayment mechanism to allow users to pay for anonymity, implicitly providing an incentive for operators to offer such a service. In contrast, our approach uses payments to provide explicit incentives for high peer availability and low group turnover in a P2P system. The proposed scheme integrates payment with data forwarding to provide full anonymity and untraceability and shares some features with our protocol. However, the scheme in [18] require nodes to have access to tamper-proof hardware, an undesirable requirement in a P2P system.

8 Summary

Limited peer availability is a performance threat to many P2P systems, but is particularly harmful to P2P anonymity systems since its effect—a reduced average number of peers in the system— has a direct impact on primary metrics of anonymity. Free-riding due to low availability is likely to be more pervasive than non-compliance, as the latter requires users to obtain and execute modified software, while in the former users simply shut down their application.

We have explored the use of payments to construct an incentive mechanism that attaches a real monetary cost to free-riding, analytically evaluating the conditions under which such mechanism can increase peer availability. We also show that it is feasible to implement an efficient payment mechanism that does not compromise anonymity and that can be readily integrated with a class existing P2P anonymity systems.

The payment-based incentive mechanism proposed addresses the problem of low peer availability, which cannot be addressed with other mechanisms that are largely focused on compliance. It would be incorrect to infer from our work that compliance-enforcing mechanisms are not needed in anonymous systems. Rather, a comprehensive incentive solution for anonymous communication must address both compliance and availability, likely with separate design components each built upon the most suitable foundations. We have argued here, that payments are a particularly suitable mechanism for the availability component.

Acknowledgments

This research has been supported in part by the NSF under grant awards ANI-0085848 and EIA-0080119 and by CAPES (Brazil). Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

We thank Sonja Buchegger for fruitful comments on early drafts of this work.

References

[1] A. Acquisti, R. Dingledine, and P. Syverson. On the economics of anonymity. In *Proc. Seventh International Financial Cryptography Conference - FC03*, Jan 2003.

[2] B. Wilcox-O'Hearn. Experiences deploying a large-scale emergent network. In *Proceedings of the First International*

[12] No publicly available document describing this system was located. A discussion about its successes and failures can be found in [2].

Workshop on Peer-to-Peer Systems (IPTPS '02), Cambridge, MA, March 2002.

[3] A. Back, U. Moller, and A. Stiglic. Traffic analysis attacks and trade-offs in anonymity providing systems. *Proc. Workshop on Information Hiding (IH2001)*, 2001.

[4] N. Ben Salem and M. J. L. Buttyan, J.P. Hubaux. A charging and rewarding scheme for packet forwarding. In *Proc. of MobiHoc*, June 2003.

[5] K. Bennett and C. Grothoff. Gap – practical anonymous networking. *Proc. Workshop on Privacy Enhancing Technologies (PET)*, 2003.

[6] S. Brands. Untraceable off-line cash in wallets with observers. In *Proc. on Advances in cryptology (CRYPTO'93)*, volume 773, pages 302 – 318. Springer-Verlag, 1995.

[7] S. Buchegger and J.-Y. L. Boudec. The effect of rumor spreading in reputation systems for mobile ad-hoc networks. In *Proc. WiOpt'03 (Modeling and Optimization in Mobile, Ad Hoc and Wireless Networks)*, 2003.

[8] L. Buttyan and J.-P. Hubaux. Stimulating cooperation in self-organizing mobile ad hoc networks. *ACM/Kluwer Mobile Networks and Applications (MONET)*, 8(5), October 2003.

[9] D. Chaum. Untraceable electronic mail, return addresses, and digital pseudonyms. *Communications of the ACM*, 24(2):84–90, 1981.

[10] D. Chaum, A. Fiat, and M. Naor. Untraceable electronic cash. In *Proc. on Advances in cryptology (CRYPTO'88)*, volume 403, pages 319–327. Springer-Verlag, 1990.

[11] G. Debreu. *Handbook of Mathematical Economics*, volume 2, chapter 15 - "Existence of Competitive Equilibrium". Elsevier, 1984.

[12] R. Dingledine, N. Mathewson, and P. Syverson. Reputation in P2P anonymity systems. In *Workshop on Economics of Peer-to-Peer Systems*, Berkeley, CA, 2003.

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

[14] R. Dingledine and P. Syverson. Reliable MIX cascade networks through reputation. In *Proc. Sixth International Financial Cryptography Conference (FC02)*, Mar 2002.

[15] M. Feldman, K. Lai, J. Chuang, and I. Stoica. Robust incentive techniques for peer-to-peer networks. In *Proc. of ACM Conference on Electronic Commerce (EC'04)*, June 2004.

[16] N. Ferguson. Single term off-line coins. *Advances in Cryptology—EUROCRYPT '93, Lecture Notes in Computer Science*, 765:318–328, 1994.

[17] D. R. Figueiredo, J. K. Shapiro, and D. Towsley. Incentives for cooperation in anonymity systems. Technical Report MSU-CSE-05-10, Michigan State University, Dept. of Computer Science and Engineering, 2005. http://www.cse.msu.edu/∼jshapiro/anon-payments-tr.pdf.

[18] E. Franz, A. Jerichow, and G. Wicke. A payment scheme for mixes providing anonymity. In *Proc. Trends in Distributed Systems for Electronic Commerce (TREC'98)*, volume 1402 of *Lecture Notes in Computer Science*, pages 94 – 108. Springer-Verlag, 1998.

[19] M. J. Freedman and R. Morris. Tarzan: A peer-to-peer anonymizing network layer. In *Proc. of the 9th ACM Conference on Computer and Communications Security (CCS 2002)*, Washington, DC, November 2002.

[20] I. Goldberg. Zeroknowledge to discontinue anonymity service. Slashdot. URL: http://slashdot.org/comments.pl?sid=22261&cid=2388977, October 2001.

[21] P. Golle. Reputable mix networks. In *Proc. Workshop on Privacy Enhancing Technologies*, 2004.

[22] Y. Guan, X. Fu, R. Bettati, and W. Zhao. An optimal strategy for anonymous communication protocols. In *Proc. 22nd IEEE International Conference on Distributed Computing Systems (ICDCS 2002)*, Jul 2002.

[23] B. N. Levine and C. Shields. Hordes: A protocol for anonymous communication over the internet. *ACM Journal of Computer Security*, 10(3), 2002.

[24] P. Linga, I. Gupta, and K. Birman. A churn-resistant peer-to-peer web caching system. In *Proc. 1st ACM Workshop on Self-Survivable and Regenerative Systems*, Oct. 2003.

[25] M. G. Reed, P. F. Syverson, and D. M. Goldschlag. Anonymous connections and onion routing. *IEEE Journal on Selected Areas in Communication Special Issue on Copyright and Privacy Protection*, 1998.

[26] M. K. Reiter and A. D. Rubin. Crowds: anonymity for Web transactions. *ACM Transactions on Information and System Security*, 1(1):66–92, 1998.

[27] M. Rennhard and B. Plattner. Introducing morphmix: Peer-to-peer based anonymous internet usage with collusion detection. In *Proc. of the Workshop on Privacy in the Electronic Society (WPES)*, Washington, DC, November 2002.

[28] P. Resnick, K. Kuwabara, R. Zeckhauser, and E. Friedman. Reputation systems. *Communications of the ACM*, 43(12):45–48, 2000.

[29] S. Rhea, D. Geels, T. Roscoe, and J. Kubiatowicz. Handling churn in a dht. In *Proceedings of the USENIX Annual Technical Conference*, June 2004.

[30] R. L. Rivest and A. Shamir. Payword and micromint–two simple micropayment schemes. In *Proc. International Workshop on Security Protocols*, volume 1189 of *Lecture Notes in Computer Science*, pages 69 – 87. Springer-Verlag, 1996.

[31] A. Serjantov and G. Danezis. Towards an information theoretic metric for anonymity. *Proc. Workshop on Privacy Enhancing Technologies (PET)*, 2482, 2002.

[32] M. Wright, M. Adler, B. N. Levine, and C. Shields. Defending anonymous communication against passive logging attacks. In *Proc. of the IEEE Symposium on Security and Privacy*, Oakland, CA, May 2003.

[33] S. Zhong, J. Chen, and Y. Yang. Sprite: A simple, cheat-proof, credit-based system for mobile ad hoc networks. In *Proc. of Infocom 2003*, 2003.