Incentives for Cooperation in Anonymity Systems

Daniel R. Figueiredo Jonathan K. Shapiro Don Towsley *
Department of Computer Science
University of Massachusetts at Amherst

Computer Science Technical Report 03-21

June 24, 2003

Abstract

Like many peer-to-peer applications, anonymous communication systems are vulnerable to free-riders, peers who use the system while providing little or no service to others. To complicate matters, the identity of the free-rider is obscured by the very service anonymity systems are designed to provide, which limits the efficacy of conventional approaches for promoting cooperation (e.g., reputation mechanisms). We argue that the design constraints imposed by anonymity systems lead naturally to the notion of a currency that can be exchanged for service in order to provide incentives for cooperation. Based on this idea, we propose a novel technique to allow anonymous digital cash payments to be made to those who provide service. We incorporate this technique into a well-known peer-peer anonymous protocol (onion routing) while introducing modest message delay overheads and preserving its architectural simplicity. Finally, we formulate an abstract model of self-interested users in such a system and show that a payment based incentive mechanism can significantly improve the degree of anonymity by fostering greater cooperation among peers.

1 Introduction

The fundamental goal of anonymous communication is to disguise the identity of one or both parties involved in a bi-directional communication from each other and from any potential eavesdropper. This concept has been discussed in the literature for over two decades and has received attention from numerous researchers, who have targeted different aspects of such systems, such as formal definitions of anonymity, practical communication protocols and performance and service degradation under different malicious attacks. The recent increase in the interest in anonymous protocols can be related to the increasing need to address user privacy and security in on-line Internet services offered today. A fairly new direction of research within anonymous communications addresses the need for incentive mechanisms that will improve the performance and viability of such systems.

As noted in [4], anonymity is a property of communication that cannot be provided by the sender (or receiver) alone. The sender must rely on one or more nodes that will cooperate to disguise its identity. A

*This research has been supported in part by the NSF under different awards and in part by CAPES (Brazil)

number of peer-to-peer anonymous systems have been proposed in the literature [7, 8, 6]. In such systems a group of peers collectively obscure the identity of a message initiator by forwarding the message randomly among themselves an arbitrary number of times before sending it to its intended recipient and returning the recipient's response along the reverse path. In such mechanisms, the message initiator is anonymous within the group of collaborating peers. That is, from the recipient's perspective, all group members are equally likely to have initiated the message. Furthermore, peers who forward messages cannot distinguish between the true message initiator and an intermediate peer along the forwarding path. An important property of anonymous protocols is the *degree of anonymity* they provide and their robustness to certain types of malicious attacks [6, 13]. These metrics are usually monotonic in the number of peers in the group; having more peers confers a higher degree of anonymity and higher robustness to attacks.

Like many other peer-to-peer applications, anonymity systems are vulnerable to free-riders, nodes that consume the service without providing service to other nodes in the system. Free-ridership has been observed in many peer-to-peer systems. In file-sharing systems [2, 5], free-riders download files without making them available to other users. In ad hoc networks [3], free-riders decline to forward packets for others while still expecting their own packets to be forwarded. In anonymity systems, free-riders join the system when they need to establish anonymous communications, provide service while they are joined, but then leave the system once their needs are fulfilled. This behavior has the potential to undermine such systems because of two undesirable consequences. First, the presence of free-riders tends to reduce the overall number of peers in the group at any particular point in time, which consequently reduces the degree of anonymity. Second, the frequent turnover in group membership caused by free-riders joining and leaving the system imposes a high group maintenance overhead and can facilitate certain types of malicious attacks [13].

In this work, we consider a novel technique that uses digital cash to provide explicit incentives for cooperation in peer-peer anonymous systems. The main contributions of this paper are:

- We argue that payment based mechanisms that uses digital cash are very well suited to provide incentives for peers to provide service to other peers while preserving the fundamental anonymity properties and the architectural simplicity of anonymity systems. We show that digital cash has many properties that makes it well-suited for providing incentives in such systems, such as hiding the identity of the payer to the payee.

- We propose two payment based mechanisms that are readily coupled with the operation of an existent anonymous system (onion routing). The key idea of these mechanisms is to provide the initiator the ability to embed in each message sent small anonymous payments destined to those peers who forward the message along its path. Peers who desire service, can either join the system and accumulate cash by providing service to others, or can purchase service with an infusion of cash into the system. We argue that the overhead imposed by these mechanisms in terms of message latency are modest.

- We formulate an abstract optimization model of self-interested peers that are subject to the costs of using the proposed payment based anonymity system. We solve the model to obtain the optimal fraction of time a peer remains joined to the system and the optimal price it is willing to pay per message sent. Using these results, we show that the incentive mechanisms proposed can significantly improve the degree of anonymity by reducing the amount of free-riding. Interestingly, our results indicate that peers very sensitive to purchasing service can still use the system by providing service to others and accumulating revenue for its own demands.

A common approach to minimize impact of free-ridership is to design some incentive mechanism into the functionality of the system. Much recent work in this area has focused on reputation mechanisms, whereby peers—individually or collaboratively—identify free-riders and punish them by declining service to peers with bad reputations [3]. While reputation mechanisms are a promising approach to reducing the incentive to free-ride in many types of systems, they clearly require peers to know each other's identities. In anonymity systems, however, the identity of a free-rider is obscured by the very service the system is designed to provide. Thus, incentive mechanisms that do not require the identification of free-riders must be adopted.

In a system without explicit incentives, there is no clear reason for a peer to stay on-line after its anonymous communication is finished. A broader discussion of the economics (costs and benefits) behind anonymity has been discussed in [1]. Intuitively, if there are explicit benefits for a peer to stay joined to the system beyond the minimum amount of time needed to satisfy its own needs, then one would expect a more reliable and robust anonymous system.

The remainder of this paper is organized as follows. Section 2 provides a background of the issues involved in our system design and a brief discussion of the related work. In Section 3 we discuss why a payment mechanism is suited for anonymous systems. Section 4 presents an overview of onion routing followed by the incentive mechanisms proposed and a brief discussion of its performance. In Section 5 we analyze the incentive mechanism using a optimization model to capture peers' cost. Finally, Section 6 concludes the paper.

2 Background and Related Work

In the proposed incentive mechanism that follows, we make extensive use of digital cash and some of its more fundamental properties. Digital cash is a mature field of research and still a promising practical idea, despite the failure of initial attempts to bring digital cash to the public domain [12]. A reason for this failure was the meagered demand for such systems attributable to the lack of applications that required its use. Although digital cash infrastructures do not currently exist in the public domain, systems that make extensive use of digital cash, such as the one proposed here, could increase the demand and accelerate its deployment.

Thus, we will assume the existence of a publicly accessible authority issuing digital cash, which we will call the *Bank*, and a digital cash mechanism with fairly standard properties. Note that the *Bank* need not be a centralized entity nor be under single administration, as long as different entities recognize and value each others' currency. The most important property of the digital cash mechanism is that it renders the relationship between a payer and its purchases untraceable by either the Bank or by the payee. Ideally, a transaction should not even reveal partial information about the payer. Also, the digital cash mechanism should be implementable solely in software and not rely on tamper-proof hardware such as a smart card. Payments using digital cash can be either *off-line* or *on-line*. On-line payments require an interaction with the Bank for each transaction as it occurs. This allows the Bank to prevent malicious users from *double-spending*—repeatedly spending the same unit of currency with different payees. Off-line payments do not involve the Bank at the time of transaction and can only detect double-spending after the fact. Most off-line payment schemes provide a disincentive to double-spend by using cryptographic payment protocols which

reveal the identity of the payer if a unit of currency is double-spent, but not if it is spent only once [9].[1] Off-line payment protocols require a challenge-response interaction between payee and payer but do not involve the Bank at the time of payment. We do not require other generally desirable properties of digital cash, like efficient transferability and divisibility, although such properties could potentially be exploited to make our scheme more efficient. Both the on-line and off-line digital cash systems that will be used in the incentive mechanism that follows have been fully designed and their description can be found in [9] and references therein.

There have been recent efforts in applying reputation mechanisms to anonymous systems. In particular, Dingledine et. al ([4]) gives an overview of two different systems that were enhanced with reputation mechanisms to provide better service to its users. In the Remailer Networks system, a reputation mechanism is used to provide a more *reliable* service to users but it is not clear if it can prevent free-riding or improve the degree of anonymity (a remailer can join the system and behave adequately while receiving service and leave the system as soon as it finishes). The authors also describe their attempts to couple the Anonymous Publishing system with a reputation mechanism to prevent free-riding (publishing content without providing any reliable storage space). In both cases, the complexity of the system was significantly enlarged by the addition of reputation mechanisms, as acknowledged by the authors. The difficulties encountered suggest that conventional reputation mechanisms that rely on gossip might not be suitable to prevent free-ridership in anonymous systems.

A broad discussion of the costs and benefits behind anonymous systems is presented by Acquisti et. al [1]. The authors enumerate the economic forces that would drive users participating in an anonymous communication system and balance them in a model to understand the implications of user actions. Using a simplified model they show that under some circumstances the system is not feasible as the cost of anonymity exceeds its benefits. The authors also suggest that some form of payment could be used as an alternative incentive mechanism, but do not elaborate on the relationship between reputation mechanisms and possible payments nor provide any scheme to exploit such alternatives.

3 Motivations for Payment-based Incentive Mechanisms

In this section, we argue for a payment-based incentive mechanism for anonymous communication systems. As we will see, the idea of exchanging payment for service arises very naturally from an attempt to define a reputation mechanism that can operate effectively under the primary design constraints imposed by such systems—that the identity of a message initiator remain concealed from even the peers who provide service to it. It is worth noting that this constraint does not rule out the use of traditional reputation mechanisms in anonymity systems to influence behavior that can be detected through direct interaction between a non-cooperative peer and its neighbors along a forwarding path. However, we are concerned here with a particular type of non-cooperative behavior where a peer uses the system to deliver its anonymous messages but provides little or no forwarding service to others. For a reputation mechanism to be effective against this type of free-riding, it must be possible to recognize cooperative peers who are not directly connected.

Since a peer's reputation cannot be associated with its identity without compromising anonymity, the initiator of an anonymous message must instead present some evidence of cooperation within the messages

[1] It is reasonable to assume that among users of an anonymous communication protocol, such a disincentive would be particularly effective in preventing double-spending.

it is trying to send, and this evidence must not reveal its identity. Peers should be able to obtain this evidence in an anonymous and decentralized way as a byproduct of service provision. In a natural realization of such a scheme, evidence might take the form of a token that could be provided in exchange for service.

Tokens in such a mechanism must have certain key properties. First, tokens must be *irrefutable*, which typically requires that they be issued by a trusted third party. Second, tokens must be *issued anonymously*, by which we mean that token issuer should not be able to link the token itself to the peer's use of it to acquire service. Furthermore, each anonymous token should be used only once to acquire service; the attempt to use it twice should reveal information about the identity of the sender. The third property is that tokens must be *transferable* in the following sense: a token presented to the issuer can be exchanged for a new token that the presenter can then use to satisfy its own service requirements. This transferability property allows decentralized operation. Since possession of the token itself is evidence that the presenter has performed some useful service for another peer, the token issuer need not observe the actual service.

Mechanisms for the issue and exchange of such tokens already exist in the form of several digital cash schemes developed over the last decade. It is therefore reasonable to consider to what extent these mechanisms can be easily adopted to provide incentives in anonymous communication systems. There are two possible ways in which digital cash could be used to implement anonymous tokens. We could adopt useful aspects of digital cash without attaching to the tokens all of the attributes normally associated with money. Or, alternatively, we could simply use actual digital cash as tokens.

As the preceding discussion demonstrates, tokens clearly must possess irrefutability and transferability properties similar to money.[2] A property of money which may or may not be desirable in tokens is *fungibility*. A token is fungible if it can be exchanged for other forms of currency.

We argue that the properties of transferability and fungibility are inseparable. Although we may intend for tokens to be transferred in exchange for service, we cannot prevent tokens from being transferred in exchange for other things, like real money. Such unintended transfers make tokens fungible.[3] Thus, regardless of whether fungibility is a useful property of tokens, it is one that we must accept to achieve the necessary transferability.

Fungibility might first appear to be undesirable because it allows peers to be free riders by acquiring tokens for money without providing service for the system. Thus a payment-based scheme would not strictly prevent free-ridership. Such a scheme would, however, attach a real monetary cost to free-riding which could be avoided by providing service to others, thereby providing a financial incentive for cooperative behavior.

More important, fungibility provides a natural way to bootstrap the system if the token issuer agrees to exchange tokens for money. Nodes who pay would become a source of new tokens, which would propagate into the system to be collected by cooperative peers and used to acquire service free of charge.

We observe that the token issuer must perform many of the essential functions of a real electronic bank: issuing digital currency (tokens) and exchanging multiple forms of currency (money and tokens). It is therefore reasonable to consider whether a real electronic bank might best provide these functions. Security

[2] Although the *efficient* transfer of real money would not involve the bank in each transfer, the inefficient transferability described above is sufficient to implement our incentive mechanism. Digital tokens can be also made efficiently transferable. We leave a thorough evaluation of the design tradeoffs necessary to achieve efficient transfer in our proposed system for future work.

[3] A real-world example of this phenomenon is the well-documented E-Bay auctions of virtual entities used in the Ultima Online multi-player game [11]. The clearing price of such auctions establishes a rate of exchange for the item—which can be considered an digital token granting its holder certain abilities in the virtual world of the game.

and availability concerns suggest that building on an existing public digital cash infrastructure is a promising approach.

4 System Design

Before describing the proposed incentive mechanism, we first review the operation of a simplified onion routing protocol, as this is a fundamental aspect of our mechanism. In onion routing, anonymous paths are constructed from the peer that generates the anonymous message M (initiator node), through a set of collaborating peers in the system to some destination D. The initiator randomly selects the intermediary peers of the path and using their respective public keys, it constructs a message known as an onion, which has the form:

$$O = S_1, \{S_2, \{S_3, \{\ldots, \{S_L, \{R, D\}_{K_L^+}\}_{K_{L-1}^+}\} \ldots\}_{K_2^+}\}_{K_1^+} \quad (1)$$

where S_i is the address of the i-th peer in the path, L is the path length, and $\{X\}_{K_i^+}$ denotes message X encrypted with public key K_i^+.

This message is then forwarded to the first peer of the path, S_1. Each intermediary peer i in the path will have access to a payload after decrypting the message with its private key. The payload contains the address of the next hop S_{i+1} and an encrypted payload to be passed to this next hop. Eventually, the message reaches the last node in the path, which then forwards message R to destination D. A response message from D traverses the reverse path by using connection state installed when the onion is first forwarded.

The essence of our proposed scheme is to embed a small digital cash payment in each hop of the anonymous path. The initiator includes a payment C_i for the i-th hop of the path within the encrypted payload destined for that hop. The inherent source-routing mechanism (initiator determining the anonymous path) provided by onion routing is particularly well suited for integration with a payment mechanism, as the initiator can safely embed a payment for each hop in the path. We consider two possible extensions of the onion routing protocol to support digital cash payments—one using on-line payment and the other using off-line payment (see Section 2).

4.1 Protocol With On-Line Payments

In the on-line protocol, each intermediate node must contact the Bank to verify the validity of its payment. This will determine whether the received cash payment has been previously spent. To discourage double spending, each peer should validate its payment before returning any response messages along the reverse path, buffering data if necessary.

To prevent intermediate nodes from receiving payment without providing service, an acknowledgment mechanism is used to finalize the payment only after the payload has been properly forwarded. The payment for node i is encrypted with a symmetric key generated by the initiator and accessible only to node $i+1$. On receiving a message, node i sends the encrypted symmetric key to its predecessor in an acknowledgment message enabling the previous hop to decrypt the key and its payment. On receiving an acknowledgment from its successor, node i decrypts its own payment C_i and contacts the Bank to validate it. Figure 1 illustrates the operation of the on-line protocol.

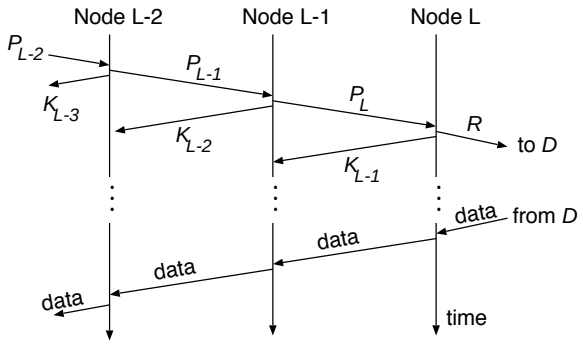

Figure 1: Message exchange for the last three hops of a path in the on-line protocol. Not shown in the diagram is an interaction with the Bank for each node to claim its payment.

The payload received by node i is

$$P_i = \{S_{i+1}, P_{i+1}, \{C_i\}_{K_i}, \{K_{i-1}\}_{K_{i-1}^+}\}_{K_i^+} \tag{2}$$

K_{i-1} is the symmetric key and the notation $\{X\}_K$ denotes a message X encrypted with symmetric key K.

The last peer in the path forwards the unencrypted request R to destination D. Since D is not assumed to participate in the anonymous protocol, the key to its payment is provided in the message payload to this final peer. Thus, the last payload is given by:

$$P_L = \{D, R, \{C_L\}_{K_L}, K_L, \{K_{L-1}\}_{K_{L-1}^+}\}_{K_L^+} \tag{3}$$

This protocol is vulnerable to malicious peer behavior along the path. For example, a peer may prevent its predecessor from receiving payment by neglecting to return an acknowledgment. It is worth emphasizing, however, that since no additional cash can be gained through this behavior, a self-interested peer is unlikely to engage in this action. Moreover, this is a scenario where conventional reputation mechanisms could be used to gossip about neighbors, and nodes with bad reputations would be less likely to appear in the paths, reducing the amount of money they receive from the system.

4.2 Protocol With Off-Line Payments

The off-line protocol differs from the on-line version mainly in how it prevents double spending. Instead of having to interact with the Bank when a payment is received, a node can accumulate payments and redeem them in batch at a later point in time, for example, when the node is idle or about to leave the system. By enabling batch redemption of payments the off-line protocol has the potential to greatly reduce the overheads associated with the inefficient transfer of digital cash, which can be burdensome in the on-line version.

Recall that the detection of double-spending using off-line digital cash requires the payee to issue a challenge to the payer. Since the payer must remain anonymous we provide a mechanism to forward the challenge from intermediate nodes back to the initiator along the reverse path. Challenges are suitably encrypted so that only the initiator can read them.

After forwarding the message to destination D, the last hop creates a message containing an encrypted challenge $\{q_L\}_{K_L}$ and sends it along the reverse path. Note that the key used to encrypt the challenge is the

symmetric key that was generated by the initiator to encrypt the payment for this node. Each node i on the reverse path appends its encrypted challenge $\{q_i\}_{K_i}$ as this message travels toward the initiator, who thus receives the following message containing encrypted challenges from all intermediate nodes:

$$\{q_L\}_{K_L}\{q_{L-1}\}_{K_{L-1}}\cdots\{q_1\}_{K_1}$$

After receiving the message with the challenges, the initiator constructs an onion containing a response r_i for each hop. When the final hop on the path receives its response r_L, it can then forward the data from the destination back toward the sender. Note that the final hop on the path will immediately send the message to the destination but should buffer any reply until r_L is received from the initiator.

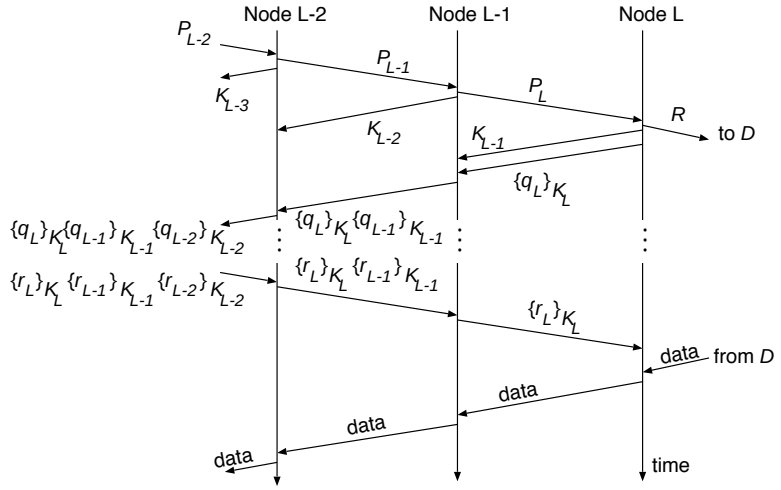

Figure 2: Message exchange for the last three hops of a path in the off-line protocol.

Like the on-line version, the off-line protocol also uses encrypted cash with encrypted keys that are returned through acknowledgments so that nodes must forward requests in order to receive their payment. Figure 2 illustrates the operation of the off-line protocol.

4.3 Discussion

The incentive mechanisms proposed above will introduce additional delay for each message sent through the system. In the discussion that follows, we characterize these delays and compare the two protocols with the unmodified onion routing protocol.

We first consider the on-line version of the protocol. Define T_L as the average round trip time from the initiator to the L-th hop in the path, and T_D as the average round trip time from node L to destination D. Let T_B be the average time required by any given peer to interact with the Bank. We assume Bank interactions occur in parallel with negligible slowdown. Since Bank interactions take place while waiting for destination D to reply to the message, the average response latency is at most:

$$T_L + T_D + T_B$$

Comparing this result with the unmodified onion routing protocol, which has an average response time of $T_L + T_D$, we observe that the delay overhead imposed by the on-line protocol is simply T_B.

The off-line version of the protocol requires an additional traversal through the anonymous path to pass challenges back to the sender and respective responses to the nodes along the path. This additional pass can be considered the performance cost of off-line transactions. Since this additional traversal can occur while waiting for destination D to reply to the message, the average response latency is at most:

$$2T_L + T_D$$

Note that since peers do not interact with the Bank at the time of the transaction, the overhead, T_B, is amortized over many messages and its contribution to the delay is negligible.

5 Incentive Model

In this section we present and evaluate a model for the incentive mechanism proposed in the previous section. The goal of this analysis is to provide some qualitative confirmation of our intuition that the proposed mechanism increases the level of cooperation in the system when users are sensitive to paying money or waiting to receive service. This increase in cooperation reduces the amount of free-riding and improves the overall degree of anonymity provided by the system. The approach taken here is to model each user of the system as a self-interested agent interested in optimizing a function of a small number of local decision variables. Since the optimal values of these variables will depend on the values chosen by other users, the system equilibrium will be defined by a fixed-point at which all users have simultaneously optimized their local objectives. In this section, we establish the existence of a system equilibrium and evaluate the conditions under which the level of cooperation at equilibrium will be higher than what would be achieved with no incentive mechanism

Consider the anonymity system described in the previous section, where a user must embed a payment of amount q for each hop of the forwarding path. For simplicity, we will assume that all paths through the system are of a fixed length L, thus the total cost to send a message is Lq. Let N denote the total number of users that are willing to use the system. We assume that user $i, 1 \leq i \leq N$, generates messages at a fixed rate λ_r^i that must be delivered anonymously through the system. In order to satisfy its anonymous message request, a peer must join the system for a minimum amount of time s, during which it cooperates by forwarding traffic.

Each individual user i optimizes over two decision variables—its level of cooperation c_i and the amount p_i that it pays from external funds for each message sent through the system. The level of cooperation c_i takes a value in the range $[s\lambda_r^i, 1]$ where $s\lambda_r^i$ is the minimum level of cooperation user i can have. c_i can be interpreted as the fraction of time the user is joined to the system. Note that c_i is bounded away from zero since a user must be joined to the system in order to receive service. Thus, letting s denote the time required for the system to deliver an anonymous message, we have that $s\lambda_r^i$ is the minimum fraction of time that a user is joined. Note that s is much smaller than λ_r^i, in particular, we assume $s\lambda_r^i$ to be much less than one. Although no assumptions are made with respect to the amount paid to send a message, there is no rational reason for a user to pay more than Lq—the total cost to send a message through the system.

As with the system defined in Section 4, users accumulate revenue while joined by forwarding messages generated by other peers. Let λ_c be the rate at which a joined user obtains revenue from the system. The value of λ_c is simply the aggregate rate of anonymous message requests generated by users divided by the

average number of users that are joined to the system. Thus,

$$\lambda_c = \frac{Lq \sum_{i=1}^{N} \lambda_r^i}{\sum_{i=1}^{N} c_i}. \tag{4}$$

Because a user only accumulates revenue while it is joined to the system, the rate at which user i accumulates revenue is $\lambda_c c_i$.

Of the total amount Lq of revenue required to send a message, user i will contribute p_i from external funds. The remaining balance $Lq - p_i$ must be collected by serving other users in the system. If the user has not yet collected this balance, then he/she will have to serve additional requests from other users in order to accumulate the necessary funds, thereby incurring a waiting time before his own message can be dispatched. The average time spent waiting per message sent will be denoted by $w(p_i, c_i)$. Observe that the waiting time is clearly a function of the local decision variables and rate parameter. For example, if a user is willing to pay the full price to send each of its anonymous messages ($p_i = Lq$), then his waiting time is zero. Similarly, if a peer is permanently joined to the system ($c_i = 1$) and has a relatively low message request rate (λ_r^i small), one might expect his waiting time to be near zero even if it is unwilling to pay anything ($p_i = 0$).

Each user then performs a local optimization to minimize a the weighted sum of the costs involved in participating in the system. In particular, we consider the following three costs:

Level of cooperation: We assume that users suffer some cost for being cooperative for a variety of reasons 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. Because these costs are difficult to quantify, we model them simply as a cost linear in the decision variable c_i.

Net cash flow rate: Users naturally value money, so there is a cost associated with paying money in order for a user to send messages anonymously through the system. Because users are both paying out and receiving money, each user will see a net cash flow rate, $\lambda_r^i p_i - \lambda_c c_i$, which we will treat as a cost.[4]

Average waiting time: Since users are also sensitive to delays, there is a cost associated with the average waiting time $w(p_i, c_i)$ before a message can be sent. An analytical expression for the average waiting time can be determined in a several ways; we consider a queueing model that captures the user's behavior within our system.

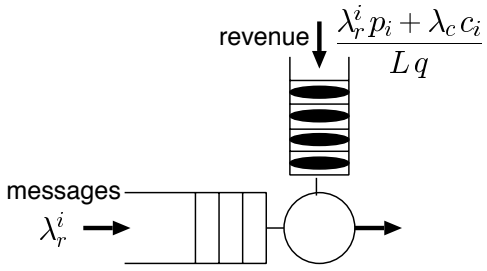

Figure 3: Leaky bucket—a model for the waiting time of a given user before it can send out an anonymous message.

The interaction between anonymous messages generated by a user and the amount of revenue accu-

[4]Note that if a user receives money at a higher rate than it spends, the net cash flow will take a negative value.

mulated is well captured by what is know as a *leaky bucket* model [10], illustrated in Figure 3. A leaky bucket is composed of two separate queues that are used to store messages and tokens, respectively. When a message arrives to an empty message queue and a token is available in the token queue, the message is immediately dispatched and one token is consumed. If no tokens are present when a message arrives, the message has to wait until a token is generated and only then is it dispatched. Both messages and tokens are generated according to fixed rates. In our case, messages represent anonymous messages generated by users at rate λ_r^i. Tokens map directly to revenue accumulated by the user, which is generated from his own payment and from being joined to the system. The overall rate at which a user accumulates revenue is given by $(\lambda_r^i p_i + \lambda_c c_i)/(Lq)$, which is normalized to match the cost of sending a single anonymous message. Assuming that both queues have infinite storage capacity and that both messages and tokens arrive according to Poisson processes, the waiting time of a message in the leaky bucket before it is dispatched, can be approximated by the waiting time of an M/M/1 queueing system [10]. Using a well-known result for the M/M/1 queue, the average waiting time for user i, is:

$$w(p_i, c_i) = \begin{cases} 0 & \text{if } p_i \geq Lq \\ 1/\left(\frac{\lambda_r^i p_i + \lambda_c c_i}{Lq} - \lambda_r^i\right) & \text{otherwise} \end{cases} \tag{5}$$

Note that we explicitly model the fact the waiting time is zero when $p_i \geq Lq$.

A necessary condition for stability in the leaky bucket model above is that the token arrival rate must be greater than the message arrival rate. This condition makes intuitive sense, since failure to satisfy it would mean that the user is unable to send all of his messages.[5]

In a diverse group of users, we expect each user to have different sensitivities to each of the costs enumerated above. For example, one user might be very sensitive to paying to send its anonymous messages while another might be more concerned with the waiting time. To capture the heterogeneity among users, we assign user-dependent weights α_i, β_i and γ_i to each of the costs.

Having laid the necessary groundwork, we may now write the local optimization for user i

$$\min_{p_i, c_i} \ \alpha_i w(p_i, c_i) + \beta_i(\lambda_r^i p_i - \lambda_c c_i) + \gamma_i c_i \tag{6}$$

subject to

$$\lambda_r^i \, p_i \sum_{j=1}^N c_j - \lambda_r^i \, L \, q \sum_{j=1}^N c_j + c_i \, L \, q \sum_{j=1}^N \lambda_r^j > 0 \tag{7}$$

$$s\lambda_r^i \leq c_i \leq 1, \ 0 \leq p_i. \tag{8}$$

The feasible region for this optimization is defined by the leaky bucket stability constraint which is expanded in equation (7) and the bounds on the decision variables (8). It can be shown that the objective function (6) is convex within this feasible region and that the feasible region is itself a convex set. By these convexity properties, we can conclude that there is a unique feasible optimal solution for each user. However, due to the nonlinearity of constraint (7) and the objective (6) (note the product $p_i \, c_i$), solving this optimization problem analytically is challenging and we therefore rely on numerical solution methods. However, a simple approximation obtained by ignoring the nonlinear term could be considered in targeting an analytical solution, although we leave this for future investigations.

[5] An alternate modeling option here would be to allow the user to reduce its anonymous message rate to match available revenue. We do not consider this option here.

5.1 Solving the model

To make our model tractable, we aggregate users into a fixed number of classes, M, where users within each class have identical behavior. (Note that subscript i will now denote a class and not a user). We will also assume that each class contains an identical number of users, although this assumption can readily be relaxed by associating weights with the different classes. Using a numerical solution technique, we can solve the resulting fixed point problem by solving the local optimization problem for each class sequentially and iterating until the solutions for all classes stabilize. This stable set of decision variables defines the system equilibrium.

We consider two classes of users, $M = 2$, where users in class 1 are very sensitive to paying but are relatively insensitive to waiting, meaning that payments from external funds have more value than waiting for service or remaining joined to the system. This behavior is modeled accordingly by adjusting the weights that balance these different costs; a larger weight value indicates higher sensitivity. Users in class 2 have the opposite behavior and very sensitive to waiting for service while being more willing to pay for service. In the first case study that follows, both classes are equally sensitive to remaining joined to the system, while in the second case, class 1 is more sensitive to being joined[6]. Using these two classes, we investigate how the system equilibrium given by the optimal choices for the price paid per message $p^* = \{p_1^*, p_2^*\}$ and for the level of cooperation $c^* = \{c_1^*, c_2^*\}$, will differ between each class under different demands for anonymous message (λ_r^1 and λ_r^2).

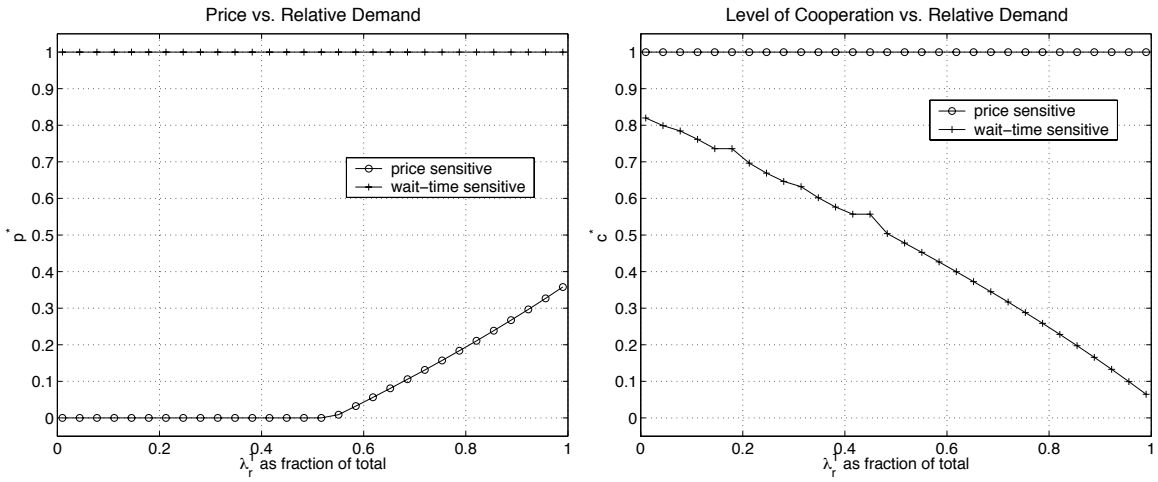

Figure 4: Optimal prices and levels of cooperation for classes 1 and 2 as a function of the fraction of total demand for the two classes.

We start by inspecting the system equilibrium as a function of the fraction of demand generated by the two classes. Figure 4 illustrates both p^* and c^* when the total message demand is kept constant ($\lambda_r^1 + \lambda_r^2 = 1$), but the fraction of the demand generated by each class varies (x-axis denotes the fraction generated by class 1). Note that users from class 2, who are insensitive to making payments, will always pay the full price to send their messages ($p^* = Lq = 1$), independent of the fraction of demand they generate. In contrast, users from class 1 will not pay if the fraction of messages they generate is below 0.5, and will only start paying if

[6]Parameters for this model are: $\alpha_1 = 5$, $\beta_1 = 50$, $\gamma_1 = 5$, $\alpha_2 = 50$, $\beta_2 = 5$, $\gamma_2 = 5$, ($\gamma_2 = 4.5$), $L = 10$, $q = 0.1$, $s\max\{\lambda_r^1, \lambda_r^2\} \leq 0.01$

this fraction is larger.

Contrary to users from class 1 who remain joined to the system 100% of the time independent of the fraction of their demand, users from class 2 increase their fraction of time joined to the system as their fraction of demand increases, as illustrated by Figure 4. Since class 2 peers have a non-negligible sensitivity in paying for service, a user can recover part of its cost by serving message requests from others by remaining joined to the system when their fraction of demand is high. An interesting observation is that the revenue provided by users that are more likely to free-ride (being insensitive to payments), can under certain conditions, enable price sensitive peers to use the service free of charge and even generate profit by remaining joined to the system.

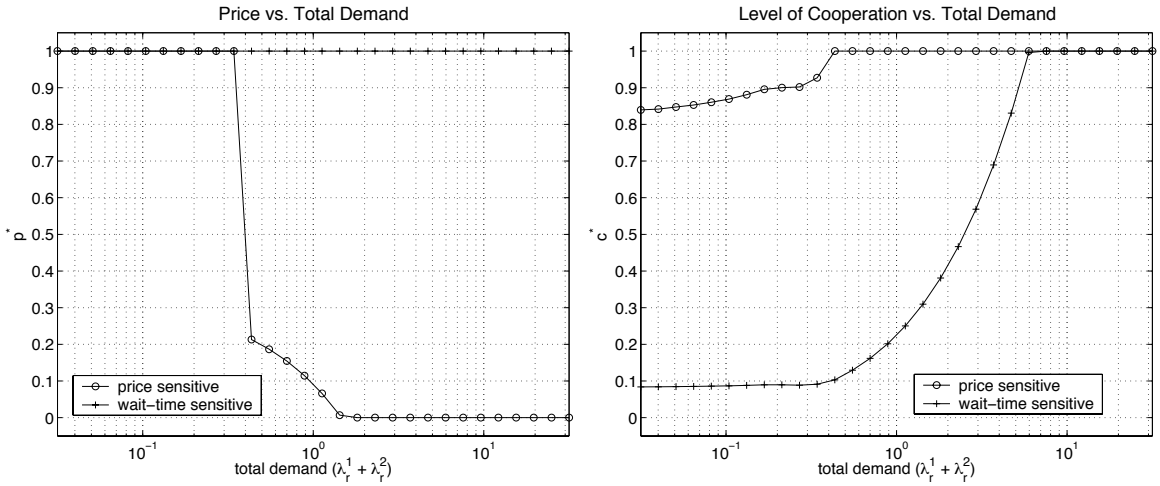

Figure 5: Optimal prices and levels of cooperation for classes 1 and 2 as a function of the total demand on the system. For these plots, the ratio for of demand for the two classes is fixed at $\lambda_r^1/\lambda_r^2 = 1$.

Figure 5 illustrates the system equilibrium as a function of the total demand on the system, when the ratio λ_r^1/λ_r^2 is kept constant (equal to one). As with the previous result, we see that class 2 pays full price to send its messages independent of its demand. However, class 1 pays only if it is not able to collect revenue from the system, which occurs when demands are low. When demands increase, class 1 stops paying altogether to send its messages. In particular, class 1 takes advantage of the willingness of class 2 to pay for immediate service to accumulate revenue for its own messages.

Class 1 always has a high level of cooperation, however, at low demand rates, the cost of remaining joined to the system is larger than the benefits of the additional revenue that could be collected, which keeps class 1's level of cooperation below 100%. Although class 2 has a low level of cooperation when the demand is low, it is interesting to note that even this low level of cooperation is above the minimum level required to satisfy its own demand. As demand increases, the level of cooperation increases, as class 2 sees the opportunity to recover part of its costs by serving other users. This behavior is illustrated in Figure 5.

In both case studies above, we note that incentive mechanisms induce a much higher level of cooperation than otherwise would occur. In particular, if no incentive mechanism is adopted, a user would be joined solely for the time required to satisfy its own demand. This fraction of time is the ratio between s, the time to service a single request, and $1/\lambda_r^i$, the message interarrival time. In practice we expect s to be at least one order of magnitude lower than $1/\lambda_r^i$, which leads to a level of cooperation of at most 0.1 assuming that users

act in pure self-interest. Although this assumption may not hold in real systems, the results shown above illustrate that incentive mechanisms such as the ones proposed, can provide a significant improvement on the level of cooperation, and hence, the degree of anonymity provided by the system.

6 Conclusion

This paper presented a novel technique to provide incentives for cooperation in peer-peer anonymous communication protocols. We argue that payment based mechanisms that use digital cash are very well suited to provide such incentives while preserving the fundamental anonymity properties and the architectural simplicity of one such system.

We propose two extensions to the onion routing protocol that make use of on-line and off-line infrastructures for digital cash, respectively. The key idea of such mechanisms is to provide the initiator the ability to embed anonymous payments to those peers who perform forwarding services. The overheads imposed by these mechanisms in terms of message latency are modest.

We formulate an abstract model of self-interested users that are subject to the costs of using a payment based anonymity system, such as the one proposed here. Using this model we demonstrated that the incentives provided can significantly improve the degree of anonymity by fostering greater cooperation among peers.

Although we show the existence of a system equilibrium using a centralized solution technique, it is not yet clear how users would achieve this equilibrium in a decentralized setting. In addition, we have assumed that the price paid and received for forwarding a message a single hop, q, is determined a priori, whereas in practice this price would likely be set by market mechanisms. An interesting question is to understand the impact that a fluctuating price hop price will have on the system equilibrium.

Our approach relies heavily on the source-routing and information hiding properties of the onion routing protocol to make payments accessible only to designated nodes on the path. A challenge for future work is to consider how digital cash might be used in a lighter weight anonymity protocol, such as Crowds [8], which lacks these properties.

References

[1] Alessandro Acquisti, Roger Dingledine, and Paul Syverson. On the economics of anonymity. In *Proc. Seventh International Financial Cryptography Conference - FC03*, 2003.

[2] Eytan Adar and Bernardo A. Huberman. Free riding on gnutella. *First Monday*, 5(10), 2000.

[3] S. Buchegger and J.-Y. Le 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.

[4] Roger Dingledine, Nick Mathewson, and Paul Syverson. Reputation in P2P anonymity systems. In *Workshop on Economics of Peer-to-Peer Systems*, Berkeley, CA, 2003.

[5] Philippe Golle, Kevin Leyton-Brown, Ilya Mironov, and Mark Lillibridge. Incentives for sharing in peer-to-peer networks. *Lecture Notes in Computer Science*, 2232, 2001.

[6] Brian Neil Levine and Clay Shields. Hordes: A protocol for anonymous communication over the internet. *ACM Journal of Computer Security*, 10(3), 2002.

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

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

[9] Bruce Schneier. *Applied Cryptography*. John Wiley & Sons, 2nd edition, 1996.

[10] Mischa Schwartz. *Broadband Integrated Networks*. Prentice-Hall Inc., 1996.

[11] Unattributed. Where there's mud, there's brass. *Economist Magazine*, July 8 2000.

[12] Peter Wayner. Electronic cash for the net fails to catch on. *The New York Times*, November 28 1998.

[13] Matt Wright, Micah Adler, Brian N. Levine, and Clay Shields. An analysis of the degradation of anonymous protocols. In *Proc. ISOC Network and Distributed System Security Symposium (NDSS 2002)*, Feb 2002.