Designing Incentives for P2P Multimedia Sharing

Yu Zhang
Electrical Engineering
University of California, Los Angeles
Los Angeles, United States
yuzhang@ucla.edu

Mihaela van der Schaar
Electrical Engineering
University of California, Los Angeles
Los Angeles, United States
mihaela@ee.ucla.edu

Abstract—**The design of incentive schemes for P2P multimedia sharing networks is challenging due to the unique features exhibited by such networks: large populations of anonymous peers interacting infrequently, asymmetric interests of peers, network errors, and multiple concurrent transactions. In this paper, we design and rigorously analyze a new family of incentive protocols that utilizes social norms for multimedia sharing. In particular, we show that, given the network and peers' characteristics, social norms can be effectively designed to deter free-riders by reducing their reputations and thus, the services which they receive from the P2P network. Unlike existing research, which deploys ad-hoc reputation schemes in P2P networks, our proposed framework rigorously determines the optimal social norm and associated reputation scheme to be used by a particular P2P system, characterized by its specific network characteristics. We also investigate how the design of the optimal social norms needs to change to account for the impact of altruistic and malicious peers. Our results show that optimal social norms are capable of providing significant improvements in the sharing efficiency of multimedia P2P networks. Specifically, depending on the network environment, the proposed social-norm based P2P protocols are able to outperform conventional Tit-for-Tat protocols by up to 8dB in terms of video quality.**

***Keywords-** Multimedia sharing; Peer-to-Peer networks; Incentive design; Indirect reciprocity; Social norms; Reputation schemes*

I. INTRODUCTION

Multimedia sharing over a variety of P2P wired and wireless networks is expanding [1][2]. However, P2P networks are well-known to being vulnerable to intrinsic incentive problems. Since peers contributing their content or resources do not receive direct benefits to providing their services, peers tend to avoid uploading while trying to download content from other peers, a behavior commonly known as *free-riding* [3][4].

Many of the existing mechanisms to design incentives to encourage cooperation and mitigate free-riding rely on game-theoretical approaches such as pricing and reciprocity [4]. Pricing mechanisms incentivize peers to share their content by rewarding them with virtual currency for uploading and charging them for downloading. However, such solutions are often very cumbersome to deploy because they require an accounting infrastructure to track the transactions of peers [5]. In reciprocity mechanisms, the peers' past reciprocative behavior is rewarded or punished in future interactions by differential service schemes. Depending on how a peer's rating is generated, reciprocity-based protocols can be classified as direct reciprocity (also known as personal reciprocation) and indirect reciprocity (also referred to as societal reciprocation) [5].

In direct reciprocity, each peer rates a specific peer individually [3]. Though easy to implement, direct reciprocity requires frequent interactions between two peers in order to establish accurate mutual ratings, which is restrictive in P2P networks characterized by high churn or asymmetry of interests. Due to the random matching feature of large P2P networks, indirect reciprocity based on reputation mechanisms becomes a more appropriate mechanism in designing incentive protocols [7][8]. In P2P systems operating based on indirect reciprocity, a peer is globally rated with a reputation calculated by its past behavior in the network. In order to make a decision, a peer does not need to know the entire reciprocation history (actions taken by a peer) but only the reputation of its opponent. However, existing research on indirect reciprocity in P2P networks is mainly experimental in nature and does not provide a rigorous foundation for the design of optimal protocols. As will be shown in the simulation section of this paper, designing such protocols in an ad-hoc manner results in a significantly reduced performance for the P2P network.

In this paper, we propose a theoretical framework for indirect reciprocity based on *social norms* [9], in order to design efficient incentive protocols for P2P multimedia sharing services. Hence, we design social norms, which consider the unique features and constraints of P2P multimedia sharing services:

- *Asymmetry of interests among peers*. To accommodate the fact that the peers' interests are asymmetric, we model the interaction between a pair of matched peers as a gift-giving game, instead of a prisoner's dilemma game, which assumes mutual interests between a pair of peers.
- *Service errors*. In contrast to the existing literature [7][8], our work explicitly takes into consideration that the exchange of multimedia data between peers may be subject to service errors and considers how protocols can be efficiently designed given the level of network errors.
- *Multiple connections*. In multimedia sharing applications over P2P networks, peers can engage in multiple simultaneous connections with other peers to exchange data in order to increase the download efficiency. Hence, we accommodate sharing using multiple connections and explicitly analyze how the number of connections will impact the peers' incentives and the social welfare of the P2P network.
- *Altruistic and malicious peers*. We also rigorously determine the impact of altruistic peers (who always provide upload services to other peers) as well as malicious peers (who upload corrupted data to other peers) on the protocol design and the P2P network's performance.

The remainder of the paper is organized as follows. In Section II, a game-theoretic model for P2P multimedia sharing is proposed. In Section III, the problem of designing the optimal incentive protocol is formalized and the structure of the optimal protocol is studied. Section IV explicitly investigates the impact of altruistic and malicious peers on the performance and robustness of the incentive protocols. After showing the simulation results with illustrative examples in Section V, we conclude the paper in Section VII.

II. SYSTEM MODEL

A. Considered P2P networks

We consider a P2P multimedia sharing network such as CoolStreaming [1], where peers would like to associate themselves with other peers that possess media content in which they are interested. The shared media content is coded and divided into media chunks by the content creator. Here we define the value (benefit) of a chunk as its dependency factor on other chunks, which represents the video distortion reduction on the peer who receives this chunk [16]. In general, both the value and the size of a chunk may depend on the priority class to which it belongs to (e.g. they could be base and enhancement layers like in a scalable video coder; the I, P and B-frames of an H.264/AVC codec etc.). To make the analysis tractable, we assume that the multimedia chunks are of equal size and have the same value (benefit) by appropriately selecting the video packets that form the various chunks using the methods discussed in [16].

At any instance, a peer buffers an amount of chunks that can be shared with others, and the trackers maintain and update periodically the buffer maps recording the content possession of each peer. We consider a discrete-time model, in which time is divided into periods representing the interval between two updates of the buffer maps by the trackers. We assume that there is a continuum of peers in the network, which is a good practical model for large-scale P2P networks [7]. When a peer wants to download a certain chunk, it sends a search request to the tracker from which it receives a response with the list of peers who have the requested content [1]. Then the peer randomly selects a peer from the list to send a service request. The selection is uniformly random such that all peers on the list have an equal probability to be chosen [10][11]. At any instance, an individual peer can support simultaneously a fixed number of b download connections, from which it downloads chunks it requests from others [1].

B. The stage game played by a pair of peers

One transaction between a pair of connected peers exchanging a chunk can be modeled as a one-stage asymmetric gift-giving game. To avoid confusion, the peer who requests the downloading of a chunk is called a *client* and the peer who is being requested is called a *server*. In one transaction, the server has the choice of selecting its action a from the set $\mathcal{A} = \{S, NS\}$, where S (Serve) implies that the server responses to the client's request to upload the chunk; whereas NS (Not Serve) implies that the server refuses to upload the chunk.

- If $a = S$, the server consumes an upload cost of c , and the client receives a benefit of r . In our formal analysis, we consider that c and r are constant for each chunk, with $r > c$ such that the sharing service provided by the P2P network is socially valuable.
- If $a = NS$, both the server and client receive a utility of 0.

Since each peer can maintain multiple simultaneous connections, the utility it receives in one period is the sum utility from all the transactions in which it is involved. The social welfare of the network is quantified by the social utility U that is defined as the average utility of all peers in one period. The social utility is maximized when all servers choose $a = S$ in their transactions. In contrast, a self-interested server will choose $a = NS$ to maximize its stage-game utility myopically, which gives rise to an undesirable zero social utility.

C. Social norms

We adopt a repeated game formulation to model the subsequent interactions among peers. P2P protocols based on social norms are considered in order to improve the inefficiency of the myopic equilibrium. Social norms define the rules that the group of peers uses to reward or punish appropriate and inappropriate behaviors in the P2P network. Since we focus on protocols that are based on social norms, we use the two terms "protocol" and "social norm" interchangeably in the rest of the paper.

In the repeated game, each peer is tagged with a reputation θ from the finite set $\Theta = \{0,1,2,\cdots,L\}$, representing its social status. For notational convenience, a peer of reputation θ is referred to as a θ -peer. The reputation of the peers is maintained and updated by a trustworthy third-party device, e.g. the tracker.

A social norm κ , is determined by the P2P protocol designer, which is composed of a social strategy σ and a reputation scheme τ . σ is a *reputation-based behavioral strategy* represented by a mapping $\sigma : \Theta \times \Theta \to \mathcal{A}$. It specifies what action $\sigma(\theta,\tilde{\theta}) \in \mathcal{A}$ should a server of reputation $\theta \in \Theta$ select when meeting a client of reputation $\tilde{\theta} \in \Theta$. τ updates a peer's reputation at the beginning of each period in our framework. Specifically, the tracker reviews all upload transactions of a peer with the result of the review recorded in a variable $x \in \{0,1\}$. At the beginning of a period, x is reset to 0. Then in each transaction, there is a mapping ϕ which maps the reputations of the peer and its client as well as the peer's action during one transaction into a binary value as $\phi : \Theta \times \Theta \times \mathcal{A} \to \{0,1\}$. If the action is in accordance to the social strategy, ϕ outputs 0 indicating that the peer behaves well in this transaction; otherwise, ϕ outputs 1 indicating that the peer does not comply with the social norm. After the transaction, x is updated by an OR-operation as $x := x \vee \phi$. Hence, after one period, $x = 0$ if and only if the peer complies with the social norm in all of its upload transactions. Based on the peer's current reputation and x , τ then determines its new reputation as $\tau : \Theta \times \{0,1\} \to \Theta$. The mapping rule is as follows: if $x = 0$, τ rewards the good behavior of the peer with an increased reputation; on the other hand, if $x = 1$, τ punishes the peer for not uploading sufficient content in this period with a decreased reputation.

In our framework, a peer's upload action in one transaction is reported by its client. We here assume that the client always makes a truthful report [1]. However, we do consider the impact of network (service) errors. With the probability ε ($0 < \varepsilon \ll 1$), a peer which intends to upload a chunk in one transaction fails to do so due to a connectivity error.

We restrict our attention to a set of threshold-based strategies Γ . Every strategy $\sigma \in \Gamma$ can be characterized by a service thresholds $h(\sigma) \in \{1,\cdots,L\}$ [2], which can be specified as follows

$$\sigma(\theta,\tilde{\theta}) = \begin{cases} S & if\ \theta \geq h(\sigma)\ and\ \tilde{\theta} \geq h(\sigma) \\ NS & otherwise \end{cases}. \tag{1}$$

By adopting σ , peers with reputation being at least $h(\sigma)$, which are called "active peers", will mutually help each other, while peers with reputation lower than $h(\sigma)$, referred to as "inactive peers", cannot download chunks from others and do not need to upload chunks to others. To avoid confusion, the prescribed social strategy is denoted as σ_o and the corresponding prescribed service threshold of the social strategy is denoted as h_o .

To keep the initial design of the P2P protocol simple, we consider a reputation scheme τ that provides the harshest punishments to peers when they do not comply with the social strategy. The reputation update rule can be written as follows

$$\tau(\theta,x) = \begin{cases} \min\{L,\theta+1\} & if\ x = 0 \\ 0 & if\ x = 1 \end{cases}. \tag{2}$$

[1] The extension to the untruthful report from clients is discussed in [13].

[2] Here the strategies with the service threshold being 0 and $L+1$ are not considered.

A schematic representation of a social norm is provided in Figure 1. , with (a) illustrating the decision process of a social strategy, where $\tilde{\theta}$ denotes the reputation of the client in one transaction, and (b) illustrating the decision process of a reputation scheme.

D. Utilities

In this work, we assume that each active peer generates chunk requests at a constant rate [7]. In one period, each peer generates a constant request of λb chunks, where λ can be interpreted as the rate at which each connection is utilized per period [6]. Once the download request is rejected, the peer immediately redirects this request to another peer on the list provided by the tracker until it is matched with a peer who accepts its request. Hence, an active peer can always download λb chunks in one period. Due to the random matching feature of the network, the chunks uploaded by an active peer per period is also λb. In summary, the expected one-period utility of a peer with $\theta \geq h_o$ can be expressed as

$$v_\kappa(\theta) = \lambda b[(1-\varepsilon)r - c]. \tag{3}$$

We evaluate a peer's expected overall utility as the sum of its expected one-period utility in the current period and its discounted expected overall utility starting from the next period, i.e.

$$v_\kappa^\infty(\theta^{t_0}) = v_\kappa(\theta^{t_0}) + \delta\sum_{\theta'} p_\kappa(\theta' \mid \theta) v_\kappa^\infty(\theta'). \tag{4}$$

$p_\kappa(\theta' \mid \theta)$ denote the transition probability of a peer's reputation across periods when following κ as follows

$$p_\kappa(\theta' \mid \theta) = \begin{cases} 1-\alpha, \ \theta \geq h_o \ \ and \ \theta' = \min\{L, \theta+1\} \\ \alpha, \ \theta \geq h_o \ \ and \ \theta' = 0 \\ 1, \ \theta < h_o \ \ and \ \theta' = \theta + 1 \\ 0, \ otherwise \end{cases} \tag{5}$$

where $\alpha = 1-(1-\varepsilon)^{\lambda b}$ is the probability that an active peer who complies with the social norm is falsely punished due to the service error.

The social utility of the network is regarded the average expected one-period utility over all peers and hence depends on the reputation distribution of the peer population, which is denoted by $\{\eta(\theta)\}_{\theta=0}^{L}$ with each term $\eta(\theta)$ representing the fraction of peers in the total population holding a reputation θ. Due to the reputation update in each period, $\{\eta(\theta)\}$ evolves dynamically over time. Since we are interested in the long-term utilities of peers, we study the stationary distribution of reputations, which does not change over time and is computed from [13] as follows with $\mu_\kappa = 1/(1+\alpha h_o)$,

$$\begin{aligned} &\eta_\kappa(L) = 1-(1+h_o\alpha)\mu_\kappa + (1-\alpha)^{L-h_o}\mu_\kappa \\ &\eta_\kappa(\theta) = (1-\alpha)^{\theta-h_o}\alpha\mu_\kappa, \ h_o+1 \leq \theta \leq L-1 \\ &\eta_\kappa(\theta) = \alpha\mu_\kappa, \ 0 \leq \theta \leq h_o \end{aligned} \tag{6}$$

Therefore, the social utility of the network is defined as the expected one-period utility averaged over all peers when the reputation distribution is stationary

$$U_\kappa = \sum_\theta \eta_\kappa(\theta) v_\kappa(\theta) = \lambda b \mu_\kappa[(1-\varepsilon)r - c]. \tag{7}$$

III. Optimal Design of Social-Norm Based Protocols

A. Defining sustainability in P2P networks

Since we consider a non-cooperative scenario, in order to ensure that a peer has no incentive for deviating unilaterally from the social norm, we need to check whether a peer can improve its expected overall utility by deviation. Particularly,

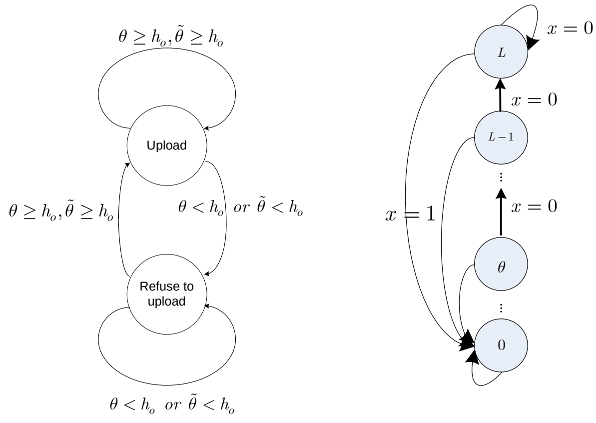

(a) (b)

Figure 1. The schematic representation of a social norm

we would like to design protocols which are social norm equilibria as defined in [9]. If a protocol is a social norm equilibrium, the sum of a peer's instant utility and its expected future utility thereafter by complying with the protocol is always larger than the sum of utilities by deviating to any behavioral strategy other than the social strategy.

In [13], we proved that the sufficient and necessary condition for a protocol to be a social norm equilibrium is the one-shot deviation principle, which is stated next. (The proof is omitted here due to space limitations.)

Lemma 1 (One-shot Deviation Principle). κ is a social norm equilibrium if and only if for any θ, there is no profitable one-shot deviation, i.e.

$$\begin{aligned} &c_\sigma(\theta) - c_\kappa(\theta) \leq \\ &\delta\left[\sum_{\theta'} p_\kappa(\theta' \mid \theta) - \sum_{\theta'} p(\theta' \mid \theta, \sigma)\right] v_\kappa^\infty(\theta') \end{aligned} \quad \text{for all } \sigma, \tag{8}$$

where $c_\kappa(\theta)$ denotes the one-period cost consumed by a θ-peer following the social norm κ; $c_\sigma(\theta)$, $p(\theta' \mid \theta, \sigma)$, and $v_{\kappa,\sigma}^\infty(\theta)$ are a peer's incurred cost per period, its reputation transition probability from θ to θ' and its expected overall utility, respectively, when it plays σ and the protocol designer implements the social norm κ.

Using Lemma 1, we can derive incentive constraints that characterize social norm equilibrium. There are two cases that need to be considered. When an active peer with reputation θ receives upload requests from another active peer with reputation $\tilde{\theta}$, the protocol should provide the θ-peer incentives to choose $a = S$ over $a = NS$. Thus, the resulting expected overall utility of the θ-peer is given by

$$V_\theta(S) = -\lambda bc + \delta[(1-\alpha)v_\kappa^\infty(\min\{L, \theta+1\}) + \alpha v_\kappa^\infty(0)]. \tag{9}$$

In contrast, if the peer deviates by refusing to upload the requested chunk and play $a = NS$, its expected overall utility is $V_\theta(NS) = \delta v_\kappa^\infty(0)$. As (8) specifies, we have that

$$\delta(1-\alpha)[v_\kappa^\infty(\min\{L, \theta+1\}) - v_\kappa^\infty(0)] \geq \lambda bc. \tag{10}$$

When $\tilde{\theta} < h_o$, the social norm should provide θ-peers incentives to choose $a = NS$ over $a = S$, with the resulting incentive constraint being

$$\delta(1-\alpha)[v_\kappa^\infty(\min\{L, \theta+1\}) - v_\kappa^\infty(0)] \geq -c. \tag{11}$$

B. Design problem of optimal sustainable social norm

Based on the above discussion, a P2P protocol can be designed by selecting three parameters: the punishment length L, the service threshold h_o, and the maximal number of concurrent connections b. The protocol designer aims to choose a protocol that maximizes the social utility (i.e. sharing level among peers) among the candidate protocols that can be sustained as social norm equilibria and thus, the problem of designing the optimal protocol in this paper can be formalized as follows (we call this problem "optimal social norm equilibrium - OSNE")

$$\begin{aligned}&\underset{(L,h_o,b)}{\text{maximize}} \quad U_\kappa = \lambda b \mu_\kappa [(1-\varepsilon)r - c] \\ &\text{s.t. } \delta(1-\alpha)[v_\kappa^\infty(\min\{\theta+1,L\}) - v_\kappa^\infty(0)] \geq \lambda bc, \ \forall \theta \geq h_o, \\ &\quad \delta(1-\alpha)[v_\kappa^\infty(\min\{\theta+1,L\}) - v_\kappa^\infty(0)] \geq -c, \ \forall \theta < h_o.\end{aligned} \quad \text{(OSNE)}$$

We have proved in [13] that the highest social utility that can be achieved for a social norm equilibrium always monotonically increases with L. Therefore, we consider the design problem of (h_o, b) given a value of L which can be selected based on the desired complexity of the protocol.

C. Designing and characterizing the optimal social norm

In this section, we explicitly analyze how the design parameters (h_o, b) will impact the social utility as well as the peers' incentives to comply with the prescribed protocol. This analysis enables us to characterize the optimal design, denoted as (h_o^*, b^*), which maximizes the social utility while providing peers sufficient incentive to comply with the protocol.

First, we can verify from Problem (OSNE) that U_κ monotonically increases with b. On the other hand, since $\mu_\kappa = 1/(1+\alpha h_o)$ monotonically decreases with h_o, we can also conclude that U_κ monotonically decreases with h_o.

We then provide the following proposition to establish what conditions should (h_o, b) fulfill (i.e. how should these parameters be selected by the protocol designer) in order to sustain the resulting protocol as a social norm equilibrium.

Proposition 1. A protocol $\kappa = (\sigma_o, \tau)$ can be sustained as a social norm equilibrium if and only if

(1) its service threshold h_o is larger than a constant H_o that is defined as

$$h_o \geq H_o \triangleq \ln\left[1 - \frac{(1-\delta)c}{(1-\alpha)\left[(1-\varepsilon)r - c\right] - \delta\alpha c}\right] / \ln \delta \,; \quad (13)$$

(2) the maximum number of concurrent connections b is smaller than a constant B, which is the solution of the following equation set

$$\begin{aligned}&\delta(1-\alpha)[1-\delta^{h_o}]\frac{[(1-\varepsilon)r - c]}{1-\delta(1-\alpha)-\alpha\delta^{h_o+1}} = c \\ &\alpha = 1-(1-\varepsilon)^{\lambda B}\end{aligned}. \quad (14)$$

Proof: See [13]. ■

Proposition 1 provides a guideline for selecting the parameters (h_o, b) of a P2P reciprocation protocol which can be sustained as a social norm equilibrium. As the proof shows, increasing the service threshold h_o provides the peers increased incentives to comply with the prescribed protocol. On the other hand, increasing the value of b reduces the peers' incentives to comply in general. Given the trade-off between the social efficiency and the incentive to comply with the prescribed protocol, the design problem now becomes selecting the smallest h_o and the largest b for which the incentive constraints in Problem (OSNE) are satisfied. A detailed algorithm for determining (h_o^*, b^*) can be found in [13].

IV. Protocol Designs for Networks with Altruistic and Malicious Peers

So far we have focused on "reciprocative" peers since they will provide services if recipients of the services are likely to return the favor [12]. In this section, we particularly consider the impact of altruistic peers and malicious peers on the design of optimal social norm equilibrium. Altruistic peers have the entire media file and provide upload services by playing $a = 1$ in response to any request it receives, regardless of the peer's reputation where the request comes from. Meanwhile, they do not request any chunk from others. We assume that the number of upload services that an altruistic peer can provide in one period is limited to a maximum number of λb services per period due to the bandwidth limitation. We assume that altruistic peers can be identified by the system and will be assigned a reputation L by the protocol constantly regardless of whether its upload is successful or not.

On the contrary, there are also malicious peers, whose goal is to cause damages to other peers and attack the system. The most common attacks include incomplete chunk attacks and pollution attacks [15]. In an incomplete chunk attack, a malicious peer agrees to send the entire requested chunk to its client, but sends only portions of it or no data at all. In a pollution attack, a malicious peer corrupts the media chunks, renders the content unreadable, and then makes this polluted content available for sharing with other peers. In both cases, the client of a malicious peer wastes its download connection and has to request the same chunk again in a separate transaction. Hence, a malicious peer is regarded to be playing $a = 0$ in any upload transaction it is engaged. Here we assume that malicious peers are treated by the protocol as regular reciprocative peers.

Let p_C, p_D, and p_R denote the fractions of altruistic, malicious and reciprocative peers, respectively, with $p_C + p_D + p_R = 1$. First, we analyze the impact of malicious peers by assuming $p_C = 0$. Similar to the previous sections, we are also interested in the long run stationary distribution of reputation, denoted as $\{\eta_D(\theta)\}_{\theta=0}^L$, with μ_D denoting the fraction of peers that can receive services according to the protocol. Let $\{\omega_D(\theta)\}_{\theta=0}^L$ denote the stationary distribution of the malicious peer population, it can be expressed as follows:

$$\begin{aligned}&\omega_D(\theta) = 0, \ h_o + 1 \leq \theta \leq L \\ &\omega_D(0) = \omega_D(h_o) \ \textit{and} \ \omega_D(\theta) = \omega_D(\theta-1), \ 1 \leq \theta \leq h_o\end{aligned}. \quad (15)$$

Let $\{\omega_R(\theta)\}_{\theta=0}^L$ denote the stationary distribution of reciprocative peers, and it can be computed using (6). Summing up, the stationary distribution can be solved as follows with $\mu_D = \sum_{\theta=h_o}^{L} \eta_D(\theta)$,

$$\eta_D(\theta) = (1-p_D)\omega_R(\theta) + p_D\omega_D(\theta), \ 0 \leq \theta \leq L. \quad (16)$$

We prove in the following proposition that the presence of malicious peers does not only decrease the social welfare, but also the incentives of reciprocative peers to comply with the protocol.

Proposition 2. Given a protocol $\kappa = (\sigma_o, \tau)$ and the fraction of malicious peers p_D, the reciprocative peers' incentive to comply with κ monotonically decrease with p_D.

Proof: See Appendix A. ■

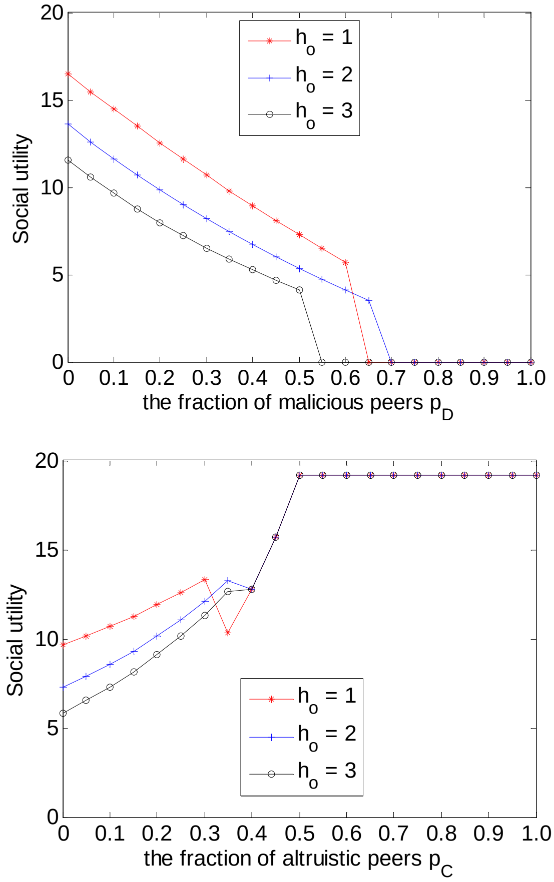

Figure 2. Average utility of reciprocative peers against p_D and p_C

Next, we investigate how altruistic peers impact the reciprocative peers' utilities and incentives with. Similarly, let $\{\eta_C(\theta)\}_{\theta=0}^{L}$ denote the corresponding stationary distribution and μ_C denote the fraction of peers that can receive services according to the protocol.

Since an altruistic peer is assigned a constant reputation of L by the system, $\{\eta_C(\theta)\}_{\theta=0}^{L}$ can be computed as follows with $\mu_C = \sum_{\theta=h_o}^{L} \eta_C(\theta)$,

$$\begin{aligned} \eta_C(\theta) &= (1-p_C)\omega_R(L) + p_C \\ \eta_C(\theta) &= (1-p_C)\omega_R(\theta),\ 0 \le \theta \le L \end{aligned}. \tag{17}$$

We show that p_C cannot be too large in order to sustain a reciprocative peer's incentive to comply with the protocol.

Proposition 3. Given a protocol $\kappa = (\sigma_o, \tau)$ and the fraction of altruistic peers p_C, κ can be sustained as a social norm equilibrium if and only if p_C is below certain threshold $\overline{p}_C \le 0.5$.

Proof: See [13]. ■

Proposition 3 provides an essential result: it is not always good to increase p_C in the network. Although having more altruistic peers allows more upload services in the network, they in turn harm a peer's incentive to comply with the protocol, which reduces the cooperative sharing behavior among reciprocative peers.

Figure 2. plots the average utility of reciprocative peers in the network against p_D and p_C. It shows that the utility monotonically decreases with p_D. When p_D increases above a certain threshold, peers lose their incentive to follow the protocol and the network collapses with the average utility falling to 0. Meanwhile, the utility does not monotonically increases with p_C, since reciprocative peers lose the incentive to comply with the protocol at certain point of p_C. However, as p_C approaches 1, the average utility finally reaches the optimal value $\lambda b(1-\varepsilon)r$ since all peers' download requests can be fully served by altruistic peers.

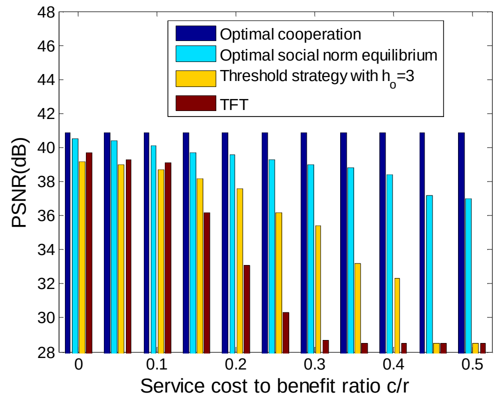

Figure 3. The PSNR of different protocols

Table 1. Decoded Video Quality

	Decoded Video Quality in PSNR (dB)			
	C1	C2	C3	C4
$h_o = 1$	39.2	38.2	37.1	36.6
$h_o = 2$	38.7	38.6	36.7	36.2
$h_o = 3$	38.4	36.9	36.9	36.5
TFT [14]	39.1	33.9	32.2	28.5

C1: $c/r = 0.1$ $\varepsilon = 0.1$; C2: $c/r = 0.3$ $\varepsilon = 0.1$; C3: $c/r = 0.1$ $\varepsilon = 0.3$; C4: $c/r = 0.3$ $\varepsilon = 0.3$.

V. ILLUSTRATIVE EXAMPLES

In this section, we illustrate the impact of the proposed social norm based protocols on P2P multimedia sharing networks using the simulator built [14]. A number of 200 peers are deployed in the network. All peers have the same download rate of 1Mbps. In each experiment, peers exchange a single video file of approximate size 100 Mbits, at CIF (352×528) resolution and 30 frames per second. The video is encoded using H.264/AVC codec and divided into chunks of 0.1s. All peers join the network at the same time. In the experiments, we deploy a reputation set $\Theta = \{0,1,2,3\}$, i.e. $L = 3$. We keep Θ and L fixed during the experiments.

We explicitly compare the average PSNR of the decoded video among all peers using different protocols. The exchanged video content is the well-known "Foreman" sequence repeated multiple times to create a long sequence. Besides the protocols studied in this paper, the performance of the Tit-for-Tat (TFT) protocol is also analyzed [14]. To make the TFT applicable to networks with random matching features, we slight change the protocol, which is defined as follows.

- The reputation set is binary as $\Theta = \{0,1\}$.
- The social strategy σ_{TFT} is defined as: $\sigma_{TFT}(\theta,\tilde{\theta}) = S$ if $\tilde{\theta} = 1$; $\sigma_{TFT}(\theta,\tilde{\theta}) = NS$ if $\tilde{\theta} = 0$.
- Using the same rule to calculate the statistic x as in Section II, $\tau(\theta, x) = 0,\ \forall\theta \in \Theta$. The reputation scheme is defined as: $\tau(\theta, x = 0) = 1$ and $\tau(\theta, x = 1) = 0$.

Table 1 presents the results given the parameters $b = 5$, $\lambda = 1$, and $\delta = 0.8$ for all peers. We also deploy 10% of altruistic peers in the network as seeds, i.e. $p_C = 0.1$. Various combinations of $(c/r, \varepsilon)$ are considered, with four strategies being implemented: the threshold-based strategy with $h_o = 1,2,3$ respectively, and TFT. When the service cost to benefit ratio and the service error are low, TFT delivers a PSNR which is comparable to those of threshold-based strategies. Meanwhile, as TFT only has two reputation levels and less peers being falsely punished, its PSNR is higher than those of threshold-based strategies with $h_o = 2,3$. On the other hand, our social norm based protocols are more robust by using threshold-based strategies, which deliver performances that are more insensitive to the variation on network conditions.

Figure 3. illustrates the advantage of optimal social norm equilibrium over fixed protocols which are selected in ad-hoc by explicitly comparing the performances of the following protocols

- Protocol 1: all peers cooperate unconditionally without considering the incentive constraints.
- Protocol 2: the optimal social norm equilibrium with h_o, b and p_C being optimized.
- Protocol 3: a fixed social norm with $h_o = 3$, $b = 5$, and $p_C = 0.3$.
- Protocol 4: TFT with $b = 5$ and $p_C = 0.3$.

Since all peers provide full services in Protocol 1, the performance it delivers remains to be constant and serves as the Pareto boundary of the performance that an incentive protocol can possibly achieve. Using this as a benchmark, Figure 3. shows that the optimal social norm equilibrium leads to significant improvements in terms of PSNR over Protocol 3 and 4, both of which adopt fixed strategies. As the PSNR delivered by the optimal social norm equilibrium remains roughly constant against the variation of c/r, the PSNR delivered by Protocol 3 and 4 drastically decrease with c/r. When c/r exceeds 0.25, the network adopting TFT collapses; while such collapse also happens in the network adopting Protocol 3 when c/r exceeds 0.45. In both cases, the reciprocative peers lose their incentive to follow the protocols and do not mutually provide upload services at all. Hence, there only exist minimum upload services in the network which are provided by the altruistic peers.

VI. CONCLUSION

We introduced a theoretical framework for analyzing and designing incentive protocols based on indirect reciprocity for P2P multimedia sharing applications. We designed optimal social norms which are sustainable and thus, under which no peer gains by deviating from the prescribed social strategy and thus have no incentive to deviate deliberately. We analyzed the structures of optimal incentive protocols and identified the trade-off between efficiency and incentives. We also discussed the impact of altruistic and malicious populations on the design and performance of optimal incentive protocols. Our simulation results verify that our social norm based protocol can deliver better performance than traditional incentive protocols.

APPENDIX A (PROOF OF PROPOSITION 2)

With simple calculation, an active reciprocative peer's expected one-period utility can be formalized as

$$v_\kappa(\theta) = \lambda b \frac{\mu_D - 1/(h+1)p_D}{\mu_D + h/(h+1)p_D}[(1-\varepsilon)r - c], \text{ for } \theta \ge h_o. \quad (18)$$

which monotonically decreases with p_D. Moreover, the expected overall utilities can be represented recursively as follows:

$$\begin{aligned} v_\kappa^\infty(\theta) &= v_\kappa(\theta) \\ &+ \delta[(1-\alpha)v_\kappa^\infty(\min\{L, \theta+1\}) + \alpha v_\kappa^\infty(0)] \end{aligned} \text{ for } \theta \ge h_o; \quad (19)$$

and

$$v_\kappa^\infty(\theta) = \delta v_\kappa^\infty(\min\{L, \theta+1\}), \text{ for } \theta < h_o. \quad (20)$$

Substituting (18) into (19) and (20), the incentive constraints for a protocol to be sustained as a social norm equilibrium can be written as

$$\delta(1-\alpha)[1-\delta^{h_o}] \frac{\lambda b \frac{[\mu_D - 1/(h+1)p_D]}{\mu_D + h/(h+1)p_D}[(1-\varepsilon)r - c]}{1 - \delta(1-\alpha) - \alpha\delta^{h_o+1}} \ge \lambda bc. \quad (21)$$

The left-hand side of (21) monotonically decrease with p_D. Hence, if a protocol can be sustained as a social norm equilibrium for some p_D, it can also be sustained for any $p_D' < p_D$. ■

REFERENCES

[1] X. Zhang, J. Liu, B. Li, and T. Yum, "CoolStreaming/DONet: A Data-driven Overlay Network for Efficient Live Media Streaming," *PROC. INFOCOM*, pp. 2102 – 2111, 2005.

[2] X. Jiang, Y. Dong, D. Xu, and B. Bhargava, "GnuStream: A P2P Media Streaming System Prototype," Proc. of 4th Int'l Conf. Multimedia and Expo., 2003.

[3] B. Cohen, "Incentives Building Robustness in BitTorrent," *Proc. P2P Econ. Workshop*, Berkeley, CA, 2003.

[4] O. Loginova, H. Lu, X. H. Wang, "Incentive Schemes in Peer-to-Peer Networks," *The B.E. Journal of Theoretical Economics*, Oct. 2008.

[5] P. Antoniadis, C. Courcoubetis, B. Strulo, "Comparing Economic Incentives in Peer-to-Peer Networks," *Computer Networks: The Int'l Journal of Computer and Telecommunications Networking*, vol. 46, No. 1, pp. 133-146, 2004.

[6] T. Ozbilgin, M. Sunay, "On the Capacity of Wireless Peer-to-Peer Networks with Transmitter and Receiver Cooperation," *Proc. IEEE Int'l Symp. on Personal, Indoor, and Mobile Radio Communications*, Athens, Greece, 2007.

[7] M. Feldman, K. Lai, I. Stoica, J. Chuang, "Robust Incentive Techniques for Peer-to-Peer Networks," *Proc. of the 5th ACM Conf. on Elec. Commerce*, Session 4, pp. 102-111, 2004.

[8] S. Kamvar, M. T. Schlosser, H. G. Molina, "The Eigentrust Algorithm for Reputation Management in P2P Networks," *Proc. 12th Int'l Conf. on World Wide Web*, pp. 640 – 651, 2003.

[9] M. Kandori, "Social Norms and Community Enforcement," *Rev. Economic Studies*, vol. 59, No. 1, pp. 63-80, Jan. 1992.

[10] Y. Zhou, D. M Chiu, J. Lui, "A Simple Model for Analyzing P2P Streaming Protocols," IEEE Int'l Conf. on Networking Protocols, pp. 226 – 235, 2007.

[11] L. Massoulie, M. Vojnovic, "Coupon Replication Systems," *IEEE/ACM Trans. on Networking*, vol. 16, no. 3, pp. 603 – 616, 2005.

[12] M. Nowak, K. Sigmund, "The Dynamics of Indirect Reciprocity," *Journal of Theoretical Biology*, vol. 194, no. 4, pp. 561 – 574, 1998.

[13] Y. Zhang, M. van der Schaar, "Peer-to-Peer Multimedia Sharing based on Social Norms," arXiv:1102.1503v1 [cs.MM]

[14] R. Izhak-Ratzin, H. Park, and M. van der Schaar, "Reinforcement Leraning in BitTorrent Systems," *PROC. INFOCOM 2011*, accepted and to appear.

[15] W. Lin, H. Zhao, and K. J. Ray Liu, "Incentive Cooperation Strategies for Peer-to-Peer Live Multimedia Streaming Social Networks," *IEEE Transaction on Multimedia*, vol. 11, no. 3, pp. 396 – 412, 2009

[16] F. Fu, M. van der Schaar, "Structural Solutions to Dynamic Scheduling for Multimedia Transmission in Unknown Wireless Environments," *UCLA Technical Report*, http://medianetlab.ee.ucla.edu/papers/TechReport_2010_fwfu_mvanderschaar.pdf