A Case for Taxation in Peer-to-Peer Streaming Broadcast *

Yang-hua Chu
Carnegie Mellon University
yhchu@cs.cmu.edu

John Chuang
UC Berkeley
chuang@sims.berkeley.edu

Hui Zhang
Carnegie Mellon University
hzhang@cs.cmu.edu

ABSTRACT

Most existing research on peer-to-peer (p2p) has been on file sharing applications. In this paper, we focus on p2p streaming applications. In particular, we argue that the *Bit-for-Bit* model, widely adopted in p2p file sharing, is not applicable in p2p streaming. In p2p streaming, the bottleneck resource is the upstream bandwidth capacity. Our empirical experience with p2p streaming indicates that a large percent of peers on the Internet have limited upstream bandwidth capacity, and the Bit-for-Bit model severely limits the amount of bandwidth these resource-poor peers can receive. To address this issue, we propose a *taxation* model.

In the taxation model, resource-rich peers contribute more bandwidth to the system, and subsidize for the resource-poor peers. This redistribution of wealth improves social welfare. Such a model is applicable in the streaming context because the publisher of the video stream has the *means* to enforce taxation on peers and the *will* to maximize their collective social welfare. We design a simple linear taxation scheme and incorporate it in a distributed streaming protocol. Our simulation results indicate that taxation can significantly improve social welfare without incurring a significant overhead to the system.

Categories and Subject Descriptors: C.2 [Computer Systems Organization]: Computer Communication Networks

General Terms: Performance, Design, Economics

Keywords: Peer-to-peer, Video Streaming, Incentive, Taxation

*This research was sponsored by DARPA under contract number F30602-99-1-0518, and by NSF under grant numbers Career Award NCR-9624979 ANI-9730105, ANI-9814929, ITR Awards ANI-0085920 and ANI-0085879. Additional support was provided by Sloan Research Fellowship and Intel. Views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of DARPA, NSF, Intel, or the U.S. government.

1. INTRODUCTION

Although many incentive mechanisms have been proposed in the peer-to-peer (p2p) file sharing context, there is little understanding on how to incorporate incentives into p2p streaming systems. While audio/video dissemination represents an important class of Internet applications, no satisfactory solutions exist. The high cost of bandwidth required for server-based solutions or content delivery networks, and the fundamental limitations of IP Multicast are two main factors that have limited broadcasting to a very small number of rich publishers. P2P architecture presents an interesting alternative in which participants in the broadcast also contribute forwarding capacity (upstream bandwidth), thereby significantly increasing the scalability of the system while lowering the cost incurred by the publisher. The ability for users to receive content that they would otherwise not have access to provides a natural incentive for them to contribute resources to the system.

As a proof of concept, we have implemented and deployed a p2p streaming system [10]. The system has been in operation for the past 18 months. It has been used to successfully broadcast many events including academic conferences such as ACM SIGCOMM and SOSP, DARPA Unmanned Vehicle Grand Challenge, CMU Commencement, and CMU SCS Distinguished Lectures. These broadcasts have reached thousands of users in residential, academic and commercial environments, spread over four continents [5]. Our experience demonstrates the great potential of using p2p architecture to provide cost-effective video streaming over the Internet. In our current system, we maximize the overall system performance by utilizing as much forwarding capacity as possible from the peers. In this paper, we consider the scenario in which all peers are *strategic* agents, i.e., they contribute more resources only if they see clear benefits in doing so.

One natural candidate is the *Bit-for-Bit* scheme, where peers receive only as much as they contribute. This scheme has been adopted by protocol designers [3, 18, 7] for its simplicity and apparent fairness. However, our deployment experience predicts that it will perform poorly in today's Internet environment.

In a p2p environment for streaming, bandwidth is the bottleneck resource. In our experience, a reasonable quality audio/video stream takes at least 350Kbps to encode. However, many of the peers are behind DSL and cable modems. These resource-poor peers have asymmetrical bandwidth capacity with low forwarding capacity (100-200Kbps) and high receiving capacity (600-1200Kbps). In our broadcasts, up to

80% of the peers are resource-poor. If we adopted the *Bit-for-Bit* model, these resource-poor peers would only receive lower video bitrate (100-200Kbps), even though they have enough capacity to receive a much higher bitrate. The net result is that these resource-poor peers would *not* participate in the broadcast due to the poor quality.

To accommodate resource-poor peers, who would otherwise be unable to participate, we would like to design a system that incentivizes resource-rich peers to contribute more bandwidth and *subsidize* for the the resource-poor peers. If we model the system as a game, and map the bandwidth resource to the *wealth* of peers, then the objective of the game is to achieve an efficient and equitable redistribution of wealth. We believe *taxation* offers a compelling approach to this design objective.

Taxation has one important precondition. There must exist an *asymmetry of roles*, where one entity (role) is empowered to enforce tax payment on individuals according to a predefined tax schedule. We believe p2p streaming is in an unique position to satisfy this precondition. In p2p streaming, the *publisher* of the video stream is the natural empowered entity. The publisher owns the content and can choose the *means* in which peers participate in the system (via proprietary software). In other words, the publisher can freely design a game and enforce the rules of the game. Peers participating in this game are strategic. They individually own the bandwidth resources and are strategic in minimizing the cost of contributing resource while maximizing the benefit of the video quality received.

In this paper, we model a publisher who is interested in maximizing the total welfare of all peers (*social welfare*). To achieve this objective, the publisher specify a taxation game where peers contribute and receive bandwidth according to a tax schedule. The taxation scheme facilitates a subsidy from the resource-rich peers to the resource-poor peers, improving social welfare. We believe a welfare-maximizing publisher is a reasonable model. By improving the performance of the resource-poor peers (80% in one event), these peers have more incentive to participate. For publisher, participation translates to visibility, advertising revenue, and market share.

We note that taxation provides a *direct* mapping between contribution and benefit, in contrast to other incentive mechanisms based on currencies [24, 22, 12, 8] and/or reputation [2, 15, 13], which provide *indirect* mappings between contribution and benefit. The indirection is necessary if a peer's contribution and consumption are temporally separated. However, this is not the case for p2p streaming. Therefore, the adoption of a taxation scheme avoids the overhead and security vulnerabilities of maintaining persistent state (e.g., tokens).

We study the performance and implementability of standard taxation schemes from the public finance literature. We find that a linear tax schedule, with a single marginal tax rate and a demogrant, provides significant social welfare improvements over the Bit-for-Bit scheme, especially with heterogeneous populations. Furthermore, the taxation scheme is implementable. By employing techniques such as multiple description codec (MDC), priority, and preemption, the scheme works well under dynamic peer environments. Evaluations show that the taxation scheme achieves high efficiency and high compliance without incurring significant overhead.

2. MODEL OF P2P STREAMING

There are two entities in the p2p streaming system we consider: publisher and peers. A publisher makes a live video stream available on the Internet, and peers who are interested in the stream join the p2p system. Peers construct an overlay structure in a distributed fashion and disseminate the video stream along the overlay. Both entities have mutual incentive to use the p2p system. By delegating the task of data forwarding to the peers, publishers can avoid the costly bandwidth provisioning to support a large number of viewers. The peers have an incentive to help the publisher in exchange for enjoying the video that might otherwise be unavailable.

In this scenario, bandwidth is the valuable resource. High quality real-time video streaming requires the availability of high bandwidth (at least hundreds of Kbps) that is persistent over time. Thus we consider the peers' bandwidth capacity as their "wealth" for taxation.

We model the bandwidth capacity of a peer i with two parameters: forward capacity (F_i) and receive capacity (R_i). F_i and R_i represent the upper bound bandwidth that a peer can contribute to and receive from the p2p system, respectively. We do not model congestion in the core of the network, as congestion happens mostly at the access links on the Internet today. Therefore, a peer can forward traffic to another peer as long the capacity bounds on both ends are met. A peer may choose to utilize only a portion of its capacity. We denote the actual bandwidth a peer contributes and receives as f_i and r_i.

A peer gains benefit b_i when it receives bandwidth from the broadcast system, and incurs cost c_i when it contributes bandwidth to the system. When faced with a taxation scheme that specifies the contributed bandwidth f_i as a function of the received bandwidth r_i, a strategic peer will choose an optimal r_i to maximize its utility u_i, subject to $r_i \leq R_i$ and $f_i \leq F_i$.

$$u_i(r_i, f_i, F_i) = b_i(r_i) - c_i(f_i, F_i) \tag{1}$$

The benefit function (b_i) captures the user-perceived video quality. The function should be *concave* or S-shaped in the receive bandwidth (r_i) to capture the diminishing returns of increased video bitrate on the perceived video quality. The benefit function should be indepedent of a peer's forward capacity (F_i) to reflect the equal desire of viewers to watch a high quality video stream.

The cost function (c_i) captures the cost of forwarding data. The function should be concave in F_i to capture economies of scale in bandwidth but convex in f_i (or more specifically f_i/F_i) to capture the effects of link congestion. We will consider specific functional forms for user benefits and costs in Section 5.

Finally, we model a publisher who is interested in maximizing the social welfare of the system, which is simply the summation of individual utilities of the peers. The publisher may also choose a tax schedule to maximize some other objective function, e.g., system throughput, demogrant.

3. PROPOSED TAXATION SCHEME

In this section, we show how to incorporate taxation into a p2p streaming system. We first construct a suitable taxation model for p2p streaming in Section 3.1 based on the public finance literature. The main departure from the traditional taxation is that p2p streaming is less tolerant to a budget deficit. In Sections 3.2 and 3.3, we propose a simple taxation schedule that we use throughout the paper. This schedule is based on linear taxation, which is widely studied in the optimal income taxation literature. We modify the semantic to ensure a budget balance. Finally in Section 3.4, we show how the publisher sets the tax schedule in practice.

3.1 Model Taxation in P2P Streaming

For peer i, the taxable income is r_i and the tax payment is f_i. In other words, a tax schedule specifies how much bandwidth a peer must contribute (f units) in order to receive r unit of bandwidth. In the following we list a set of requirements for taxation in the public finance literature [21].

• *Asymmetry of roles and power*: There must exist an entity empowered to set and collect taxes from the individuals. In p2p streaming, the publisher is the natural entity. The publisher owns the broadcast rights of the video stream, and has control over the means (the software) in which the video is distributed. The peers are assumed to be strategic. Each peer chooses its optimal contribution level f_i, and receive a corresponding amount r_i determined by the published tax schedule, in order to maximize its utility. Peers are also assumed to satisfy the *participation constraint*, i.e., their utility from participation and adherence to the taxation scheme exceeds a reservation utility derived from not doing so (e.g., not participating in the system, or modifying the software to engage in tax evasion).

• *Public and fixed tax schedule*: A tax schedule should be fixed and made public such that peers can adjust their strategy to maximize their utility. The tax schedule should not change (or should change at a very large time scale) to minimize system instability due to peers reacting to the changes in tax schedule.

• *Fair*: There are two types of fairness: horizontal and vertical. Horizontal fairness requires that individuals with similar wealth should bear similar tax liability. We adopt the same requirement in p2p streaming. Vertical fairness requires that individuals with different wealth should bear (potentially) different tax liability. In public finance, vertical fairness is more a matter of public opinion, where taxpayers can influence the tax schedule (through voting). In p2p streaming, the tax schedule is determined solely by the publisher. In this paper, the publisher aims to maximize social welfare as described in Section 5.1.

• *Budget balanced*: Budget is not balanced (budget deficit) if the budget expenditure exceeds the revenue from taxation. In p2p streaming, this means that $\sum f_i$ (budget expenditure) must be greater or equal to $\sum r_i$ (tax revenue). This is intuitive because every byte of bandwidth received by a peer must be contributed by another peer. The requirement for budget balance is more stringent in p2p streaming than in public finance, because taxed money collected in public finance is a *persistent* resource but the taxed bandwidth collected in live p2p streaming is perishable. Thus in p2p streaming, it is not possible to “store” the bandwidth resource and “use” it later on.

3.2 Linear Tax Schedule

We choose a linear tax schedule, which takes on two parameters: (i) t, *marginal tax rate*, and (ii) G, lump sum grant, also known as *demogrant*. Note that if a peer does not contribute any bandwidth ($f = 0$, or free ride), it would still receive a demogrant ($r = G$). The publishers sets only the first parameter. The second parameter is dynamically inferred from the peer environment to achieve budget balance. This is a departure from the traditional literature in linear taxation where both parameters are configured simultaneously.

$$f = max(t * (r - G), 0) \quad (2)$$

The tax rate (t) must be at least 1, otherwise fundamentally the budget cannot be balanced even if the demogrant is 0. When $t = 1$, the tax schedule becomes *Bit-for-Bit* and $G = 0$. This is because when $f_i = r_i$, there is no extra tax expenditure for demogrant. When $t > 1$, the demogrant may be greater than 0. If a peer contributes more than it receives ($f_i > r_i$), the bandwidth difference goes to a *demogrant pool*. This pool of bandwidth is then evenly distributed among all peers as demogrant.

Linear taxation has been widely studied in the optimal income taxation literature. Despite its simplicity over nonlinear taxes, linear taxes provide surprisingly robust results in many settings [16, 20]. Our evaluation results in Section 5.4 are aligned with this observation. We note that there are many other taxation schedules and budget balance strategies, and defer their investigation to future work. Our goal in this paper is to demonstrate the existence of one taxation scheme that is effective and implementable in a p2p environment.

3.3 Budget Balance Strategy

Given a fixed tax rate, a fixed demogrant, and a fixed peer environment, the tax budget (tax revenue minus the expenditure) can be in one of the three conditions: (i) balance, (ii) surplus, or (iii) deficit. Ideally we want the budget be balanced. Given the taxed bandwidth is a perishable resource, a surplus wastes resource (and hurt social welfare) and a deficit makes the system infeasible to operate.

To deal with a highly dynamic peer environment, we tune the demogrant to ensure budget balance. Deriving the demogrant value in a centralized manner takes several rounds. In each round j, the planner announces the demogrant value (G_j) to all peers. Each peer determines the best (utility-maximizing) strategy and tells the publisher its f_i. The planner then finds the highest demogrant (G_{j+1}) without causing a budget deficit. The algorithm stops when $G_{j+1} = G_j$. In Section 4, we show how to derive demogrant dynamically in a distributed protocol.

3.4 Setting the Tax Schedule

How does the publisher set the tax schedule? In theory, the publisher can set the optimal tax schedule once it knows the distribution of user types. In many distributed system settings, user types are private information and users may not truthfully reveal their types to the system. A possible response is to design *strategyproof* mechanisms to induce truthful revelation by the users [11, 17]. A publisher can design an incentive compatible tax schedule such that user types can be inferred from user action.

In the case of p2p streaming, the user types are their bandwidth capacities (F and R), and the user actions are the actual amounts of forwarding and received bandwidths (f and r). Bandwidth capacities are static host characteristics

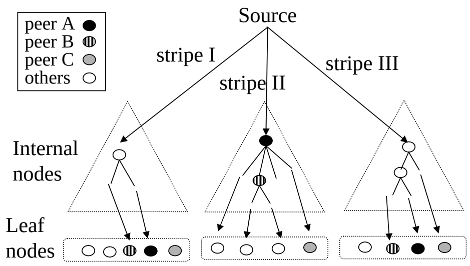

Figure 1: An example of a multiple disjoint tree structure.

	f	r^+	1st stripe	2nd stripe	3rd stripe
Peer A	4	2	interior node (0)	entitled leaf (0)	demogrant leaf (1)
Peer B	2	1	interior node (0)	demogrant leaf (1)	demogrant leaf (2)
Peer C	0	0	demogrant leaf (1)	demogrant leaf (2)	demogrant leaf (3)

Figure 2: An example depicting three peers with different forward bandwidth. Peers join each tree either as an interior node or a leaf. The numbers in bracket are their priority in receiving nth stripe. The shaded blocks are their entitled bandwidth.

that can be easily determined by the software agent running on the peer host, so strategyproofness is not a major concern. More interestingly, we find that accurate knowledge of user type distribution may actually be unnecessary in practice. In Section 5.4 we show that a fixed linear tax schedule is surprisingly robust against changes in type distributions, and good social welfare outcomes can be obtained for a wide range of values for the marginal tax rate t.

4. INCORPORATE TAXATION INTO A DISTRIBUTED PROTOCOL

This section shows how a linear taxation scheme can be incorporated in a distributed protocol for video broadcast. We extend the protocol in [6] to express the linear taxation semantic and allow peers to change strategy dynamically. Both [6] and this protocol build a multiple disjoint tree structure, which was first proposed in [3, 18].

4.1 Multiple Tree Protocol

In the original multiple tree proposal, the source splits the video stream into m stripes using MDC, and multicasts each stripe along a separate tree. Each peer selects *one tree* at random, and joins the trees as an *interior node*. It joins all other tree as *leaf nodes*. Figure 1 illustrates an example structure that a multiple tree protocol constructs. Note that a peer contribute bandwidth to one of the m trees, and receive bandwidth from all the trees it joins, including the one it contributes bandwidth to.

We use MDC for a different purpose than originally proposed. MDC helps to support receiver heterogeneity, improve resilience to machine failure and network congestion. In taxation, MDC offers peers the flexibility to contribute and receive small increments of bandwidth. Such flexibility helps us express the taxation schedule in finer granularity.

Now we attempt to incorporate the taxation semantic into this protocol and motivate the need to extend the protocol. To forward f unit bandwidth, a peer would configure the fanout of the interior node to be f. To receive r unit bandwidth, the peer would join r trees. The issue is that r depends on the available demogrant (G), and it requires global information to infer G. In the following, we describe the technique to infer G in a distributed fashion.

4.2 Distributed Bandwidth Allocation

Conceptually, the distributed protocol incrementally allocates r to the peers (by increasing G) until the budget is balanced. First, the heuristic assumes G is 0, and allocates the receive bandwidth to only the peers who pay tax. For peer i, this amount is equal to f_i/t. We call this the *entitled bandwidth* (r^+). A peer is entitled to receive r^+ even in a worse-case environment ($G = 0$). After all peers receive their share of r^+, the leftover bandwidth in the system is the demogrant pool. Next, the protocol iteratively increase G by 1, until the demogrant pool is exhausted. With every unit increment of G, all peers are allowed to join one more tree.

We leverage two techniques, *priority* and *preemption*, to facilitate this join order. These techniques are first described in [6] and we provide a short description below. Each peer assigns a *priority* value for each of the m trees it joins. The peer marks the first r^+ joins with the highest priority. Then the peer iteratively marks all other joins with decreasing priority. Table 2 shows an example of priority assignment for the three peers. To receive the entitled bandwidth of 2, A sets a high priority (priority = 0) to join the first 2 trees, in which one is an interior node and the other is an leaf node. A has a lower priority (priority = 1) to join the third tree until the demogrant becomes 1.

Now that the nodes in each tree has proper priority values that reflect the tax schedule, the protocol needs to perform admission control based on the priority. Nodes that are accepted in the tree should have higher priority than those rejected. To achieve this, each interior node individually runs a *preemption* rule on the joining peers. If the fanout bound is reached, a peer with a higher priority can preempt existing peers with lower priority. With dynamic peer environment, a peer that was previously preempted may become eligible. Thus a preempted peer periodically (every 30 seconds) attempt to rejoin the tree and get its fair share of the demogrant.

The peer may change its strategy (by changing f) depending in part by the available demogrant. Each peer can passively estimate G by counting the number of trees it joins. However, this estimate is not reliable due to the transient condition in the distributed protocol. To increase the accuracy of the estimate, each peer periodically (every 30 seconds) queries a subset of other peers (20 peers) about their estimates of G, and merge with its own estimate.

5. EVALUATION

Our evaluation seeks to answer the following three questions:

• Does taxation yield good social welfare outcome under various peer environment? We compare our proposed linear taxation scheme with two benchmark schemes. The lower bound benchmark is a Bit-for-Bit scheme, and the upper bound benchmark is a socially optimal scheme where peers are obedient (or altruistic).

• Is the proposed tax schedule effective in maximizing tax schedule

f	r
0	1
2	2
4	3

outcome

	F	R	f	r	U
Peer A	10	3+	4	3	1.45
Peer B	4	3+	2	2	1.16
Peer C	1	3+	0	1	1.0

(a) Linear taxation with t=2.0, G=1: **S=1.20**

tax schedule

f	r
0	0
1	1
2	2
3	3

outcome

	F	R	f	r	U
Peer A	10	3+	3	3	1.54
Peer B	4	3+	2	2	1.16
Peer C	1	3+	1	1	0.25

(b) Bit-for-Bit with t=1.0, G=0: **S=0.98**

Figure 3: An example illustrating the two tax schedules and their impact on the strategy and utility of the three peers. In this example where peers are heterogeneous, taxation (with $t = 2.0$) has higher social welfare than Bit-for-Bit.

social welfare? We compare our fixed linear tax schedule with dynamic and non-linear tax schedules.

• What is the performance implication when incorporating taxation in a distributed streaming protocol? We quantify the performance implications using three metrics: utilization, compliance, and stability.

We seek answers to these questions using two different simulation setups. The first two questions are about the fundamental efficacy of taxation, and we conduct (*static*) simulation with a fixed group of peers using a centralized algorithm. The last question is about implementation feasibility, and we conduct trace simulation using the proposed distributed protocol. This section first presents the simulation setup and the utility functions used in the evaluation, and then presents the simulation results to answer the three questions in turn.

5.1 Utility Functions and Example

To quantify the potential benefit of taxation, we consider a simple set of utility functions below. We acknowledge that the utility functions are based on intuition. However, we do believe the shape of the curve (convex vs. concave) is accurate. We are in the process of collecting measurements from the real p2p environment to refine the utility functions and their parameters.

We approximate the benefit function (b_i) as a square root function to the received bandwidth (r_i). This concave function represents the diminishing benefit for the perceived video quality as the bitrate increases.

$$b_i(r_i) = \sqrt{r_i} \tag{3}$$

The cost function (c_i) captures the cost of forwarding data, as shown in Equation 4. The forwarding cost (f_i) is modeled as a fraction (p_i) of the (dollar) cost in purchasing the access bandwidth F_i. With economy of scale, the cost of F_i is concave, and is modeled as a square root of F_i. We added the α parameter to calibrate the cost function with the benefit function. α should be less than 1, indicating the desire to view the video content. In the evaluation, α is set to 0.75.

$$c_i(f_i, F_i) = \alpha * \sqrt{F_i} * p_i(f_i, F_i) \tag{4}$$

p_i is the fractional cost to use f_i given F_i, and the value is between 0 and 1. It is modeled as a weighed average (β) between the two components, as shown in Equation 5. The first component models the direct forwarding cost, where the cost of forwarding f_i is linear to the cost of F_i. The second component models the congestion cost. Congestion happens when f_i is close to F_i. Since the access links on the Internet are typically shared by different users or applications, link congestion will affect the performance (utility) of other users and applications. We model this effect as the fourth power of the linear fraction. The β parameter is set to 0.5, where both components have equal weight in the cost function.

$$p_i(f_i, F_i) = \beta * (\frac{f_i}{F_i}) + (1 - \beta) * (\frac{f_i}{F_i})^4 \tag{5}$$

We now provide a concrete example in Figure 3 to illustrate that taxation has the potential to improve social welfare over Bit-for-Bit. In this example, there are 3 peers with heterogeneous bandwidth capacities. The maximum source rate is 3 bandwidth units. Figure 3(a) shows a tax schedule with a marginal tax rate of 2.0 and a demogrant of 1 bandwidth unit. The outcome of the three peers are shown in the adjacent table. The outcome includes the bandwidth capacity (F, R), the chosen strategy (f, r), and the marginal utility (U). For example, A is a resource-rich peer, which can contribute up to 10 units of bandwidth. It is incentivized to contribute 4 units, in order to receive the full source rate of 3 units. This translates to a utility of 1.45. Figure 3(b) shows the Bit-for-Bit tax schedule and the corresponding outcome of the three peers.

Although both tax schedules collect the same tax revenue ($\sum f_i = 6$), taxation provides better social welfare than Bit-for-Bit. This improvement can be explained from two inter-related angles: (i) The cost of raising the same tax revenue is reduced with taxation. The reduction comes from a shift of tax liability from the poor (C) to the rich (A). A has lower marginal cost (1.54-1.45=0.09) of contributing one additional bandwidth unit compared to C (0.75). (ii) There is a benefit to redistributing the tax expenditure. Node C receives a higher marginal benefit from receiving the additional bandwidth unit than node A. The use of demogrant facilitates this redistribution.

5.2 Evaluation Environment

To realistically model the simulation environment, we use data and traces collected from several live events using a p2p broadcast system [4]. Due to space constraints, we show the results of one trace (Slashdot) in detail. The Slashdot event is the largest among all the traces, and attracted 1316 peers. The mean and median stay time is 18 minutes and 3 minutes, respectively. The trace lasts for 8 hours. The peak group size is 160. The trace contains the TCP throughput measurements (upstream and downstream) between each peer and a well-provisioned server (in our university). We use the bandwidth measurements to model the forward and receive capacity of the peers (F_i and R_i).

Static Simulation: We conduct the first two parts of the

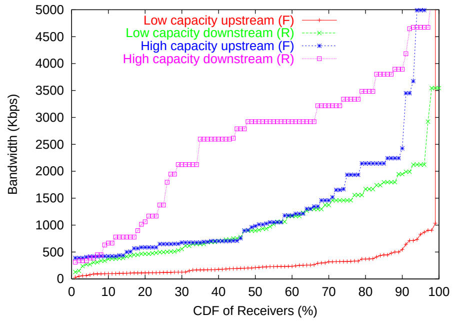

Figure 4: Measured TCP throughput of peers in Slashdot. Peers are categorized tinto low and high capacity, and the distribution is used to model environment heterogeneity in the simulation.

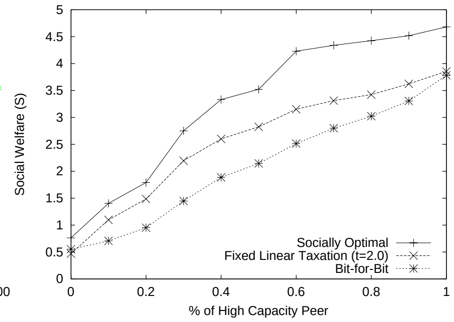

Figure 5: Social welfare as a function environment heterogeneity for the taxation scheme and the two benchmark schemes. Taxation has a social welfare outcome in between the two benchmarks.

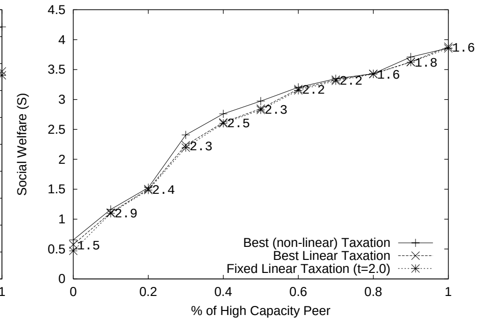

Figure 6: Social welfare as a function of environment heterogeneity for three different tax schedules. The proposed fixed linear tax schedule is surprisingly robust against environment changes.

evaluation with a fixed group size of 100 peers. To simulate a wide range of peer environments, we systematically vary the bandwidth capacity of the 100 peers. We consider two categories of peers, high capacity and low capacity, and model a range of peer environment by varying the *composition* of these peers. Thus, a homogeneous peer environment contains either 100% high capacity peers or 100% low capacity peers. To assign the bandwidth capacity of a peer in a category, we draw from a distribution derived from the Slashdot trace. To derive this distribution, we first categorize each peer in the Slashdot trace using its DNS name and other access bandwidth measurements. A peer is categorized as low capacity if it is behind DSL, cable modem, or the access measurement is below T1. A peer is categorized as high capacity otherwise. Then, the bandwidth distribution is derived from all the peers under the same category. We note that in the Slashdot trace, about 20% of the peers are high capacity and 80% are low capacity. Figure 4 shows a CDF of peer TCP upstream (F_i) and downstream (R_i) bandwidth. We observe significant heterogeneity among peers, and the upstream bandwidth is significantly lower than the downstream bandwidth.

Trace Simulation: We conduct the last part of the evaluation by playing back the peer join and leave sequence in the Slashdot trace. The peers are assigned the bandwidth capacity as recorded in the trace. The simulator captures the overlay tree changes due to peers joining and leaving the group, but not due to network congestion. So peers would switch parents only if they are preempted by other peers, or if their parents leave the group. To find a parent, a peer probes a small number of other peers (up to 5) that are the interior nodes in the tree. This limit bounds the overhead in maintaining each tree.

Common Parameters: Both simulators assume uniform delay between any pair of peers and no packet loss. The maximum streaming rate is 1600Kbps. The stream is evenly divided among 32 stripes using MDC (so there are 32 trees). So each stripe (bandwidth unit) is 50Kbps.

5.3 Social Welfare of Taxation

We expect that taxation should improve social welfare compared to Bit-for-Bit, but it will not be *socially optimal*. We define a socially optimal scheme where peers are obedient in contributing as much bandwidth as dictated by the system designer. The system designer can then decide how much bandwidth each peer should contribute and receive to maximize overall social welfare. We use the following algorithm to derive a socially optimal outcome given a fixed peer environment:

Socially Optimal Scheme: The algorithm iterates on the total amount of forward bandwidth (called W) in the system. Initially W is 0, and is incremented by 1 (unit bandwidth) until peers contribute all of their forward capacity. In each iteration, the algorithm minimizes the aggregate cost of raising W from the peers (called C_W) and maximize the aggregate benefit of using W among the peers (called B_W). The social welfare of W (S_W) is B_W - C_W. The socially optimal outcome then is the bandwidth distribution of W such that S_W is the highest. To get B_W, the algorithm allocates W evenly among peers (because the benefit curve is concave). To get C_W, the algorithm incrementally raise bandwidth from the peer who has the lowest marginal cost.

Figure 5 shows the social welfare of the three schemes under various peer environments. 20% on the x-axis means the peer environment is composed of 20% of high capacity peers and 80% low capacity peers. Each curve represents one scheme. For the linear taxation scheme, we choose a tax rate of 2.0. This is an optimized choice which will become clear in Section 5.4. We make two observations. (i) Taxation is a *strongly dominating* strategy over Bit-for-Bit when peers are *heterogeneous*. For example, when the composition ratio is 20% (Slashdot environment), social welfare improves from 1 to 1.5, a 50% improvement. (ii) Taxation is a *weakly dominating* strategy when peers are *homogeneous*. This is because with similar forwarding capacity, peers will elect similar strategy in selecting f and r, and the degree of redistribution becomes minimal. (iii) The taxation scheme is still considerably worse than the socially optimal scheme. This is expected because the socially optimal scheme demands a large contribution from the high capacity peers (because the per-unit cost is lower) but allocate the same r as all other peers. Such a behavior is not possible to enforce in a taxation scheme where peers are strategic.

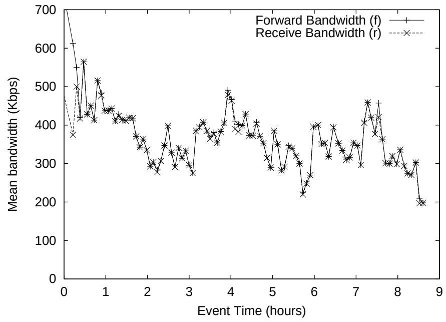

Figure 7: Mean bandwidth of peers vs. time in the Slashdot trace simulation.

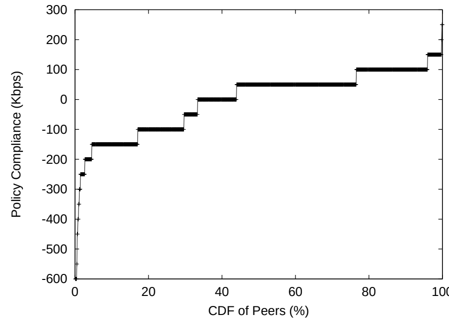

Figure 8: CDF of policy compliance for the peers in the experiment.

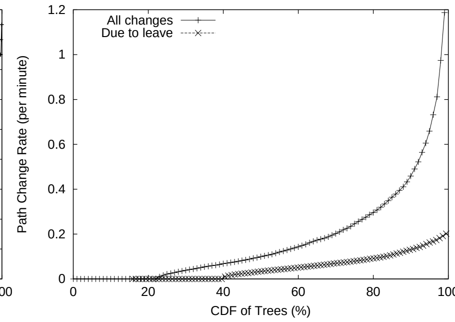

Figure 9: CDF of path change rate for all the trees peers participate.

5.4 Effectiveness of Linear Taxation

The linear taxation schedule we propose is (i) linear and (ii) fixed (i.e. t is constant across time and event). However, this choice of tax schedule may not be effective in maximizing social welfare. In this part of the evaluation, we quantify the potential penalty of this design choice. We compare the proposed tax schedule with the following two schemes:

Best Linear Tax Scheme: In this scheme, the tax schedule is linear but the rate can be dynamically adjusted to optimize S. To find the best tax rate (t_b), the scheme varies the tax rate from 1.0 to 4.0 with an increment of 0.1 (41 possible rates), and selects the rate that maximizes S.

Best (Non-Linear) Taxation Scheme Here the tax schedule can be non-linear *and* the rate can be adjusted dynamically. This scheme represents the ideal tax schedule that maximizes S. Unfortunately, we are unable to come up with an algorithm to find this ideal tax schedule directly. A brute force solution seems infeasible. With a source rate of 32 unit and 41 possible tax rates, the number of possible tax schedule is as large as 32^{41}. To approximate the ideal tax schedule, the heuristic performs 20 rounds of hill climbing. At the beginning of each round, the heuristic chooses a random tax schedule. Then it iteratively adjust the 32 entries in the tax schedule, until no single adjustment can yield a higher S. Finally, the heuristic picks the tax schedule that yields the highest S among the rounds.

Figure 6 shows the social welfare of the three schemes under various peer environments. Each curve corresponds to one scheme. The lowest curve is the fixed linear taxation with a rate of 2.0. The middle curve is the best linear tax schedule, with the rate numbered for each environment. As expected, a tax schedule is more effective when it is non-linear and dynamic. However, the difference is not significant, as the three curves are very close to each other. This is an indication that a *fixed* and *linear* taxation scheme is effective under a variety of peer environments, and there is marginal benefit when tuning this parameter specific to a peer environment. Our further analysis indicates that the 2.0 linear tax rate is not a magic number. In fact, a tax rate between 1.5 and 2.5 yields similar results.

5.5 Distributed Protocol Performance

The distributed taxation protocol (with a tax rate of 2.0) is evaluated with the following three performance metrics:

Utilization: Ideally, an efficient protocol should utilize all the bandwidth peers contribute to the overlay. A distributed protocol cannot be as efficient because it takes some time for peers to (i) find unsaturated trees and parents, or (ii) find parents who preempt other children because the joining peers have higher priority. Figure 7 shows the mean bandwidth metric of peers as a function of experiment time. The top curve is the mean forward bandwidth (f) and the bottom curve is the mean receive bandwidth (r). The closer the two curves are, the higher the efficiency is. Our protocol is quite efficient, which utilizes at least 95% of the contributed bandwidth most of the time.

Compliance: A good protocol should not only utilize the contributed bandwidth efficiently, but also allocate the right amount of bandwidth to peers in compliance with the taxation policy. To measure compliance, at every time interval (5 seconds) we compare the difference between the received bandwidth allocated by our protocol and the bandwidth the peer should ideally receive in a centralized setting. Figure 8 shows a cumulative distribution of that difference for all peers during the experiment. A negative difference means the protocol allocates less bandwidth than it is supposed to. We observe that a majority of time (80%), peers see a difference no more than 100Kbps, indicating the protocol is fair in allocating bandwidth to peers.

Stability: Since peers join, leave, and change their strategy dynamically (by changing f), the protocol must also dynamically adjust the bandwidth allocation (r) to the peers. Our protocol achieves high compliance and high utilization through preemption. However, preemption incurs a cost in performance. Peers that are preempted may experience transient data loss before finding another parent. Worse yet, if interior nodes are preempted, their descendants are also affected. To capture this cost, we count the rate of *path changes* for each tree that peers participate during the experiment. Some of the path changes are fundamentally unavoidable. Specifically, the path to a peer will change if any one of their ancestors leaves the group. Figure 9 shows the cumulative distribution of path change rate for each tree that peers participate in the experiment. We note that the total cost of implementing the policy is about twice the fundamental cost in maintaining the structure. This is shown as the difference between the two curves.

6. RELATED WORK

Incentive Mechanisms: To our knowledge, this is among the first work to incorporate the concept of taxation into an incentive mechanism for p2p systems. Most of the prior work adopts either a token-based or a reputation-based incentive mechanism. In token-based schemes, users earn system-specific currencies such as *mojo* [24] or *karma* [22] which can be used for redeeming service. The use of market currencies, supported via micropayments, has been studied analytically in [12, 8]. In reputation-based schemes, the transactional history of each peer is used to compute its reputation [2, 15], which in turn dictates the level of service obtainable by the peer. BitTorrent is the first system that explicitly adopts the *Bit-for-Bit* scheme for p2p file sharing [7].

P2P Streaming Protocols: Most of the p2p streaming protocols today assume a "cooperative" (i.e. obedient) peer behavior. This implies that the system designer has complete control over the behavior of individual peers. Assuming peers are not strategic, these work focus on protocol designs to achieves the best outcome based on the level of cooperation considered. In ESM [4] and Bullet [9], peers are completely cooperative in contributing resource to the p2p system. The objective of both protocols is to deliver the highest multicast throughput to all peers by harnessing the bandwidth effectively. By assuming peers are obedient, the game designer (the protocols) can achieve optimal outcome that meets its objective. SplitStream [3] and CoopNet [18] assume minimal cooperation from peers and implements a *Bit-for-Bit* outcome. We leverage many of the ideas there for the protocol design. Lately, [6] devises a new protocol which allows the publisher to specify the spectrum of cooperative policy to meet different performance objectives. Finally, a reputation-based incentive mechanism has been proposed that leverages service-differentiated peer selection in many-to-one p2p streaming sessions to encourage user contribution [13].

7. SUMMARY

We identify three contributions in this paper. (i) We leverage the uniqueness of the p2p streaming context and propose *taxation* as an incentive mechanism to achieve a desirable outcome. The enabling observation is the asymmetry of power, where the publisher has the power to design a taxation game and enforce the rules on the participating peers. We believe the concept of taxation is novel and has not been introduced in other p2p context. (ii) We show that taxation is an effective means to maximize social welfare when peers are strategic in a heterogeneous p2p environment. In an environment consisting of 80% low capacity peers, the social welfare improves by 50% compared to Bit-for-Bit. (iii) We demonstrate that linear taxation can be implemented efficiently in a distributed streaming protocol with reasonable overhead. The protocol consistently utilizes above 95% of the resource and the allocation is compliant to the tax schedule within 2 bandwidth unit 80% of the time. However, this comes at a cost of structural instability due to an increase in path rate changes. We show the cost is twice the fundamental cost in maintaining the structure.

One possible objection to the taxation scheme is that mandatory taxation may displace voluntary contributions by altruistic peers. It has been hypothesized that government grants (financed through taxation) *crowd out* private philanthropy [23, 1]. However, empirical studies have provided conflicting evidence on this matter [14, 19]. Therefore, it may be prudent to design the system such that altruistic peers, if they exist, are not prohibited from contributing more than is required by the taxation scheme.

8. REFERENCES

[1] T. Bergstrom, L. Blume, and H. Varian. On the private provision of public goods. journal of public economics 29: 25-49, 1986.

[2] C. Buragohain, D. Agrawal, and S. Suri. A game theoretic framework for incentives in p2p systems. In *International Conference on Peer-to-Peer Computing*, 2003.

[3] M. Castro, P. Druschel, A. Kermarrec, A. Nandi, A. Rowstron, and A. Singh. SplitStream: High-bandwidth Content Distribution in Cooperative Environments. In *Proceedings of SOSP*, 2003.

[4] Y. Chu, A. Aganjam, T. Ng, S. Rao, K. Sripanidkulchai, J. Zhan, and H. Zhang. Early Experience with an Internet Broadcast System. Technical report CMU-CS 03-214, Carnegie Mellon University, December 2003.

[5] Y. Chu, A. Ganjam, T. Ng, S. Rao, K. Sripanidkulchai, J. Zhan, and H. Zhang. Early experience with an internet broadcast system based on overlay multicast. In *USENIX Symposium on Internet Technologies and Systems*, jun 2004.

[6] Y. Chu and H. Zhang. Considering altruism in peer-to-peer internet streaming broadcast. In *ACM NOSSDAV*, June 2004.

[7] B. Cohen. Incentives build robustness in bittorrent. In *1st Workshop on Economics of Peer-to-Peer Systems*, June 2003.

[8] C. Courcoubetis and P. Antoniadis. Market models for p2p content distribution. In *1st International Workshop on Agents and Peer-to-Peer Computing*, July 2002.

[9] J. A. D. Kostic, A. Rodriguez and A. Vahdat. Bullet: High Bandwidth Data Dissemination Using an Overlay Mesh. In *Proceedings of SOSP*, 2003.

[10] End system multicast toolkit and portal. http://esm.cs.cmu.edu/.

[11] J. Feigenbaum and S. Shenker. Distributed algorithmic mechanism design: Recent results and future directions. In *6th International Workshop on Discrete Algorithms and Methods for Mobile Computing and Communications, pp. 1-13*, 2002.

[12] P. Golle, K. Leyton-Brown, and I. Mironov. Incentives for sharing in peer-to-peer networks. In *ACM Conference on Electronic Commerce*, 2001.

[13] A. Habib and J. Chuang. Incentive mechanism for peer-to-peer media streaming. In *IEEE International Workshop on Quality of Service (IWQoS'04)*, 2004.

[14] B. R. Kingma. An accurate measurement of the crowd-out effect, income effect, and price effect for charitable contributions. journal of political economy 97: 1197-207, 1989.

[15] H. Kung and C.-H. Wu. Differentiated admission for peer-to-peer systems: Incentivizing peers to contribute their resource. In *1st Workshop on Economics of Peer-to-Peer Systems*, June 2003.

[16] J. Mirrless. An exploration in the theory of optimum income taxation, review of economic studies 38 (1971): 175-208.

[17] C. Ng, D. Parkes, and M. Seltzer. Strategyproof computing: Systems infrastructures for self-interested parties. In *1st Workshop on Economics of Peer-to-Peer Systems*, June 2003.

[18] V. Padmanabhan, H. Wang, and P. Chou. Resilient Peer-to-peer Streaming. In *Proceedings of IEEE ICNP*, Nov. 2003.

[19] A. A. Payne. Does the government crowd-out private donations? new evidence from a sample of non-profit firms. journal of public economics 63: 323-45, 1998.

[20] N. Stern. On the specification of models of optimum income taxation. journal of public economics 6 (july-august 1976): 123-162.

[21] M. Veseth. Public finance, reston publishing, 1984.

[22] V. Vishnumurthy, S. Chandrakumar, and E. Sirer. Karma: A secure economic framework for p2p resource sharing. In *1st Workshop on Economics of Peer-to-Peer Systems*, June 2003.

[23] P. G. Warr. Pareto optimal redistribution and private charity. journal of public economics 19: 131-8, 1982.

[24] B. Wilcox-O'Hearn. Experiences deploying a large-scale emergent network. In *1st International Workshop on Peer-to-Peer Systems*, Mar. 2002.