Natural Selection in Peer-to-Peer Streaming: From the Cathedral to the Bazaar

Vivek Shrivastava
University of Wisconsin-Madison, USA
viveks@cs.wisc.edu

Suman Banerjee
University of Wisconsin-Madison, USA
suman@cs.wisc.edu

ABSTRACT

Success of peer-to-peer applications in many cases is attributed to user altruism, where a user contributes some of its own resources to facilitate performance of other users. This observation has been corroborated with some experimental evidence. In this paper we make a first attempt to demonstrate that there are many scenarios where peer-to-peer resource sharing is a natural behavior that selfish users can use to improve their own performance. In particular we examine such natural incentives that exist in a streaming media application which lead such greedy users to cooperate and share resources with each other in forming an efficient overlay multicast tree. We define a freestyle *Bazaar* environment in which streaming media receivers interact with each other and cooperatively construct an overlay tree for improving their perception of media streams from a single server. Through simulations we demonstrate the efficacy of our proposed environment.

Categories and Subject Descriptors: C.2.4 Distributed Systems - Distributed Applications
General Terms: Design, Experimentation, Performance
Keywords: media streaming, peer-to-peer networks, market model

1. INTRODUCTION

Peer-to-Peer (p2p) networking paradigms offers new possibilities for content distribution over the network. Recent research has shown that the level of peer altruism is a key factor for the success of p2p applications [10]. Kazaa and Gnutella [1] are examples of popular file sharing applications that rely on peer altruism. However a recent study by Saroiu et. al. [14] has shown that majority of the participating peers in such p2p applications are free riders. It is perceived that there is no incentive for peers to contribute resources in p2p applications. This is why new applications like BitTorrent [6] employ a *tit-for-tat* rule to facilitate file sharing. But as we show in this paper, there are scenarios in *p2p streaming media applications* where resource sharing is a natural behavior that can facilitate the formation of efficient multicast trees in the absence of external rules or incentives.

Altruism vs Incentives in p2p Streaming: In a server based approach for streaming (Figure 1(a)), a high load is imposed on the server streaming the content as all the interested parties directly join the server for downloading the content. While in a p2p streaming scenario (Figure 1(b)), the participating peers can also forward the content to other interested peers, thereby reducing the bandwidth constraints for the server. There are two distinct entities in a peer-to-peer streaming environment: the *publisher* and the *interested peers*. The publisher makes the content available on the network and the peers interested in the content form an overlay for receiving and distributing data in a distributed fashion [9].

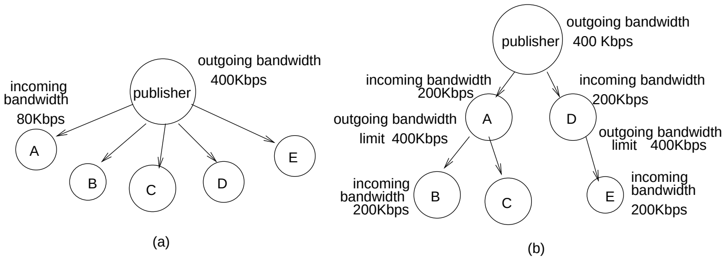

Figure 1: Illustration of a bandwidth degradation problem in the presence of selfish nodes and a possible solution for the same

In the best scenario, all the participating peers are altruistic and freely contribute the available content to other interested peers in the overlay. Past research has shown that upstream bandwidth is the bottleneck resource in p2p environment [14, 17]. A typical audio/video stream takes 350 Kbps to encode. Often, most of the peers are behind DSL and cable modems and are characterized by asymmetrical bandwidth capacities with low forwarding capacity and high receiving capacity [9]. Furthermore making content accessible to others reduces a peer's access bandwidth and degrades its network access performance. So a strategic peer will essentially avoid streaming content to other peers unless the cost incurred due to data forwarding is offset by other factors (as described later).

Most of the existing research in p2p streaming [9, 17, 8, 7, 5, 1, 11, 4], have been focused on introducing various centrally imposed rules and incentives to motivate the peers to contribute forward bandwidth to other peers in the system. We call such centrally imposed rule-based mechanisms (each peer is expected to follow a same set of rules), the *Cathedral* approach. In this paper, we make a key observation that the p2p streaming environment has inherent *natural incentives* for participating peers to contribute bandwidth to the peer community. We present a simple architecture in which such natural incentives can be exploited and allow natural selection guide the formation of an efficient overlay tree for data streaming. Our approach marks a shift from the *Cathedral* style mechanism to a free *Bazaar* mechanism, where no rules are being imposed on the peers and resource sharing takes place naturally as peers try to maintain their perceived data *utility*.

Utility: The utility of the streaming content for a participating peer i can be modeled as:

$$\text{benefit}(bandwidth_{incoming}, latency) \text{ - cost}(bandwidth_{outgoing})$$

where *benefit* function captures the perceived data quality and *cost*

---

[1] See http://www.kazaa.com and http://www.gnutella.org

NOSSDAV'05, June 13–14, 2005, Stevenson, Washington, USA.

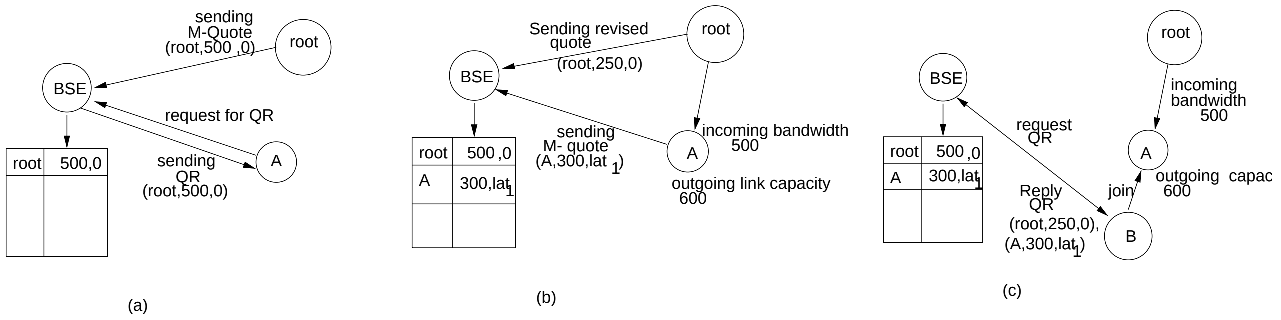

Figure 2: Working of Bazaar framework (a) *root* joins the system and sends market quote to BSE. Peer A enters the system and requests *M-Quotes* from BSE. As *root* is the only existing member of the overlay, A sends a join request to the root. (b) *root* starts streaming data to A at 500 Kbps. A being strategic advertises 300 Kbps which is greater than its parent's quote of 250 Kbps. lat_1 *is the latency incurred on the path from root to peer A along the overlay structure*.(c) peer B joins the system and receives M-Quotes of *root* and A. It selects A as its parent and send a join request to A. *All bandwidths in the figure are in units of Kbps.*

function captures the cost of forwarding data [9, 17]. The specific forms of these functions are given in Section 3.

A Motivating Example: Let us consider the following scenario depicted in Figure 1(a). A publisher with some video content decides to allocate 400 Kbps of its total outgoing link capacity for streaming the content on the Internet. Suppose there are 5 peers interested in the video. If all the interested peers try to download the data directly from the publisher and do not contribute bandwidth to other peers in the system, then the participating peers will form a *star* topology with the publisher at the center. Hence each peer receives a bandwidth of $(\frac{400}{5})$, i.e., 80 Kbps. This simple scenario manifests that *as the number of interested selfish peers directly connected to the publisher increases, the bandwidth received by each such peer decreases and so does its perceived data quality.*

Now consider a different overlay topology as shown in Figure 1(b). In this overlay tree, all the peers A,B,C and D receive 200 Kbps of streaming data without violating access bandwidth constraints. Here peer A forwards a total of 400 Kbps to its children B and C. By doing so, it reduces the load at the publisher and hence both A and D manage to receive 200 Kbps directly from the publisher. But in this case, A has to incur the cost of forwarding 400 Kbps which will reduce its own access bandwidth and may degrade its network access performance to some extent. *Hence there is a clear trade off for peer A between forwarding* 400 *Kbps and receiving* 200 *Kbps from the publisher or not forwarding any data and receiving* 80 *Kbps from the publisher*. If increase in A's perceived data utility due to the increase in incoming bandwidth (200 Kbps from 80 Kbps) offsets the loss due to data forwarding (400 Kbps), then peer A will choose to forward data to B and C, and try to form the overlay shown in Figure 1(b). Since each peer may have a different relationship between incoming bandwidth and perceived data utility, the actual trade off parameters are different for different peers and the final overlay will depend on the mix of utility functions of the participating peers. This example clearly outlines the *natural incentive* for peer A to maintain its own incoming bandwidth from the publisher by forwarding data to B and C. Also peer B now has a natural incentive to join under peer A as it receives a substantially higher streaming bandwidth (200 Kbps), as compared to bandwidth received in the *star* topology (80 Kbps).

So if we enable the peers to exploit this natural incentive and form the second overlay, then we can improve the overall performance of the system and more peers will be able to receive the content of high quality. In this paper, we make a first attempt to provide a platform which facilitates the formation of such a mutually beneficial overlay without introducing any rules or incentives in the system.

We propose a simple *Bazaar* model that takes into account the aforementioned analysis and enables efficient overlay tree construction for streaming applications. Our model does not impose any specific rules on the peers and does not assume any asymmetry of roles and power. It is based on the natural incentive of the peers to conserve their own incoming bandwidth by attracting the new entrant to join under itself rather than its parent. An appealing feature of this model is that it allows each peer to behave independently of each other, i.e., some of them can be altruistic, while others are strategic. In contrast, most other prior schemes assume that all peers are either strategic or altruistic.

Roadmap: We begin by explaining the working of our *Bazaar* framework in Section 2. Next we present details of our simulation set up in Section 3, followed by the results in Section 4. In Section 5 we provide a brief discussion on some additional issues followed by related work in Section 6. We conclude the paper in Section 7.

2. THE P2P STREAMING BAZAAR

In this section, we describe our *Bazaar* framework for p2p streaming. There are three basic entities in our bazaar:(1) regular peers, (2) the BSE (Boot Strap Entity) and (3) the root (publisher). The peers are strategic in nature and try to maximize their incoming bandwidth and minimize their outgoing bandwidth. The BSE bootstraps the nodes joining the overlay with information about other nodes in the overlay. The *root* makes the content available on the overlay and assigns a fixed fraction of its total forwarding capacity for streaming purposes. It is not strategic as it allocates a fixed amount of bandwidth for streaming purposes that remains constant for the entire session of streaming irrespective of the peer dynamics in the system. The bandwidth received by any entity is termed as its *incoming bandwidth* and the total bandwidth contributed to other peers is termed as *outgoing bandwidth* of the peer.

M-Quote. Each node in the system provides a *Market Quote (M-Quote)* of its services to attract other peers to join under it rather then its parent, so as to preserve its own incoming bandwidth. M-Quote of a peer has the following advertised components:(1) *bandwidth* that the advertising peer is willing to forward to any peer that joins as its immediate child, (2) *latency* incurred by the data to reach the peer from the root along the overlay. *Note that the quote*

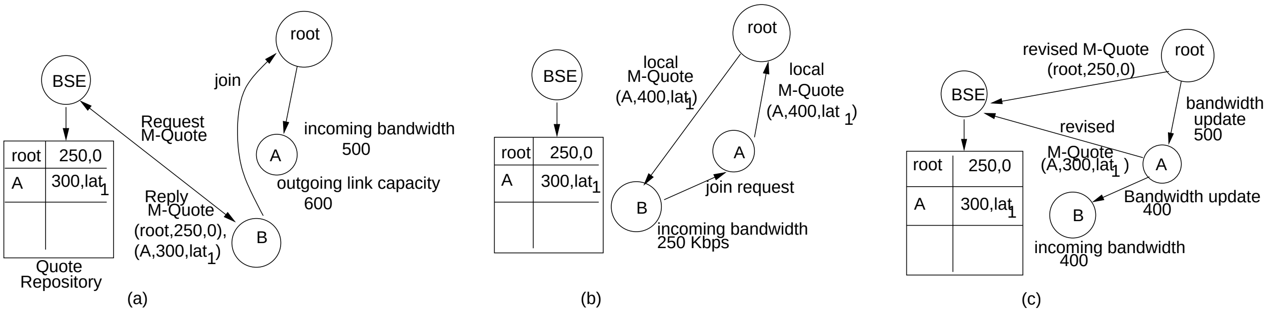

Figure 3: Working of Shuffle-1 operation (a) peer B entering the system joins the *root* as B's perceived utility may be higher for *root's* quote than A's quote. (b) peer A sends a local *M-Quote* to the *root* with a quote of 400 Kbps for peer B. The *root* forwards the quote to peer B, who accepts the quote and joins peer A. (c) peer A starts receiving full quota of 500 Kbps and forwards 400 Kbps to peer B. lat_1 *is the latency incurred in the path from root to A in the overlay. All bandwidths in the figure are in units of Kbps*

is an offer to forward data of a certain quality (determined by the bandwidth and latency values in the quote) and not a financial offer.

Quote Repository. In our scheme, BSE maintains a *quote repository (QR)* of quotes from various participating peers in the streaming environment.

In our *Bazaar* framework, peers perform the following basic actions:(1) Join/Leave the overlay dynamically, (2) Advertise their services by sending *M-quote* to BSE, (3) Participate in *Shuffle* operations to improve overlay structure (as described below). The detailed working of our *Bazaar* framework is illustrated by example in Section 2.1.

2.1 Bazaar in Action

We will describe our approach in detail with the help of following example (Figure 2). In Figure 2(a), the root (publisher) enters the system with some content and sends a M-Quote to the BSE containing the bandwidth allocated for streaming purposes (500 Kbps). The format of the M-Quote is (peer id, bandwidth advertised, latency of data at the peer). This quote is stored by the BSE in the Quote Repository(QR). In Figure 2(b) another peer A, interested in receiving the content enters the system and queries the BSE for the M-Quotes available from currently participating peers. Since A is the first peer (apart from root) to join the system, BSE replies with the M-Quote of the root only. On receiving the quote, the incoming peer sends a join request to the root, which adds the incoming peer A as its child and starts forwarding data to it at 500 Kbps as mentioned in its market quote. Next as depicted in Figure 2(c), the root recomputes its market quote and sends a revised M-Quote to the BSE. This revised quote has a bandwidth value of 250 Kbps, and signifies that if a new node joins the root now, it will receive only 250 Kbps as it will share the quota of 500 Kbps with existing child A.

As peer A starts receiving the full capacity of the root's allocated bandwidth (500 Kbps), it may decide not to forward any bandwidth to the other peers. It can therefore send a market quote stating available bandwidth as zero or may not send the quote at all. Now if another peer B joins the system and requests the M-Quotes of currently participating peers in the system, it will receive a M-Quote of 250 Kbps and 0 corresponding to root and peer A respectively. Naturally the incoming peer will decide to join the root. This will lead A and B to share the allocated bandwidth of 500 Kbps at the root. As a result, the incoming bandwidth for peer A also gets reduced to 250 Kbps and hence the perceived data utility for peer A may degrade sharply.

So in order to conserve its own incoming bandwidth, peer A will try to advertise a competitive quote in the market, better than the root's quote, so that any incoming peer will select peer A as its parent with high probability and will join under A instead of joining under root directly. The final decision of choosing a parent depends on the exact utility function of the incoming peer. Peer A also has to compensate for the higher latency associated with its data ($latency\ at\ A\ =\ link\ delay\ from\ publisher\ to\ A$). So advertising a bandwidth quote higher than its parent's quote is an attractive alternative for peer A to sustain its incoming bandwidth.

Strategic Peer Behavior. Since peers are strategic, A will not advertise its entire outgoing link capacity for streaming. It will just send a M-Quote sufficient to prevent the incoming peers from joining the root. This behavior is typical of strategic peers who try to minimize their forward bandwidth and maximize their incoming bandwidth. So as shown in Figure 2(c), peer A advertises bandwidth of 300 Kbps, which is greater than its parent's quote of 200 Kbps but less than its maximum possible streaming capacity of 500 Kbps.

On receiving M-quotes of the root and peer A, B will compute its perceived utility for both the quotes and send a join request to the node whose quote produces a higher utility. If the high bandwidth quote by A results in a better utility, then B will send a join request to A and hence A's incoming bandwidth is preserved. But if higher latency at A (compared to root) offsets A's advantage of higher bandwidth quote, then B will send a join request to the root (Figure 3(a)). The root on receiving the join request, adds peer B as its child and splits its total allocated bandwidth (500 Kbps) equally between peers A and B. In this case, the share for peer A, drops to 250 Kbps It may be the case that A has a capacity of forwarding more than 300 Kbps, but it did not advertise high bandwidth in its M-Quote to conserve its outgoing link bandwidth. So if the degradation in perceived utility motivates peer A to regain higher incoming bandwidth, it can participate in a local shuffling operation that provides an opportunity for peer A to tweak the overlay into a more favorable formation.

Local Shuffle. It is an operation performed periodically to improve the efficiency of the overlay structure. A *Shuffle* operation that involves peers who are maximum *k* levels apart in the overlay tree is termed as *Shuffle-k* operation. Figure 3 illustrates the working of *Shuffle-1* operation. As shown in Figure 3(b), peer A will now send a M-Quote to its parent with a higher bandwidth offer of

400 Kbps for the peer B. On receiving the quote, root will broadcast the quote to peer B. As peer B receives the quote, it will calculate the perceived utility as per the quote offered by A and if the new quote results in a utility greater than its current utility, it will join peer A. The root, will remove B from its list of children and hence the share for A's bandwidth gets doubled. Next as shown in Figure 3(c), peer A will add peer B as its child and start forwarding data to peer B at 400 Kbps. It is worth noting that if peer A does not forward at its advertised bandwidth of 400 Kbps, then peer B will simply revert back to the root and it will again reduce the perceived data utility for peer A. So peer A has no incentive to cheat another peer by advertising a wrong quote.

We propose a general *Shuffle-k* operation that can involve peers from k different levels in the overlay multicast tree. In this paper, we have presented our *Bazaar* framework with *Shuffle-1* implementation.

3. EVALUATION

Here we present the experimental setup for our simulations and assess the performance of our scheme in various scenarios of peer heterogeneity. We analyze the efficacy of our free *Bazaar* framework for forming an efficient overlay tree by measuring the perceived data utility achieved by participating peers.

Utility Functions: As discussed before in Section 1, strategic peers try to maximize their incoming bandwidth and minimize their forward bandwidth. We use the following model for utility function as proposed in [9].

$$U = \lambda \cdot \sqrt{r} - \alpha \cdot \sqrt{F} \cdot \mu \cdot (\frac{f}{F}) \cdot (1 - \mu) \cdot (\frac{f}{F})^4 + \frac{1}{\sqrt{l}} \quad (1)$$

where r is the incoming bandwidth,f is the bandwidth forwarded, F is the outgoing bandwidth capacity and l is the latency associated with incoming data. For our simulations, λ is set to 2, μ is set to 0.5 and α is set to 0.25. We have intentionally chosen λ to be high as incoming bandwidth is extremely crucial for video/audio streaming applications. However, we wish to point out that our proposed framework is independent of the utility functions used by the peers. Please refer to [9] for rationale regarding the shape of these functions.

Simulation Environment: The simulations were conducted on an open source peer-to-peer network simulator *myns* [13], developed at University of Maryland. We use Transit Stub topology, generated using the GT-ITM topology generator [18], as our underlying network.

We conducted our simulations with a group size of 50 peers and all results are averaged over 1000 permutations of peer join order. The publisher in our simulations allocates a fixed fraction of its outgoing bandwidth capacity for streaming. The allocated bandwidth gets shared amongst all the children joining the publisher directly. We evaluate our *Bazaar* framework for the following modes of operation:

Altruistic Here all the peers are completely altruistic and are ready to contribute their total outgoing link capacity for forwarding data to other peers in the overlay.

Strategic This is a more realistic scenario where all the peers except the publisher are strategic and try to maximize their incoming bandwidth and minimize their forwarding bandwidth.

Random Here all the peers are strategic except the publisher, but the peers join the overlay tree at random positions without taking advantage of the underlying *Bazaar* framework.

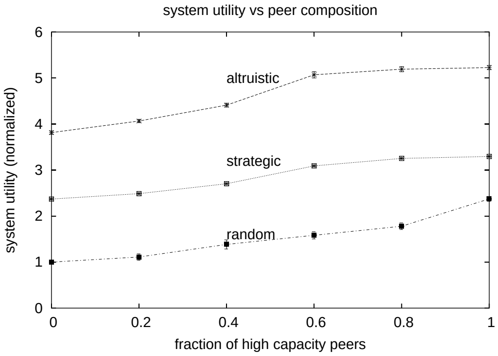

Figure 4: Comparison of total system utility for altruistic, strategic and random modes

Metrics: We compare the aforementioned modes in terms of the *throughput(incoming bandwidth)* achieved by each peer in the overlay and the *total system utility (defined as sum of perceived data utility of all the participating peers)* in the three modes of operation discussed above.

Supporting a mix of Altruistic & Strategic Peers: Existing protocols implicitly assume peers are either completely altruistic(e.g. ESM [5], Bullet [12], NICE [1]) or completely selfish(e.g.Taxation [9], StrategyProof [17]) and tailor their designs as per their assumption [10]. However, our freestyle *Bazaar* framework makes no underlying assumptions and will allow the overlay tree structure to adapt to any mixture of altruistic and strategic peers. We have performed experiments with such peers compositions as well.

4. RESULTS

Here we present the results from our Bimodal and Trace Based Simulations.

4.1 Bimodal Simulations

In these set of simulations all the participating peers are categorized as (1) high capacity peers (2) low capacity peers. In each experiment, a high capacity peer is assigned an outgoing bandwidth capacity at random from a high range of bandwidths varying from 500 Kbps - 1 Mbps. Similarly a low capacity peer is assigned outgoing bandwidth capacity at random from a range of 50 Kbps - 450 Kbps. We simulate heterogeneous peer environments by varying the fraction of high capacity peers in the environment from 0 to 1. The maximum streaming rate is 500 Kbps for this set of simulations, i.e., the publisher allocates 500 Kbps for streaming data and this bandwidth gets shared equally amongst all the children joining the publisher directly. We compare the altruistic, strategic and random modes of operation for bimodal bandwidth distribution.

System Utility: Figure 4 shows the system utility of the streaming environment under three different modes. The values on the x-axis represent the fraction of high capacity peers in the streaming environment. The y-axis represents the system utility (normalized with respect to smallest utility observed) for the corresponding peer composition. We also compute the standard deviation of the system utility over 1000 runs of simulations and find that it is a tight bound. We make two observations: (1) As expected, the system utility increases with the increase in fraction of high capacity peers in all the compared modes. This behavior stems from the fact that the cost of forwarding bandwidth is modeled as the fraction of actual forward bandwidth and the total link capacity. So as the total link capacity increases, the cost to forward same amount of

Capacity of peers	High 10 Mbps	Medium 1.5 Mbps	Low 100 Kbps
Sigcomm	76	2	22
Slashdot	22	4	74
Gnutella	08	27	65

Table 1: peer composition from real Internet broadcast events

bandwidth decreases and hence the willingness to forward data to other peers also increases. (2) The strategic mode using the bazaar framework outperforms the random mode under varying peer compositions. This is because in random scheme the incoming peer randomly joins any existing member of the overlay, which may not be a good strategy and hence it may degrade its own utility as well as the utility of the other children of the parent it joined. The altruistic mode performs even better than the strategic mode but then it makes the fundamental assumption of altruism.

Throughput: Figures 5- 7 compare the CDF of throughput for the *altruistic* mode and the *strategic* mode. As we move from Figure 5 to 7, the fraction of high capacity peers increases from 20% to 80%. These figures show that the gap in peer throughput is minimum in the case where fraction of high capacity peers is 20% and keeps on increasing as we reach peer environment with 80% high capacity peers. We can conclude from the above trend that the performance gap between the altruistic mode and the strategic mode *decreases* with the *decrease in fraction of high capacity peers*. As we will show for realistic traces in the next set of simulations, fraction of high capacity peers is indeed low in real life streaming environments and hence our framework will enable strategic peers to achieve good throughput in such real life scenarios.

4.2 Trace Based Simulations

Here the outgoing capacity distribution of peers are based on traces collected from real Internet broadcast events, namely Sigcomm, Slashdot and Gnutella [2]. In this set of simulations, a peer is assigned a outgoing bandwidth capacity of (1) 10 Mbps (2) 1.5 Mbps (3) 100 Kbps depending on the percentage composition of that trace. The maximum streaming rate is 500 Kbps for this set of simulations. The composition of these traces are shown in Table 1. In Figures 8,9 and 10, we analyze the performance of the *Bazaar* (in terms of throughput achieved by peers) by varying the fraction of strategic peers in the environment. The bandwidth of nodes in the environment are derived from respective traces, however the fraction of strategic peers in the environment is varied systematically from $zero$ (all altruistic peers) to one (all strategic peers). The top most curve in these figures correspond to the completely altruistic scenario and the the bottom most curve corresponds to the completely strategic scenario. We observe that for Sigcomm trace, the performance of the *Bazaar* degrades substantially as the fraction of strategic peers in the environment increases, while the Slashdot and Gnutella traces degrade gracefully with the increase in fraction of strategic peers. Hence we can infer that our *Bazaar* framework is particularly well suited for many of the p2p streaming scenarios, in which peers are mostly resource poor (like Gnutella, Slashdot).

5. DISCUSSION

Distributed Quote Repository: While we discuss the market architecture in context of a centralized *QR*, it is possible to maintain the *QR* in a distributed fashion, where each parent stores the quote for its descendants. Also each peer periodically forwards its local quote repository (containing its own quote and the quote of its descendants) to its parent. In this bottom-up manner, quotes of all the participating peers is collected at the *root* of the overlay, which can then broadcast the combined quote on the overlay.

Schemes	*Bandwidth allocation style*	*Incentive type*	*Peer type assumed*
Payment based VCG [17]	Cathedral	Financial	Strategic
Rule Based Taxation [9]	Cathedral	Rule Based	Strategic
Altruism based ESM [5], NICE [1]	Cathedral	None	Altruistic
Bazaar	Bazaar	Natural	Strategic or altruistic

Table 2: Comparison of various p2p mechanisms with *Bazaar* model

6. RELATED WORK

Market Based Schemes: Many market based schemes [16, 7, 15] have been proposed for p2p applications, where a peer collects revenues from other peers in return of forwarding content to them.

Altruism Based Schemes: In the past researches have proposed many schemes to optimize bandwidth allocation in p2p streaming environments. But a majority of such approaches like ESM [5], NICE [1], Bullet [12], Overcast [11], SCRIBE [4], assume complete altruistic behavior from the peers. Though such schemes enable optimal tree formation, the underlying assumption is not true in p2p streaming environment.

Rule Based Schemes: Chu et. al. [9] discuss the same issue of bandwidth allocation in streaming applications and make use of the Splitstream [3] scheme in conjunction with taxation. Their model assumes that the owner of the content, typically the publisher, determines the amount of bandwidth a peer should forward in order to receive a particular amount of bandwidth. Here it is envisioned that the publisher will enforce the tax payment on the participating peers.

Payment Based Schemes: Recently, a strategy-proof scheme [17] was proposed for overlay multicast in the presence of strategic peers. Here some financial incentives are introduced in the system to motivate them to join the overlay in such a manner that improves the utility for the entire system. While this approach tends to optimize the *social welfare* of the system, it has some inherent problems of budget overflow and has to rely on a trusted third party to manage the payments in the system.

Table 2 provides a brief comparison of our *Bazaar* model with the various class of mechanisms described above.

7. SUMMARY

We have proposed a free *Bazaar* framework for achieving improved performance in p2p streaming applications involving strategic peers. The enabling observation is that p2p streaming structure has a natural inherent incentive for peers to contribute bandwidth to the community. We propose a bazaar framework, which leverages this natural incentive of the participating peers and facilitates the formation of an efficient overlay for data streaming. Some optimizations to the basic bazaar framework are also proposed in the form of *Shuffle-k* operations. We have currently reported the performance of the *Bazaar* with *Shuffle-1* operation.

Our results show that our scheme performs quite efficiently in peer environments comprising of low fraction of high capacity peers. Also real Internet broadcast traces traces indicate that in reality fraction of high capacity peers is indeed low and our *Bazaar* is well suited for such scenarios. An additional advantage of the *Bazaar* framework is that it supports an arbitrary mix of strategic and altruistic peers and the performance (throughput achieved by peers)

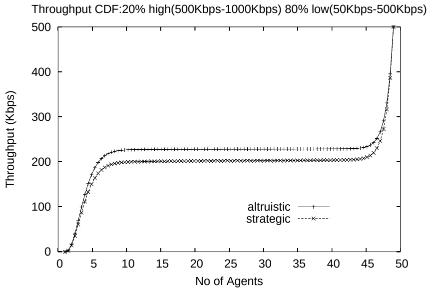

Figure 5: Comparison of Altruistic & Strategic CDF for Throughput: 20***% high capacity peers*** 80***% low capacity peers***

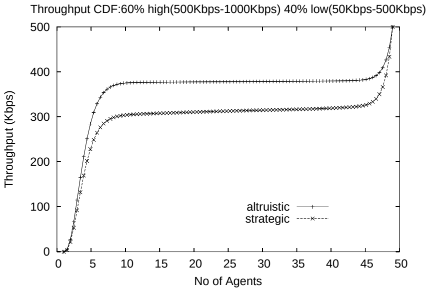

Figure 6: Comparison of Altruistic & Strategic CDF for Throughput: 60***% high capacity peers*** 80***% low capacity peers***

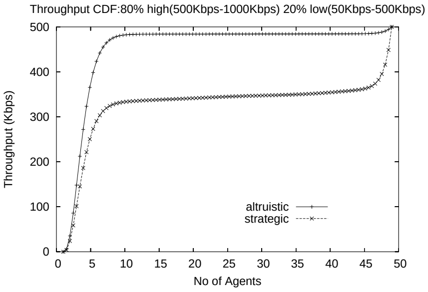

Figure 7: Comparison of Altruistic & Strategic CDF for Throughput: 80***% high capacity peers*** 20***% low capacity peers***

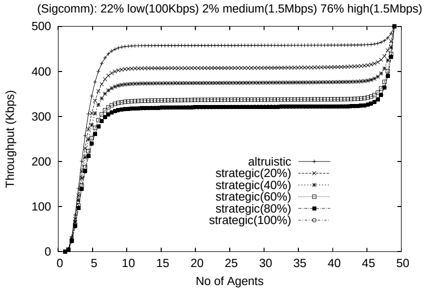

Figure 8: Sigcomm :CDF of Throughput achieved by participating peers under varying mixtures of altruistic and strategic peers

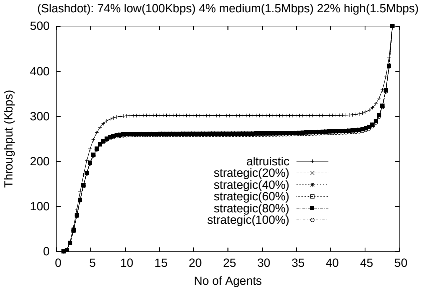

Figure 9: Slashdot :CDF of Throughput achieved by participating peers under varying mixtures of altruistic and strategic peers

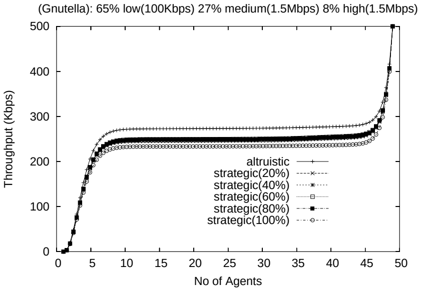

Figure 10: Gnutella :CDF of Throughput achieved by participating peers under varying mixtures of altruistic and strategic peers

degrades gracefully as the environment varies from completely altruistic to completely strategic.

8. REFERENCES

[1] S. Banerjee, B. Bhattacharjee, and C. Kommareddy. Scalable application layer multicast. In *ACM Sigcomm*, 2002.

[2] A. Bharambe, S. Rao, V. Padmanabhan, S. Seshan, and H. Zhang. The impact of heterogeneous bandwidth constraints on dht-based multicast protocols. In *International Workshop on P2P Systems*, 2005.

[3] M. Castro, P. Druschel, A. Kermarrec, A. Nandi, A. Rowstron, and A. Singh. Splitstream: High-bandwidth multicast in cooperative environments. In *ACM SOSP*, 2003.

[4] M. Castro, P. Druschel, A. Kermarrec, and A. Rowstron. SCRIBE: A large-scale and decentralized application-level multicast infrastructure. *IEEE JSAC*, 2002.

[5] Y.-H. Chu, S. G. Rao, and H. Zhang. A case for end system multicast. In *ACM Sigmetrics*, 2000.

[6] B. Cohen. Incentives build robustness in bittorrent. In *Workshops on Economics of Peer-to-Peer Systems*, 2003.

[7] Courcoubetis and Antoniadis. Market models for p2p content distribution. In *First International Workshop on Agents and Peer-To-Peer Computing (AP2PC)*, 2002.

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

[9] Y. hua Chu, J. Chuang, and H. Zhang. A case for taxation in peer-to-peer streaming broadcast. In *Workshop on Practice and theory of incentives in networked systems (PINS)*, 2004.

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

[11] J. Jannotti, D. K. Gifford, K. L. Johnson, M. F. Kaashoek, and J. W. O'Toole, Jr. Overcast: Reliable multicasting with an overlay network. In *ACM OSDI*, 2000.

[12] D. Kosti, A. Rodriguez, J. Albrecht, and A. Vahdat. Bullet: high bandwidth data dissemination using an overlay mesh. In *Proceedings of the nineteenth ACM SOSP*, 2003.

[13] myns(p2p) Simulator. www.cs.wisc.edu/∼suman/research/myns/index.html.

[14] S.Saroiu, K.P.Gummadi, and S.D.Gribble. Measuring and analyzing the characteristics of napster and gnutella hosts. *Multimedia Systems*, 9(2), 2003.

[15] B. Strulo, A. Smith, and J. Farr. An architecture for peer-to-peer economies. In *International Conference on Peer-to-Peer Computing*, 2003.

[16] S. C. V. Vishnumurthy and E. G. Sirer. Karma: A secure economic framework for p2p resource sharing. In *Workshop on the Economics of Peer-to-Peer Systems*, 2003.

[17] S. Yuen and B. Li. Strategyproof mechanisms for dynamic multicast tree formation in overlay networks. In *IEEE Infocom*, Mar. 2005.

[18] E. W. Zegura, K. L. Calvert, and S. Bhattacharjee. How to model an internetwork. In *IEEE Infocom*, 1996.