

When is P2P Technology Beneficial for IPTV Services?

Yih-Farn Chen, Yennun Huang, Rittwik Jana, Hongbo Jiang, Michael Rabinovich,
Bin Wei, and Zhen Xiao*

ABSTRACT

This paper studies the conditions under which peer-to-peer (P2P) technology may be beneficial in providing IPTV services over typical network architectures. It has two major contributions. First, we contrast two network models used to study the performance of such a system: a commonly used logical “Internet as a cloud” model and a “physical” model that reflects the characteristics of the underlying network. Specifically, we show that the cloud model overlooks important architectural aspects of the network and may drastically overstate the benefits of P2P technology by a factor of 3 or more. Second, we provide a cost-benefit analysis of P2P video content delivery focusing on the profit trade-offs for different pricing/incentive models rather than purely on capacity maximization. In particular, we find that under high volume of video demand, a P2P built-in incentive model performs better than any other model for both high-definition and standard-definition media, while the usage-based model generally generates more profits when the request rate is low. The flat-reward model typically performs better than the usage-based model - especially when most requests are for high-definition media.

Keywords : IPTV, P2P streaming, Content distribution network, FTTN, Video-on-Demand.

1. INTRODUCTION

Internet protocol TV (IPTV) promises to offer viewers an innovative set of choices and control over their TV content. Two major U.S. telecommunication companies, AT&T and Verizon, have invested significantly to replace the copper lines in their networks with fiber optic cables for delivering many IPTV channels to residential customers.

*Yih-Farn Chen, Rittwik Jana, and Bin Wei are with AT&T Shannon Research Laboratory; Y. Huang is with the Institute for Information Industry, Taiwan; H. Jiang and M. Rabinovich are with Case Western Reserve University; Z.Xiao is with IBM T. J. Watson Research Labs. The contact author is Yennun Huang: yennunhuang@gmail.com.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

NOSSDAV '07 Urbana, Illinois USA

Copyright 2007 ACM 978-1-59593-746-9/06/2007 ...\$5.00.

A viewer can receive IPTV videos in good quality if the available bandwidth satisfies the need of video encoding rate for the target resolution and frame rate. To provide sufficient bandwidth for IPTV services, Internet service providers use high speed xDSL or cable networks to deliver video content to viewers' set-top boxes. As an example, AT&T Light-Speed is using Fiber-to-the-Neighborhood (FTTN) Networks. Its architecture consists of a small number of national super head-ends (SHE) and a large number of local video hub offices (VHO). The super head-ends serve as the national content aggregation points for broadcast and video on demand encoding. The local video hub offices provide aggregation and storage of local content. Each video hub office serves as a Video-On-Demand (VOD) library and distributes video content through local access switches to the customers. We refer to this network hierarchy as the “physical” model throughout the paper. FTTN networks can provide 20-25Mbps bandwidth to each household, which is typically enough to support several high quality TV streams as well as high speed Internet and Voice over IP (VoIP) services.

A significant problem in providing IPTV services is its high deployment and maintenance cost. In addition, the capacity of the video servers can quickly become a bottleneck. One solution to alleviate the load on servers is to use peer-to-peer (P2P) systems like Skype [15] or Kontiki [10]. While early P2P systems were mostly used for file downloading, recently there have been several efforts on using the peer-to-peer approach to support live streaming [16][17][5][2][3][11] and VOD streaming[14][7][13][6]. Existing research studies that evaluate the benefits of P2P video content delivery typically do not consider the constraints of the underlying service infrastructure (e.g., [12][18]). Rather, they view the network as a “cloud”. Researchers, however, are increasingly aware of the need to reduce cross-ISP P2P traffic, while maintaining satisfactory P2P performance[4]. In this paper, we reveal the deficiency of this cloud model and investigate when P2P streaming can be beneficial in an IPTV environment. As we will see, P2P video sharing can be harmful under certain network conditions.

Another challenge for P2P streaming in an IPTV environment is the pricing strategy. Most broadband ISPs today charge a flat fee for providing bandwidth. Usage-based pricing has emerged in some markets but even in those cases it is limited to volume-based pricing. Among the limited early work on pricing strategies for P2P, Adler, et al. [1] provided a comprehensive model applicable to a variety of P2P resource economics. Implementation of peer selection algorithms in realistic networking models like the IPTV environ-




Figure 1: Cloud Model




Figure 2: Physical Model for IPTV Service

ment was not addressed. Hefeeda et al. presented a cost-profit analysis of a P2P streaming service for heterogeneous peers with limited capacity [8]. The analysis shows that the service provider can achieve more profit by providing the appropriate incentives for participating peers. However, their analysis did not consider the bandwidth constraints of the underlying infrastructure and hence cannot be easily extended to our IPTV environment.

We make the following contribution in this paper:

- We compare two network models (the “cloud” model and the “physical” model) and show that the cloud model can dramatically overestimate P2P benefits by a factor of 3 or more.
- We couple three P2P pricing models (flat-fee, usage-based, and built-in) with a “physical” model and study their trade-offs from a profit perspective.

The rest of the paper is organized as follows. We describe the physical network model and constraints for the IPTV system in section 2. Section 3 provides the insights as to why a more accurate physical network model is necessary to realize a profitable IPTV system. Three different pricing models are analyzed and simulated in section 4. Section 5 provides a conclusion and potential future work.

2. NETWORK MODELS

This section contrasts two network models that can be used in studying the performance of P2P video content delivery.

2.1 Cloud Model

Research in P2P streaming typically considers Internet at a logical level[12][18]: it represents the Internet at large as an abstract cloud and only considers the capacity of the content server and the characteristics of the access links to related hosts. We refer this view of the Internet as the “cloud model” as shown in Figure 1.

2.2 Physical Model

In contrast to the cloud model, the physical model considers the network architecture and bandwidth constraints of the underlying links and network devices. In [9], we described and analyzed the physical model of FTTN access networks for IPTV services. The model and analysis can also be applied to xDSL or Cable connections.

As shown in Figure 2, video streaming servers are organized in two levels - a local video hub office (VHO), which consists of a cluster of streaming servers or proxies to serve viewers directly, and national super head end (SHE) offices, which can distribute videos to local serving offices based on existing policies or on demand. We concentrate on video on demand (VOD) in this paper. Each local VHO office (often referred to as “local office” below) connects to a set of access switches such as xDSL, FTTN or Cable switches through optical fiber cables. Each switch connects a *community* of IPTV service customers through twisted-pair copper wires, fibers or coaxial cables. A community consists of all homes which are connected to the same access (xDSL or Cable) switch. A local VHO also includes a service router to connect to a national SHE office. These uplinks (or “north-bound links”) of local offices are implemented over high-speed optical fiber networks.

The following parameters are used throughout the paper:

- B_{0D} : Download bandwidth into a home.
- B_{0U} : Upload bandwidth out of a home.
- B_{1S} : Total capacity of south-bound links (downlinks) of a local access switch.
- B_{1N} : Capacity of the north-bound link (uplink) of an access switch determined by the total bandwidth of north-bound fibers from a switch to a local VHO and the switching capacity of the service router in the VHO.
- B_{2S} : Maximum throughput in a local VHO determined by capacities of service routers, optical network cables and/or streaming servers in the VHO.
- u : Average streaming bit rate for a video.
- N : Maximum number of concurrent viewers supported by a local VHO.

As an example, AT&T LightSpeed network allocates 20 to 25Mbps download bandwidth ($B_{0D} \leq 25Mbps$) and 1Mbps upload bandwidth ($B_{0U} \leq 1Mbps$) to each home. LightSpeed uses an FTTN switch which has a maximum of 24Gbps downlink (or “south-side”) switching capacity ($B_{1S} \leq 24Gbps$). Each FTTN switch can connect an OC-24 fiber to a service router in a local VHO ($B_{1N} \leq 1.244Gbps$). The service router in a local VHO could then connect an OC-192 fiber to national SHE offices. Each high-definition (HD) channel uses 6Mbps bandwidth and each standard-definition (SD) channel uses 2Mbps bandwidth.

2.3 Network Constraints under Physical Model

In a physical network environment, all P2P upload traffic has to traverse through the access switches and service routers that connect the peers. As a result, P2P streaming will increase the load of access switches, local offices and national offices.

Compared with the conventional IPTV services, P2P sharing within a community may not be beneficial if the south-bound link bandwidth of an access switch is the bottleneck. However, P2P sharing within a community decreases the load on the north-bound link of an access switch. Therefore, P2P sharing within a community will have the most benefit if the infrastructure bottleneck is on the north-bound link bandwidth of an access switch.

Similarly, P2P sharing among peers across communities increases the traffic on both the north-bound links and the south-bound links of access switches. If the network bottleneck is in either B_{1N} or B_{1S} , P2P sharing among peers in all communities creates more congestion for the switches and decreases the number of concurrent viewers which can be served by a local office. In this case, P2P sharing across communities is not beneficial for IPTV service providers. Also, if an IPTV service provider can apply content distribution network (CDN) technologies such as caching and replication to reduce the workload in SHE, the benefit of P2P sharing across communities in a VHO is very limited. The detailed analysis of network constraints for P2P IPTV services can be found in [9].

3. NETWORK AT THE PHYSICAL LEVEL

A key insight of this paper is that using the “cloud model” for P2P streaming is over simplistic and misleading. More reliable results can be obtained by considering the network at the physical infrastructure level. To demonstrate our point, consider the following simple P2P algorithm. The content server receives a request for a video, identifies candidate peers with that video and spare upload capacity, and selects a random set among them to collectively serve the video. If not enough candidates are available to serve the video at its encoding rate, the server tries to serve the remaining portion itself, or denies the request if it cannot.

We simulated the performance of the system under the two models. For the physical model, we used a slice of the infrastructure of Figure 2 corresponding to one local office with 20 communities and considered the situation where the content server in the local office distributes video content to the viewers in these communities. For the cloud model, we assume the same content server and viewers are connected via the Internet cloud. We assume the same behavior for every node in the community: an idle user (i.e., the user not viewing a stream already) requests a stream with probability of 2% every time tick. A time tick occurs every minute. A peer may download only one stream at a time. There are 1000 video programs available for viewing. When a peer issues a request, it selects a program according to Zipf’s popularity distribution. Each stream lasts 120 minutes and has a data rate of 6Mbps.¹ Once downloaded, the program remains available at the peer for a period called the stream time-to-live (stream TTL) with a default value of 1000 minutes. A peer may be turned off and on by its user. An


¹The HD stream encoding rate is constantly improving and we expect it to eventually reach 6Mbps.

operational peer is turned off with probability 0.1% on every time tick, and a non-operational peer is turned on with probability 0.5% on every tick. This means that on average every peer stays on five times longer than it stays off. We further assume that $B_{1N} = 0.622$ G (OC-12), and $B_{2S} = 10$ G. Each data point in the graphs throughout the paper is obtained by running the simulation program over 5000 time clicks and taking the average over the last 2500 time ticks (when the system reached a steady state in all the simulations).


The results for the cloud and physical models are shown in Figure 3. The figure also includes curves for the system that does not use P2P delivery under the physical model. Figure 3a shows the average number of concurrent viewers the system can support as the number of peers grows for fixed network and server capacities. The cloud model indicates that P2P delivery allows the system to serve more concurrent viewers and to scale to the growing number of viewers. However, the result is drastically different when the limitations of the physical infrastructure are brought into the picture. In fact, the cloud model could overestimate the benefit by a factor of 2 when there are more than 800 peers in a community as shown in Figure 3a. Not only does the P2P system serve fewer users, it does not scale with a growing number of users and has only a slight capacity advantage over the much simpler centralized delivery (which in fact turns to slight *disadvantage* for other parameter settings as seen in Figures 3b and 3c). The reason behind this drastic change is the limitations of B_{1N} , the links between the local office and individual access switches. When P2P delivery occurs across different communities, two of these links are traversed: one upstream from the serving peer to the local office, and the other downstream from the local office to the receiving peer. Overall, these links are more heavily utilized under P2P delivery and more requests are denied.

Now consider the number of concurrent viewers under varying capacity of the office-to-access-switch link (Figure 3b), when the community size is fixed at 500 viewers. The results for the cloud model are not affected by this link since the model does not consider it. However, the physical model reveals an important trend: the centralized delivery becomes quickly bottlenecked at the server and stops responding to the growing bandwidth of the office-to-access-switch link. On the other hand, with P2P delivery, improvement in this link’s capacity produces a roughly linear growth in the number of concurrent viewers served, at least within the bandwidth range studied.


More differences are seen when we increase the server capacity instead (Figure 3c). In this case, the cloud model quickly reaches the point where it serves all requested streams and stops being affected by the increase in server capacity. In particular, this result might indicate that it is highly beneficial to increase the server capacity from 10 Gbps to 20 Gbps. Under physical model, however, the number of concurrent viewers is unaffected by this change. Thus, the above investment would be useless under the simple algorithm we are considering. Comparing the P2P and centralized delivery under the physical model, the centralized delivery benefits from increased server capacity until it reaches 20 Gbps, after which the bottleneck shifts to the office-to-access-switch link. However, this bottleneck transpires later than in the P2P case. Overall, Figure 3 shows that de-



(a) Concurrent capacity vs. number of users



(b) Concurrent capacity vs. bandwidth of the office-to-access-switch link



(c) Concurrent capacity vs. server capacity

Figure 3: Cloud vs. physical model comparison




Figure 4: Profit per unit time in the no-P2P model and the three incentive models under varying B_{1N} and B_{1S} capacities (this figure should be viewed in color)

pending on whether or not the network operator plans to use P2P delivery, they should focus their investment on the office-to-access-switch link bandwidth or spread it between both server and office-to-access-switch link capacities. These trade-offs cannot be revealed under the conventional cloud model.

4. COST-BENEFIT ANALYSIS

In order to encourage viewers to make their set-top boxes available for P2P sharing, some incentive may be given to peers who upload videos to other peers. This section analyzes the cost and benefit of deploying P2P technology on a physical network and compares its maximum possible profit to that of a conventional IPTV service.

4.1 Maximum Benefit for Conventional IPTV

Let r be the fee paid by a viewer in a time unit (e.g. hours or days) for video streaming services. For conventional IPTV services, the maximum revenue in a local office per time unit is

$$R_{max} = rN$$

where N represents the total number of viewers supported by a local office - with or without P2P incentives.

The maximum profit per time unit, P_{nop2p} , is

$$\begin{aligned} P_{nop2p} &= \text{maximum income} - \text{IPTV expenses} \\ &= rN - E_{nop2p} \end{aligned}$$

where E_{nop2p} is the capital and operation expenses of the IPTV services per time unit.

4.2 P2P Incentive Models

To encourage P2P sharing among viewers, we consider three incentive models: Built-in model, Flat-reward model and Usage-based model.

4.2.1 Built-in Model

In this model, every set-top box includes P2P streaming software by default. Hence, P2P sharing is hidden from the viewers. The maximum profit per time unit is

$$P_b = rN - E_{p2p}$$

where E_{p2p} is the total operation and capital expenses per time unit for providing P2P IPTV services. It should be greater than E_{nop2p} because P2P software needs to be installed on servers and clients and hence will increase the cost of the infrastructure. Let's assume

$$E_{p2p} = E_{nop2p} + A_{p2p}$$

where A_{p2p} includes the additional software license and maintenance fees paid for P2P software and of additional hardware (such as disk storage). In the built-in model, we assume that the recurring software licence and maintenance fees and the amortization of additional hardware results in each set-top box costing t dollars extra per time unit. Therefore, $A_{p2p} = tN$. Then,

$$P_b = rN - E_{nop2p} - tN$$

4.2.2 Flat-reward Model

In this model, a viewer signs up for the video sharing feature for a flat reward. Assume $w\%$ of viewers in a community sign up for video sharing and the reward is d dollars per time unit. The total number of viewers supported by a local office is denoted to be N as before. The maximum cost of incentive per time unit for the office is dwN . Each peer who signs up for the sharing needs to install and activate the P2P software on her set-top box. We assume that a service operator incurs the P2P software license fee only for the set-top boxes which activate the P2P software. Therefore, E_{p2p} equals $E_{nop2p} + twN$. The maximum profit per time unit in this model is

$$\begin{aligned} P_f &= \text{total income} - \text{expenses} - \text{incentive} \\ &= rN - E_{p2p} - dwN \\ &= rN - E_{nop2p} - twN - dwN \end{aligned}$$

In general, w depends on d : increasing d will increase the percentage of viewers willing to share videos and hence increase w .

4.2.3 Usage-based model

In this model, a viewer who signed up for P2P sharing will get credit based on the number of bytes uploaded from its set-top box. Let q be the credit per bit uploaded from a viewer's set-top box and T be the length of a time unit in seconds. The number of bits uploaded from peers for P2P IPTV services in T seconds is $TubN$, where bN is the number of viewers downloading videos from peers among all N viewers in a local office and u is the average video streaming rate. The IPTV service provider gives incentives to these supporting peers based on their contributed bandwidth. In this model, the total reward given by an IPTV service provider to peers in a local office per time unit is $qTubN$. The maximum income per time unit in this model is

$$\begin{aligned} P_s &= rN - E_{p2p} - qbuTN \\ &= rN - E_{nop2p} - tN - qbuTN \end{aligned}$$

As an example to compare the maximum profit per time unit under the conventional, no-P2P model and the three incentive models, we assume that each viewer pays 3 dollars to watch a movie ($r=3$) and each movie lasts about two hours ($T=7200$ seconds). With download bandwidth B_{0D} of 22Mbps, upload bandwidth B_{0U} of 1Mbps, and HDTV streaming rate u of 6Mbps, each HD movie consumes 43.2Gb or 5.4GB and will require six streams from peer nodes for P2P delivery. We further assume that the capital/software/operational cost of each office is \$100 million per year and the additional cost of incorporating P2P software and hardware (disk storage) on each set-top box per time unit is 10 cents. We assume that $B_{2S} = 50$ Gbps. Note that B_{2S} is also constrained by the total streaming throughput from the server, which is about 10Gbps.

We can now plot the profit per unit time for the conventional model vs. various incentive models of VOD services with varying B_{1S} (1-30Gbps) and B_{1N} (1-10Gbps) capacities, as shown in Figure 4. The maximum number of concurrent users are estimated according to a linear optimization program as discussed in [9]. In Figure 4, upper bounds for N are used to illustrate the profit capacity surfaces. Typical values of $w = 0.5$, $t = 0.1$, $q = 0.01$ per Gb, and $d = 0.02$




Figure 5: Profit comparison of different incentive models with varying number of peers




Figure 6: Profit comparisons of the various incentive models under different request rates

were used to estimate these capacities. The profit number ramps up faster for the built-in model (with a given B_{1N}) compared to the no-P2P model as we increase the bandwidth of B_{1S} until it reaches a plateau. Such analysis helps us identify where the focus of the investments should be in increasing the overall profits. Given a fixed B_{1S} beyond certain capacity, ramping up the capacity of B_{1N} appears to add profit to the no-P2P model linearly since it allows more users in each community without relying on the P2P technology. Note that substantial investment may be needed to increase the B_{1N} capacity, which is not reflected in this graph.

4.3 Maximizing Profit Using MediaGrid Algorithm

MediaGrid algorithm [9] is a P2P sharing algorithm which selects peers for streaming or download based on physical network conditions. To study the benefit of P2P technology for an IPTV service provider under various incentive models, we performed an event-driven simulation study using the MediaGrid algorithm, with an additional enhancement allowing the stream delivery to be split between the office and peers (the original algorithm allowed split delivery only in the aftermath of peer failures). Based on the analysis in section 2, which shows that the benefit of P2P sharing among peers in different communities is very limited [9], we only consider P2P sharing within a community and simulate a system comprised of the local office and one community. We use two variations of the MediaGrid P2P sharing algo-

rithm for the simulation:

- The “peer-first” MediaGrid algorithm where peers are selected whenever the requested video can be served by peers.
- The “server-first” MediaGrid algorithm where peers are selected only when the VOD server in a VHO is overloaded.

We assume the same simulation model as described in section 3, using the physical network model. We assume that viewing each movie costs \$3 (even if it is viewed from the local cache), peer incentive in the flat-reward model is 2.5 cents per time unit (120 minutes), and peer incentive in the usage-based model is 1 cent per upload minute. Figure 5 shows the profit numbers under the conventional no-P2P model and the three incentive models for different community sizes. As the number of peers increases, all P2P incentive models clearly generate more profit than the no-P2P model, because of the increased system capacity due to P2P content delivery. However, we see large differences among the incentive models. In fact, the usage-based model underperforms the no-P2P model for small communities because it may utilize (hence, compensate) peers even when the server has spare capacity (even in the server-first algorithm, once a peer starts uploading the stream, it continues to do so for the duration of the movie regardless of the server load). In the usage-based model, the server-first strategy generates more profits for small communities since it avoids making incentive payments to peers whenever possible. As the number of users increases, the server becomes fully utilized in both approaches and their profits converge. Finally, the built-in model always generates more profits than the other incentive models. The reason is that at the request rate used in this experiment (equal to 0.02 or one request every 50 minutes from idle viewers), the system is sufficiently utilized for the built-in model to amortize its investment on the additional hardware and software.

What happens when the system is under-utilized? Figure 6 shows the effect of lower request rates on profit, for a fixed community size of 200 viewers. As Figure 6 reveals, when the request rate is low, no-P2P model is actually slightly more profitable than all P2P models except for the usage-based model with the server-first strategy. The latter stays competitive since very few payments have to be made. Once the request rate picks up, the flat-reward model and the built-in model become more profitable since they enjoy the P2P benefits without making additional payments, while the usage-based models fall behind. In fact, the server-first profits decrease to converge with the peer-first model. This is due to the fact, that as the server utilization increases, the server-first strategy serves additional movies from the peers, and peer incentives for these movies (which can reach \$7.20 for a fully peer-delivered movie under our parameters) can exceed the viewing revenue (\$3).

5. CONCLUSIONS

This paper studied the conditions under which P2P technology may be beneficial in providing IPTV services. We show that the cloud model may drastically overstate the benefits of P2P video content delivery. Thus, one must consider physical network infrastructure to obtain more reliable results. Finally, we provide a cost-benefit analysis for different pricing/incentive models. In summary, P2P may not be beneficial for IPTV services unless we employ properly

engineered algorithms and incentive strategies as discussed in this paper.

6. REFERENCES

- [1] M. Adler, R. Kumar, K. Ross, D. Rubenstein, T. Suel, D. Yao, *Optimal peer selection for p2p downloading and streaming*, Proc. of INFOCOM 2004.
- [2] S. Banerjee, Bobby Bhattacharjee, and C. Kommareddy, *Scalable application layer multicast*, in Proc. of ACM SIGCOMM 2002.
- [3] Mayank Bawa, Hrishikesh Deshpande, Hector Garica-Molina, *Transience of Peers and Streaming Media*, in ACM SIGCOMM Computer Communications Review, January, 2003.
- [4] R. Bindal, P. Cao, W. Chan, J. Medved, G. Suwala, T. Bates, A. Zhang, *Improving Traffic Locality in BitTorrent via Biased Neighbor Selection*, in Proc. of ICDCS, 2006.
- [5] H. Deshpande, M. Bawa and H. Garcia-Molina, *Streaming Live Media over a Peer-to-Peer Network*, Stanford database group technical report (2001-20), Aug. 2001.
- [6] Yang Guo, Kyoungwon Suh, Jim Kurose, and Don Towsley, *A Peer-to-Peer On-Demand Streaming Service and Its Performance Evaluation*, in Proc. of IEEE Int. Conf. on Multimedia Expo (ICME'03), 2003.
- [7] Yang Guo, Kyoungwon Suh, Jim Kurose, Don Towsley *P2Cast: P2P Patching Scheme for VoD Service*, in Proc of 12th WWW, 2003.
- [8] M. Hefeeda and A. Habib and B. Bhargava, *Cost-profit analysis of a peer-to-peer media streaming architecture*, Technical report, CERIAS TR 2002-37, Purdue University, June 2003.
- [9] Y. Huang, Y. Chen, R. Jana, H. Jiang, M. Rabinovich, A. Reibman, B. Wei, and Z. Xiao, *Capacity Analysis of MediaGrid: a P2P IPTV Platform for Fiber to the Node Networks*, in IEEE Journal on Selected Areas in Communications special issue on Peer-to-Peer Communications and Applications, January 2007.
- [10] *Kontiki - www.kontiki.com*
- [11] X. Liao, H. Jin, Y. Liu, L. M. Ni, D. Deng, *AnySec: Peer-to-Peer Live Streaming*, in Proc. of IEEE INFOCOM, 2006.
- [12] V. N. Padmanabhan, H. J. Wang, P. A. Chou, and K. Sripanidkulchai, *Distributing Streaming Media Content Using Cooperative Networking*, ACM NOSSDAV, May 2002.
- [13] Stefan Saroiu, P. Krishna Gummadi, and Steven D. Gribble, *A measurement study of peer-to-peer file sharing systems*, in Proc. of ACM/SPIE on Multimedia Computing and Networking (MMCN'02), January 2002.
- [14] Simon Sheu, Kien A. Hua, Wallapak Tavanapong, *Chaining: A Generalized Batching Technique for Video-on-Demand Systems*, in Proc. of the IEEE Int'l Conf. On Multimedia Computing and System, June 1997.
- [15] *Skype -www.skype.com*
- [16] Duc A. Tran, Kien A. Hua, Tai T. Do, *A Peer-to-Peer Architecture for Media Streaming*, in IEEE Journal on Selected Areas in Communication, Special Issue on Advances in Overlay Network, 2003.
- [17] Eveline Velos, Virgilio Almeida, Wagner Meira, Azer Bestavros, Shudong Jin, *A Hierarchical Characterization of a Live Streaming Media Workload*, in IEEE IMW'02, 2002.
- [18] Dongyan Xu and Sunil Suresh Kulkarni and Catherine Rosenberg and Heung-Keung, *A CDN-P2P Hybrid Architecture for Cost-Effective Streaming Media Distribution*, Computer Networks, Vol. 44, Issue.3, pp. 353-382, 2004.