

A New Method of Live Streaming with Receiver –Based Peer-division Multiplexing for Next Generation Networks

D.Sri Lakshmi^{#1}, V Srikanth^{*2}

^{#1}M.Tech Scholar
Department of CSE ,

^{*2} Associate Professor
Department of CSE
Pydah College Of Engineering,
Visakhapatnam ,AP,India.

Abstract

The behavior of commercial peer-to-peer systems for live streaming has been extensively studied in several measurement papers. These studies have to rely on a “black-box” approach, where packet traces are collected from a single or a limited number of measurement points, to infer various properties of traffic on the control and data planes. Next generation broadcast network concentrates on IPTV (Internet Protocol Television). The main hurdle in IPTV is streaming of audio and video signals. A number of commercial systems are built to study and analyze the behavior of live streaming of audio and video signals. Peer to Peer multiplexing (P2P) provides a good solution for this problem. In this paper a variation of P2P multiplexing is proposed which is called as receiver based P2P multiplexing. To analyze the performance of the proposed multiplexing techniques the very famous European network “Zattoo” is considered. This paper also describes the network architecture of Zattoo and uses the data collected from the provider to evaluate the performance of the proposed variation in P2P multiplexing.

Keywords:

Peer-to-peer system, live streaming, network architecture, IPTV, Multiplexing

1. Introduction

Live television, time-shifted programming, and content-on-demand are all presently available over the Internet. Increased broadband speed, growth of broadband subscription

base, and improved video compression technologies have contributed to the emergence of IPTV services.

Current generation broadcast network for TV is DTH which will be slowly replaced by the next generation Internet Protocol Television (IPTV) network. There is an emerging market for IPTV. Numerous commercial systems now offer services over the Internet that is similar to traditional over-the-air, cable, or satellite TV. Live television, time-shifted programming, and content-on-demand are all presently available over the Internet. Increased broadband speed, growth of broadband subscription base, and improved video compression technologies have contributed to the emergence of these IPTV services [1][6]. IPTV systems deliver video and audio channels to viewing devices by switching a single channel to multiple sources. IP Television networks are primarily constructed of computer servers, gateways, access connections and end user display devices. Servers control the overall system access and processing of channel connection requests and gateways convert the IP television network data to signals that can be used by television media viewers.

Peer-to-Peer (P2P) networks used in three ways: delay-tolerant file download of archival material, delay-sensitive progressive download of archival material, and real-time live streaming.

1. Delay-Tolerant File Download Of Archival Material:

In this case, the completion of download is elastic, depending on available bandwidth in the P2P

network. The application buffer receives data as it trickles in and informs the user upon the completion of download. The user can then start playing back the file for viewing in the case of a video file.

2. Delay-Sensitive Download Of Archival Material:

In this case, video playback starts as soon as the application assesses it has sufficient data buffered that, given the estimated download rate and the playback rate, it will not deplete the buffer before the end of file. If this assessment is wrong, the application would have to either pause playback and rebuffer or slow down playback. While users would like playback to start as soon as possible, the application has some degree of freedom in trading off playback start time against estimated network capacity.

3. Real-Time Live Streaming:

In this case, has the most stringent delay requirement. While progressive download may tolerate initial buffering of tens of seconds or even minutes, live streaming generally cannot tolerate more than a few seconds of buffering. The **Zattoo peer-to-peer live streaming system** was a free-to-use network serving over 3 million registered users in eight European countries at the time of study, with a maximum of over 60000 concurrent users on a single channel. The system delivers live streams using a receiver-based, peer-division multiplexing scheme.

2. System Architecture

The Zattoo system rebroadcasts live TV, captured from satellites, onto the Internet. The system carries each TV channel on a separate peer-to-peer delivery network and is not limited in the number of TV channels it can carry. Although a peer can freely switch from one TV channel to another, and thereby departing and joining different peer-to-peer networks, it can only join one peer-to-peer network at any one time. We henceforth limit our description of the Zattoo delivery network as it pertains to carrying one TV channel. Fig.1 shows a typical setup of a single TV channel carried on the

Zattoo network. TV signal captured from satellite is encoded into H.264/AAC streams, encrypted, and sent onto the Zattoo network.

The encoding server may be physically separated from the server delivering the encoded content onto the Zattoo network. For ease of exposition, we will consider the two as logically co-located on an Encoding Server. Users are required to register themselves at the Zattoo website to download a free copy of the Zattoo player application. To receive the signal of a channel, the user first authenticates itself to the Zattoo Authentication Server. Upon authentication, the user is granted a ticket with limited lifetime. The user then presents this ticket, along with the identity of the TV channel of interest, to the Zattoo Rendezvous Server. If the ticket specifies that the user is authorized to receive signal of the said TV channel, the Rendezvous Server returns to the user a list of peers currently joined to the P2P network carrying the channel, together with a signed channel ticket. If the user is the first peer to join a channel, the list of peers it receives contains only the Encoding Server. The user joins the channel by contacting the peers returned by the Rendezvous Server, presenting its channel ticket, and obtaining the live stream of the channel from them

2.1 Peer-Division Multiplexing

To minimize per-packet processing time of a stream, the Zattoo protocol sets up a virtual circuit with multiple fan outs at each peer. When a peer joins a TV channel, it establishes a peer-division multiplexing (PDM) scheme amongst a set of neighboring peers[4], by building a virtual circuit to each of the neighboring peers. Barring departure or performance degradation of a neighbor peer, the virtual circuits is maintained until the joining peer switches to another TV channel. With the virtual circuits set up, each packet is forwarded without further per-packet handshaking between peers.

The PDM establishment process consists of two phases: the search phase and the join phase. In the search phase, the new, joining peer determines its set of potential neighbors. In the join phase, the joining peer requests peering relationships with a subset of its potential neighbors. Upon acceptance of a peering relationship request, the peers become

neighbors and a virtual circuit is formed between them.




Fig. 1. Zattoo delivery network architecture.

2.2 Stream Management

We represent a peer as a packet buffer, called the IOB, fed by sub-streams incoming from the PDM. The IOB drains to (1) a local media player if one is running, (2) a local file if recording is supported, and (3) potentially other peers. Fig. 2 depicts a Zattoo player application with virtual circuits established to four peers. As packets from each sub-stream arrive at the peer, they are stored in the IOB for reassembly to reconstruct the full stream. Portions of the stream that have been reconstructed are then played back to the user. In addition to providing a reassembly area, the IOB also allows a peer to absorb some variability in available network bandwidth and network delay.

The IOB is referenced by an input pointer, a repair pointer, and one or more output pointers. The input pointer points to the slot in the IOB where the next incoming packet with sequence number higher than the highest sequence number received so

far will be stored. The repair pointer always points one slot beyond the last packet received in order and is used to regulate packet retransmission and adaptive PDM [3] as described later.




Fig 2. Zattoo/Peer system with an IOB

We assign an output pointer to each forwarding destination. The output pointer of a destination indicates the destination's current forwarding horizon on the IOB. In accordance to the three types of possible forwarding destinations listed above, we have three types of output pointers: player pointer, file pointer, and peer pointer. One would typically have at most one player pointer and one file pointer but potentially multiple concurrent peer pointers, referencing an IOB. The Zattoo player application does not currently support recording since we maintain the IOB as a circular buffer, if the Incoming packet rate is higher than the forwarding rate of a particular destination, the input pointer will overrun the output pointer of that destination. We could move the output pointer to match the input pointer so that we consistently forward the oldest packet in the IOB to the destination. Doing so, however, requires checking the input pointer against all output pointers on every packet arrival. Instead, we have implemented the IOB as a double buffer.

With the double buffer, the positions of the output pointers are checked against that of the input pointer only when the input pointer moves from one sub-buffer to the other. When the input pointer moves from sub-buffer a to sub buffer b, all the output pointers still pointing to sub-buffer b are moved to the start of sub-buffer a and sub-buffer b is flushed, ready to accept new packets. When a sub-buffer is flushed while there are still output pointers referencing it, packets that have not been forwarded to the destinations associate with those pointers are lost to them, resulting in quality degradation. To minimize packet lost due to sub-buffer flushing, we would like to use large sub-buffers.

However, the real-time delay requirement of live streaming limits the usefulness of late arriving packets and effectively puts a cap on the maximum size of the sub-buffers. Different peers may request for different numbers of, possibly non-consecutive, sub-streams. To accommodate the different forwarding rates and regimes required by the destinations, we associate a packet map and forwarding discipline with each output pointer. Fig. 3 shows the packet map associated with an output peer pointer where the peer has requested sub-streams 1, 4, 9, and 14. Every time a peer pointer is repositioned to the beginning of a sub-buffer of the IOB, all the packet slots of the requested sub-streams are marked NEEDED and all the slots of the sub-streams not requested by the peer are marked SKIP. When a NEEDED packet arrives and is stored in the IOB, its state in the packet map is changed to READY. As the peer pointer moves along its associated packet map, READY packets are forwarded to the peer and their states changed to SENT. A slot marked NEEDED but not READY, such as slot $n + 4$ in Fig.3, indicates that the packet is lost or will arrive out-of-order and is bypassed. When an out-of-order packet arrives, its slot is changed to READY and the peer pointer is reset to point to this slot. Once the out-of-order packet has been sent to the peer, the peer pointer will move forward, bypassing all SKIP, NEED, and SENT slots until it reaches the next READY slot, where it can resume sending.

Fig.4 shows an IOB consisting of a double buffer, with an input pointer, a repair pointer, and an output file pointer, an output player pointer, and two

output peer pointers referencing the IOB. Each output pointer has a packet map associated with it. For the scenario depicted in the figure, the player pointer tracks the input pointer and has skipped over some lost packets. Both peer pointers are lagging the input pointer, indicating that the forwarding rates to the peers are bandwidth limited. The file pointer is pointing at the first lost packet. Archiving a live stream to file does not impose real-time delay bound on packet arrivals. To achieve the best quality recording possible, a recording peer always waits for retransmission of lost packets that cannot be recovered by error correction.




Fig. 3. Packet map associated with a peer pointer.




Fig. 4. IOB, input/output pointers and packet maps.

2.3. Adaptive PDM

While we rely on packet retransmission to recover from transient congestions, we have two channel capacity adjustment mechanisms to handle longer-term bandwidth fluctuations. The first mechanism allows a forwarding peer to adapt the number of sub-streams it will forward given its current available bandwidth, while the second allows the receiving peer to switch provider at the sub-stream level. Peers on the Zattoo network can redistribute a highly variable number of sub-streams, reflecting the high variability in uplink bandwidth of different access network technologies. For a full-stream consisting of sixteen constant-bit rate sub streams, our prior study show that based on realistic peer characteristics measured from the Zattoo network, half of the peers can support less than half of a stream, 82% of peers can support less than a full-stream, and the remainder can support up to ten full streams (peers that can redistribute more than a full stream is conventionally known as super nodes in the literature) [5].

With variable-bit rate streams, the bandwidth carried by each sub-stream is also variable. To increase peer bandwidth usage, without undue degradation of service, we instituted measurement-based admission control at each peer. In addition to controlling resource commitment, another goal of the measurement-based admission control module is to continually estimate the amount of available uplink bandwidth at a peer.

3. Global Bandwidth Subsidy System

Each peer on the Zattoo network is assumed to serve a user through a media player, which means that each peer must receive, and can potentially forward, all n sub-streams of the TV channel the user is watching. The limited redistribution capacity of peers on the Zattoo network means that a typical client can contribute only a fraction of the sub streams that make up a channel. This shortage of bandwidth leads to a global bandwidth deficit in the peer-to-peer network Whereas bit torrent like delay-tolerant file downloads or the delay-sensitive progressive download of video-on-demand applications can

mitigate such global bandwidth shortage by increasing download time, a live streaming system such as Zattoo's must subsidize the bandwidth shortfall to provide real-time delivery guarantee.

Zattoo's Global Bandwidth Subsidy System (**or simply, the Subsidy System**), consists of a global bandwidth monitoring subsystem, a global bandwidth forecasting and provisioning subsystem, and a pool of Repeater nodes. The monitoring subsystem continuously monitors the global bandwidth requirement of a channel. The forecasting and provisioning sub system projects global bandwidth requirement based on measured history and allocates Repeater nodes to the channel as needed. The monitoring and provisioning of global bandwidth is complicated by two highly varying parameters over time, client population size and peak streaming rate, and one varying parameter over space, available uplink bandwidth, which is network-service provider dependent. Forecasting of bandwidth requirement is a vast subject in itself.

When a bandwidth shortage is projected for a channel, the Subsidy System assigns one or more Repeater nodes to the channel. Repeater nodes function as bandwidth multiplier, to amplify the amount of available bandwidth in the network. Each Repeater node serves at most one channel at a time; it joins and leaves a given channel at the behest of the Subsidy System. Repeater nodes receive and serve all n sub-streams of the channel they join, run the same PDM protocol, and are treated by actual peers like any other peers on the network; however, as bandwidth amplifiers, they are usually provisioned to contribute more uplink bandwidth than the download bandwidth they consume. The use of Repeater nodes makes the Zattoo network a hybrid P2P and content distribution network. We next describe the bandwidth monitoring subsystem of the Subsidy System, followed by design of the simple bandwidth projection and Repeater node assignment subsystem.

3.1 Global Bandwidth Measurement

The capacity metric of a channel is the tuple (D, C) , where D is the aggregate download rates required by all users on the channel, and C is

the aggregate upload capacity of those users. Usually $C < D$ and the difference between the two is the global bandwidth deficit of the channel. Since channel viewership changes over time as users join or leave the channel, we need a scalable means to measure and update the capacity metric. We rely on aggregating the capacity metric up the peer division multiplexing tree. Each peer in the overlay periodically aggregates the capacity metric reported by all its downstream receiver peers, adds its own capacity measure (D,C) to the aggregate, and forwards the resulting capacity metric upstream to its forwarding peers. By the time the capacity metric percolates up to the Encoding Server, it contains the total download and uploads rate aggregates of the whole streaming overlay. The Encoding Server then simply forwards the obtained (D, C) to the Subsidy Server.

3.2 Global Bandwidth Projection and Provisioning

For each channel, the Subsidy Server keeps a history of the capacity metric (D,C) reports received from the channel's Encoding Server. The channel utilization ratio (U) is the ratio D over C. Based on recent movements of the ratio U, we classify the capacity trend of each channel into the following four categories.

Stable: the ratio U has remained within $[S_-, S_+]$ for the past T_s reporting periods.

Exploding: the ratio U increased by at least E between T_e (e.g., $T_e = 2$) reporting periods.

Increasing: the ratio U has steadily increased by I ($I \geq E$) over the past T_i reporting periods.

Falling: the ratio U has decreased by F over the past T_f reporting period.

Orthogonal to the capacity-trend based classification above, each channel is further categorized in terms of its capacity utilization ratio as follows.

Under-utilized: the ratio U is below the low threshold L, e.g., $U_{0.5}$.

Near Capacity: the ratio U is almost 1.0, e.g., $U_{0.9}$. If the capacity trend of a channel has been "Exploding," one or more Repeater nodes will be assigned to it immediately.

If the channel's capacity trend has been "Increasing," it will be assigned a Repeater node with a smaller capacity. The goal of the Subsidy System is to keep a channel below "Near Capacity." If a channel's capacity trend is "Stable" and the channel is "Under-utilized," the Subsidy System attempts to free Repeater nodes (if any) assigned to the channel. If a channel's capacity utilization is "Falling," the Subsidy System waits for the utilization to stabilize before reassigning any Repeater nodes.

In addition to Repeater nodes, the Subsidy System may recognize extra capacity from idle peers whose owners are not actively watching a channel. However, our previous study shows that a large number of idle peers are required to make any discernible impact on the global bandwidth deficit of a channel [5]. Our current Subsidy System therefore does not solicit bandwidth contribution from idle peers.

4. Server-Side Measurements

In the Zattoo system, two separate centralized collector servers collect usage statistics and error reports, which we call the "stats" server and the "user-feedback" server respectively. The "stats" server periodically collects aggregated player statistics from individual peers, from which full session logs are constructed and entered into a session database. The session database gives a complete picture of all past and present sessions served by the Zattoo system. A given database entry contains statistics about a particular session, which includes join time, leave time, uplink bytes, download bytes, and channel name associated with the session. We study the sessions generated on three major TV channels from three different countries (Germany, Spain, and Switzerland), from June 1st to June 30th, 2008. We label those channels from Germany, Spain, and Switzerland as ARD, Cuatro, and SF2, respectively. Euro 2008 games were held during this period, and those three

channels broadcast a majority of the Euro 2008 games including the final match. See Table I for information about the collected session data sets.

TABLE I
SESSION DATABASE (6/1/2008–6/30/2008).

Channel	# sessions	# distinct users
ARD	2,102,638	298,601
Cuatro	1,845,843	268,522
SF2	1,425,285	157,639

TABLE II
FEEDBACK LOGS (6/20/2008–6/29/2008).

Channel	# feedback logs	# sessions
ARD	871	1,253
Cuatro	2,922	4,568
SF2	656	1,140

The “user-feedback” server, on the other hand, collects users’ error logs submitted asynchronously by users. The “user feedback” data here is different from peer’s quality feedback used in PDM reconfiguration. Table II describes the feedback logs collected from June 20th to June 29th. A given feedback log can contain multiple sessions (for the same or different channels), depending on user’s viewing behavior. The second column in the table represents the number of feedback logs which contain at least one session generated on the channel listed in the corresponding entry in the first column. The numbers in the third column indicate how many distinct sessions generated on said channel are present in the feedback logs.

TABLE III
AVERAGE SHARING RATIO.

Channel	Average sharing ratio	
	Off-peak	Peak
ARD	0.335	0.313
Cuatro	0.242	0.224
SF2	0.277	0.222

Table III shows the average sharing ratio in the two time periods separately. Zattoo’s usage during peak hours typically accounts for about 50% to 70% of the total usage of the day. According to the table, the average sharing ratio during peak hours is slightly lower than, but not very much different from during off-peak hours. Cuatro channel in Spain exhibits relatively lower sharing ratio than the two other channels. One bandwidth test site [6] reports that average uplink bandwidth in

Spain is about 205 kbps, which is much lower than in Germany (582 kbps) and Switzerland (787 kbps). The lower sharing ratio on the Spanish channel may reflect regional difference in residential access network provisioning. The balance of the required bandwidth is provided by Zattoo’s Encoding Server and Repeater nodes.

5. Client-Side Measurements

To further study the P2P overlay beyond details obtainable from aggregated session-level statistics, we run several modified Zattoo clients which periodically retrieve the internal states of other participating peers in the network by exchanging SEARCH/JOIN messages with them. After a given probe session is over, the monitoring client archives a log file where we can analyze control/data traffic exchanged and detailed protocol behavior. We run the experiment during Zattoo’s live coverage of Euro 2008 (June 7th to 29th). The monitoring clients tuned to game channels from one of Zattoo’s data centres located in Switzerland while the games were broadcast live. The data sets presented in this paper were collected during the coverage of the championship final on two separate channels: ARD in Germany and Cuatro in Spain. Soccer teams from Germany and Spain participated in the championship final.

5.1 Sub-Stream Synchrony

To ensure good viewing quality, peer should not only obtain all necessary sub-streams (discounting redundant sub streams), but also have those sub-streams delivered temporally synchronized with each other for proper online decoding. Receiving out-of-sync sub-streams typically results in pixilated screen on the player.

5.2 Peer Synchrony

While sub-stream synchrony tells us stream quality different peers may experience, “peer synchrony” tells us how varied in time peers’ viewing points are. With small scale P2P networks, all participating peers are likely to watch live streaming roughly synchronized in time. However, as the size of the P2P overlay grows, the viewing

point of edge nodes may be delayed significantly compared to those closer to the Encoding Server. In the experiment, we define the viewing point of a peer as the median of the latest sequence numbers across its sub-streams. Then we choose one peer (e.g., a Repeater node directly connected to the Encoding Server) as a reference point, and compare other peers' viewing point against the reference viewing point.

6. Conclusion

In this paper, we have presented a receiver-based, peer-division multiplexing engine to deliver live streaming content on a peer-to-peer network. The same engine can be used to transparently build a hybrid P2P/CDN delivery network by adding Repeater nodes to the network. By analyzing large amount of usage data collected on the network during one of the largest viewing event in Europe, we have shown that the resulting network can scale to a large number of users and can take good advantage of available uplink bandwidth at peers. We have also shown that error-correcting code and packet retransmission can help improve network stability by isolating packet losses and preventing transient congestion from resulting in PDM reconfigurations. We have further shown that the PDM and adaptive PDM schemes presented have small enough overhead to make our system competitive to digital satellite TV in terms of channel switch time, stream synchronization, and signal lag.

7. References

- [1] R. auf der Maur, "Die Weiterverbreitung von TV- und Radioprogrammen über IP-basierte Netze," in *Entertainment Law* (f. d. Schweiz), 1st ed. Stämpfli Verlag, 2006.
- [2] H. Fujisawa, "Forty Years of Research in Character and Document Recognition—An Industrial Perspective," *Pattern Recognition*, vol. 41, no. 8, pp. 2435-2446, Aug. 2008.
- [3] S. Xie, B. Li, G. Y. Keung, and X. Zhang, "CoolStreaming: Design, theory, and practice," *IEEE Trans. Multimedia*, vol. 9, no. 8, pp. 1661-1671, Dec. 2007.

[4] K. Shami et al., "Impacts of peer characteristics on P2PTV networks scalability," in *Proc. IEEE INFOCOM*, Apr. 2009, pp. 2736-2740.

[5] K. Shami et al., "Impacts of Peer Characteristics on P2PTV Networks Scalability," in *Proc. IEEE INFOCOM Mini Conference*, April 2009.

[6] X. Hei, C. Liang, J. Liang, Y. Liu, and K. W. Ross, "Insights into PPLive: A measurement study of a large-scale P2P IPTV system," in *Proc. IPTV Workshop, Int. World Wide Web Conf.*, May 2006.

8. About the Authors



D. Sri Lakshmi is currently pursuing her M.Tech (CSE) in Computer Science & Engineering Department, Pydah College of Engineering, Visakhapatnam. Her area of interests

includes Networks, Security.



V. Srikanth is currently working as an Associate Professor in Computer Science & Engineering Department, Pydah College of Engineering, Visakhapatnam. His research interests include

Data Mining & Warehousing, Networks and Security, Image Processing.