

# Scalable live video streaming to cooperative clients using time shifting and video patching

Meng Guo, Mostafa H. Ammar  
{mguo, ammar}@cc.gatech.edu

Networking and Telecommunication Group  
College of Computing, Georgia Institute of Technology

**Abstract**—We consider the problem of how to enable the streaming of live video content from a single server to a large number of clients. One recently proposed approach relies on the cooperation of the video clients in forming an application layer multicast tree over which the video is propagated. Video continuity is maintained as client departures disrupt the multicast tree, using multiple description coded (MDC) streams multicast over several application layer trees. While this maintains continuity, it can cause video quality fluctuation as clients depart and trees are reconstructed around them. In this paper we develop a scheme using the transmission of a single-description coded video over an application layer multicast tree formed by cooperative clients. Video continuity is maintained in spite of tree disruption caused by departing clients using a combination of two techniques: 1) providing time-shifted streams at the server and allowing clients that suffer service disconnection to join a video channel of the time-shifted stream, and 2) using video patching to allow a client to catch up with the progress of a video program. Simulation experiments demonstrate that our design can achieve uninterrupted service, while not compromising the video quality, at moderate cost.

## I. INTRODUCTION

We consider the problem of how to enable *live streaming* of video content from a single server to a large number of clients. Due to the bandwidth-intensive nature of video streams, deploying scalable streaming service with acceptable quality has always been a challenge. The straightforward solution of setting up a connection for every single request is obviously not scalable. The server side link can be easily congested as the client requests accumulate. Native network layer IP multicast could be a good solution for scalable media streaming applications, but it is not widely deployed.

One recently proposed approach relies on the cooperation of the video clients in forming an overlay network over which the video is propagated [10], [13], [15]. In this approach, a client currently in the overlay network forwards the content it is receiving, and serves other client's request as a server. By distributing the transmission load evenly to the clients all over the network, the video server is no longer the bottleneck. This approach is scalable in the sense that the forwarding capability of the overlay network is growing incrementally. New clients joining the network also bring in extra bandwidth capacity to the system.

The major problem for application layer multicast is the video discontinuity caused by the dynamics of membership. Since the clients in the group can leave at any time, other clients which are receiving video content from them have to suffer service disconnection. Although the disconnected clients can resume the service by rejoining the application layer multicast tree, this process can take time and can result in loss of video content and interruption of video reception. To make things worse, unsatisfied clients leaving the group can become a positive feedback process, causing more clients to leave, which ultimately makes the streaming media service unacceptable. In CoopNet [13], video continuity is maintained using multiple description coded (MDC) streams multicast over several application layer trees. CoopNet employs a server-based centralized control protocol, the server is responsible to make sure that multiple trees for different descriptions are uncorrelated with each other. A single client's departure can only disconnect a subset of descriptions for its children.

While CoopNet maintains video continuity, it can cause video quality fluctuation as clients depart and trees are reconstructed around them. Recent work in CoopNet shows the PSNR variation with different number of descriptions [14]. Our solution tries to provide continuous streaming service without video quality fluctuation. In this paper we develop a scheme using the transmission of a *single description coded* video over an application layer multicast tree formed by cooperative clients. Clients in the tree always receive full quality video stream. Video continuity is maintained in spite of tree disruption caused by departing clients using a combination of two techniques: (1) providing time-shifted streams at the server and allowing clients that suffer service disconnection to join a video channel of the time-shifted stream, and (2) using video patching to allow a client to catch up with the progress of the live video program. We also need a buffering scheme that allows the client to store the video packets, and to playout when needed.

In our design, the client buffers the initial part of the video content for a certain time before play out. If the client is disconnected from the multicast tree, it can play out full quality video from this buffer while reconnecting to the group. The client rejoins the group by connecting to both the original stream and a time-shifted stream. The time-shifted stream (patching stream) is used to retrieve the missed video. The original stream is used to catch up with the progress of the

live video. Received video packets that are not being played out are stored at the buffer for future playout. When the missed portion of the video content is fixed, the time-shifted stream is released, and the client receives from the original stream only. The use of video patching requires that a client should have enough bandwidth for two video streams: the original stream, and the time-shifted stream.

This paper is organized as follows. We give an overview of our design in Section 2. In Section 3, we describe our design in detail, Then, in Section 4 we set up the simulation environment, and show some sample results of our performance evaluation experiments. Finally, we conclude this paper in Section 5.

## II. DESIGN OVERVIEW

In this section, we first give a general description of how our system operates. Then, we discuss the potential problem of the straightforward time shifting solution. We then apply video patching in live streaming service, and describe how the design goals are met. Finally, we give an example to illustrate the system operations.

### A. Basic Operations

Our design of the system is composed of three components: 1) a time shifting video server, 2) a level-based tree management protocol, and 3) a video stream patching scheme.

A time shifting video server  $S$  broadcasts video program in  $C$  channels. Each channel can be used to transmit one video stream. There is an application layer multicast tree associated with each channel. The server serves the clients with the original stream, and  $m$  time-shifted streams. We label these streams  $s_0, s_1, \dots, s_m$ . Stream  $s_0$  is the original stream, while  $s_i$  starts after a  $i \times d$  delay. Video server is the single source of the video content, and is the root of the application layer multicast tree. It processes client requests to join, leave, and rejoin the multicast group, and is responsible for maintaining the topological structure and resource availability of the multicast tree.

When a client first joins the multicast group, it always joins a multicast tree of the original stream. If the server has free video channel available, the client connects to the server directly. Otherwise, the client joins the tree by connecting to a client already in the tree who has enough available bandwidth resources, while at the same time, has the shortest overlay path to the video server. This node join protocol guarantees that the clients in the upper level of the tree are fully loaded, before the clients in the lower level of the tree start to adopt new clients as their children. In this way, we can get a “well-shaped” wide and short multicast tree. A wide and short tree can achieve lower bandwidth consumption, and can reduce the probability of service disconnection due to ancestor node failure.

A client in the multicast tree suffers service disconnection in two cases: 1) upstream link congestion, or 2) an ancestor node’s failure. Similar to CoopNet, to detect service disconnection, the client sets a threshold value, if the packet loss

rate is above the threshold, the client deems it as a service disconnection. For the case of ancestor node failure, the client detects 100% packet loss. The client manages to rejoin the group by connecting itself to another parent node. The *node rejoin delay* for a client is the time interval between the moment when the client is disconnected and the moment when the client is reconnected. We denote the *node rejoin delay* for client  $c$  as  $r_c$ . A straightforward approach for lossless video reception is: when the client rejoins the tree, it can select to join the video channel of an appropriate time-shifted video stream, so that it will not miss any video content. For example, if at time  $t_0$ , the client is disconnected, and it manages to rejoin the group at  $t_0 + r_c$ , the appropriate stream should be the stream with  $\lceil \frac{r_c}{d} \rceil$  delay. Clients that join the same video channel form an application layer multicast tree.

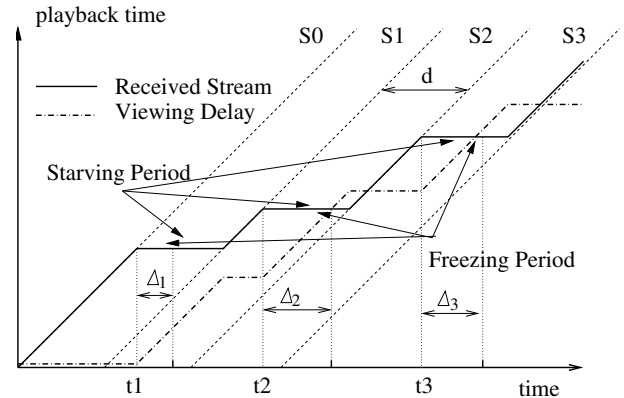


Fig. 1. An example of indefinite shifting

A client might experience multiple service disconnections during the video reception process. Figure 1 shows an example when client  $c$  suffers multiple disconnections. In this figure, the horizontal axis is the actual time, while the vertical axis is the video playback time. The server sends out video streams with equal time shifting interval  $d$ . The client experiences three service disconnections at time  $t_1, t_2$ , and  $t_3$ . *Viewing delay* means the delay between the playback time of the video stream that the client is watching and the playback time of the original stream. *Starving period* is the time interval when the client is not receiving any video. *Freezing period* refers to the time period when the client side play out is temporarily stopped.

The client joins the original stream at time 0. During  $[0, t_1]$ , the viewing delay is 0. At time  $t_1$ , it is disconnected from the application layer multicast tree. It manages to reconnect to the tree at time  $t_1 + \Delta_1$ , the missed video is  $[t_1, t_1 + \Delta_1]$ . The designated time-shifted stream is  $s_i$ , where  $i = \lceil \frac{\Delta_1}{d} \rceil = 1$  (assume  $\Delta_1 < d$ ). The client begins to receive from stream  $s_1$  at time  $t_1 + d$ . The viewing delay during  $[t_1, t_1 + d]$  increases from 0 to  $d$ . The client is disconnected again at time  $t_2$ , and rejoins the tree at  $t_2 + \Delta_2$ . Note that at this time, the missed video portion is  $[t_2 - \lceil \frac{\Delta_1}{d} \rceil \times d, t_2 + \Delta_2]$ . The designated time-shifted stream should be  $s_j$ , where  $j = \lceil \frac{\Delta_1}{d} \rceil + \lceil \frac{\Delta_2}{d} \rceil = 2$ . (again, we assume  $\Delta_2 < d$ ). The client begins to receive from stream  $s_2$  at  $t_2 + d$ , and the viewing delay is increased to

$2 \times d$ . Similar event occurs after the third disconnection. In a general case, if a client suffers  $n$  service disconnections, it has to switch to stream  $s_k$ , where  $k = \sum_{i=1}^n \lceil \frac{\Delta_i}{d} \rceil$ , and the client viewing delay will be  $k \times d$ . Here,  $\Delta_i$  denotes the rejoin delay after the  $i_{th}$  disconnection.

As shown in the Figure 1, each time the client rejoins the multicast tree, it has to join a stream with a larger time shifting value. This can result in *indefinite time shifting*, which is undesirable since the client's reception time could be much longer than the actual video program time. Also, if the time shifting value exceeds the boundary that the server can provide ( $m \times d$ ), the client will suffer video content loss. Another problem demonstrated in this example is *video freezing* caused by *video starvation*. In this example, the *starving period* and *freezing period* is same. Whenever the client is disconnected from the application layer multicast tree, the client side's playback is paused. In our design, we introduce *video patching* to tackle the *indefinite time shifting*; and use *initial buffering* to provide continuous video streaming.

### B. Video Patching in Live Streaming

Video patching is a popular channel allocation technique in Video on Demand (VoD) service. It is designed to provide better server channel utilization and thus lower server load. In video patching, video channels of the VoD server are classified in two categories: regular channels and patching channels. Regular channels transmit the entire video stream, while patching channels only transmit the initial portion of a video as needed. When the server allocates a free channel to client requests, it first scans the on-going channels. If there is no regular channel distributing the requested video, or the starting time for this video is too early for the client to patch, this channel is allocated as a regular channel. Otherwise, the channel is allocated as a patching channel. Under video patching, the clients receive data from both the regular channel and the patching channel. The data from the patching channel is used to make up for the data the clients are missing from the regular channel. While receiving from the patching channel, the data being received from the regular channel is buffered by the clients for playout when needed. When the missing portion is completely received, the patching channel is released, and the clients only receive packets from the regular channel. More detailed description of video patching can be seen in [9].

In our scheme, the server sends out the original stream, as well as multiple time-shifted streams spaced by a fixed time interval  $d$ . The time-shifted stream serves as patching stream in our system. Here is how it works: when the connection to the tree is re-established after a service disconnection, a client would have missed the video from the point it is disconnected to the point it is reconnected. The reconnected client utilizes a time-shifted stream as patching stream. The patching stream is used to retrieve the missed video portion. At the same time, this client also receives the original stream, this stream is used to catch up the progress of the video program. During the

patching period<sup>1</sup>, the client is actually receiving at twice the speed as the normal stream rate. After the progress of the video is caught up, the patching stream is released, and the client receives only from the original stream.

There are several major differences between video patching in VoD service, and video patching in our scheme. 1) Different purposes: video patching in VoD service is used to reduce the access latency, while video patching in our scheme is used to provide lossless video reception. 2) Different starting time: video patching in VoD service starts at the beginning of the service, in our scheme, video patching could happen at anytime during the video reception process. 3) Different releasing time: in VoD service, the patching channel is released when the video playback difference with another regular channel is fixed, while in our service, the patching channel is released when the missed video is retrieved.

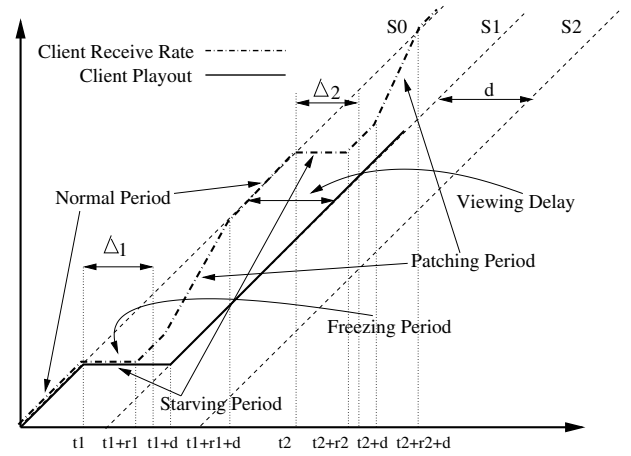


Fig. 2. Video patching in live streaming

We illustrate how the video patching scheme works, and how it eliminates *indefinite time shifting* through an example shown in Figure 2. Assume at time  $t_1$ , the client is disconnected. It rejoins the group by sending out two “rejoin” signals to the server. One of them is for the original stream, the other is for the patching stream. The client rejoins the original stream at time  $t_1 + r_1$ , the missed video portion is  $[t_1, t_1 + r_1]$ . It rejoins the patching stream at time  $t_1 + \Delta_1$  (here, we assume  $\Delta_1 < d$ ), and begins to receive the patching stream  $s_1$  at time  $t_1 + d$ . At time  $t_1 + r_1 + d$ , the missed video portion is made up, and the patching channel is released. From this time, the client receives only from the original stream. Assume at time  $t_2$ , the client is disconnected again. At this moment, the client's playing out time is  $t_1 - d$ , and its buffer stores the video portion of  $[t_1 - d, t_1]$ . The node rejoins the original stream at time  $t_2 + r_2$ , and the missed video is  $[t_2, t_2 + r_2]$ . It rejoins the patching stream at time  $t_2 + \Delta_2$ . We analyze two cases based on the value of  $\Delta_2$ . In the first case where  $\Delta_2 \leq d$ , the client starts to receive patching stream at  $t_2 + d$ .

<sup>1</sup>patching period denotes the time when a client is receiving both the original stream and the time-shifted stream

During  $[t_2, t_2 + d]$ , the client is playing out from the buffer, the client's viewing delay remains unchanged. When the client rejoins the tree, it still receives the patching stream from  $s_1$ , instead of the other streams with longer time shifting values. In the second case, where  $\Delta_2 > d$ , say  $(i-1) \times d < \Delta_2 < i \times d$ ,  $i > 1$ . The client starts to receive patching stream at  $t_2 + i \times d$ . During  $[t_2, t_2 + d]$ , the client is playing out from the buffer, the buffer is exhausted at  $t_2 + d$ . When the client rejoins the tree at  $t_2 + i \times d$ , it receives the patching stream from  $s_i$ , and the client viewing delay is  $i \times d$ . In a general case, if a client suffers the  $n_{th}$  disconnection, the client viewing delay should be determined by the formula below:

$$d_n = \begin{cases} \lceil \frac{\Delta_n}{d} \rceil \times d & n = 1 \\ d_{n-1} + \lceil \frac{\max(\Delta_n - d_{n-1}, 0)}{d} \rceil \times d & n > 1 \end{cases}$$

Solving this formula, we get:

$$d_n = \max(\lceil \frac{\Delta_i}{d} \rceil) \times d \quad i \in [1..n]$$

The client side viewing delay is determined by the maximum value of the *node rejoin delay*. Thus, the *indefinite time shifting* can be eliminated.

### C. Continuous Video Streaming

As we stated in previous sections, the combination of a time shifting server and a video patching scheme can achieve lossless streaming, and prevent *indefinite time shifting*. But there are still some periods that the client's playback is halted, when the client is temporarily disconnected from the multicast tree. This typically happens when the client is first disconnected, when the client buffer is empty; or when the client's rejoin time is longer than the playback time of the buffered video.

To avoid interruption of play out during the *starving period*, buffered video accumulated during an initial playout delay can be used. In our design, when the client first joins the multicast tree, instead of immediately playing out the video stream, it can buffer the initial part of the video for a certain time. This time interval is called *initial access delay*. By waiting for an appropriate *initial access delay*, when there is a service disconnection in the future, the client can still playout the video from its buffer instead of stopping the playout.

We illustrate our solution in Figure 3. The client connects to the server at time 0, instead of playing out immediately, it buffers the initial  $D$  time unit video. At time  $D$ , it begins to play out the video from its buffer, while it keeps receiving from the original stream. As to how large this  $D$  should be, there is a tradeoff between access latency and video continuity. Obviously, longer access latency can result in smoother video reception. In the case of time shifting without video patching, the initial delay should be  $D \geq \sum_{i=1}^n \lceil \frac{\Delta_i}{d} \rceil \times d$ . If video patching is introduced, the initial delay can be reduced to  $D \geq \max(\lceil \frac{\Delta_i}{d} \rceil) \times d$  under non overlapping failures. For detailed analysis, please refer to our technical report [12].

In Figure 3, at time  $t_0$ , the client is disconnected. Since the client buffer stores video  $[t_0 - D, t_0]$ , it can still play from the buffer while reconnecting to the server. If the initial delay

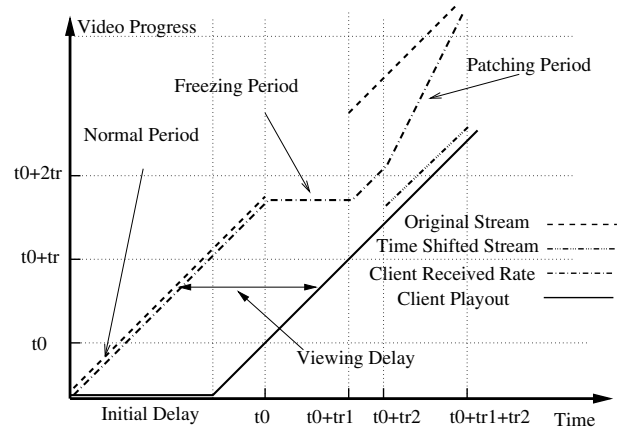


Fig. 3. Continuous video streaming with initial delay

$D$  is larger than the *node rejoin delay*, the client buffer will not be drained out before it is reconnected to the server. Once the client is reconnected, it receives the time-shifted (patching) stream as well as the original stream, until the missed video is made up. During this process, the client does not experience any pause of playout, neither does it suffer any extra delay. If  $D$  is smaller than the *node rejoin delay*, then the client side video playback has to be paused.

### D. An Example of System Operations

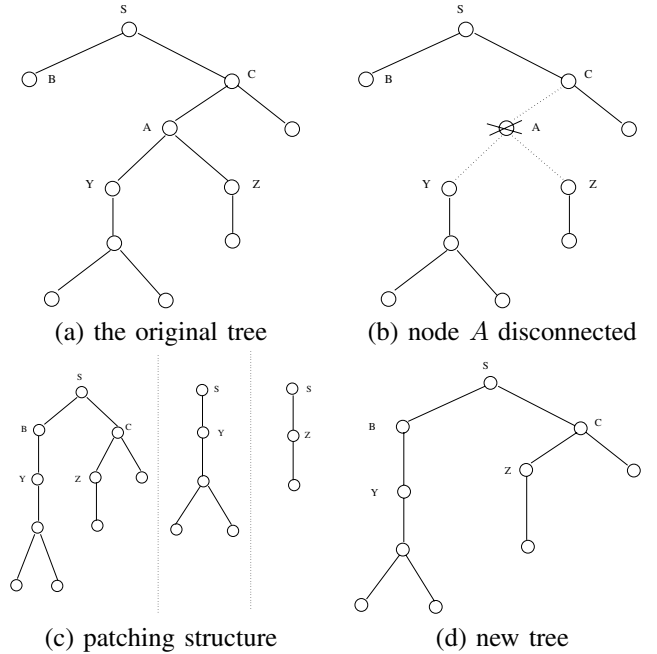


Fig. 4. The example of system operations

Figure 4 shows several snapshots of the application layer multicast tree during the video streaming process. Figure 4(a) shows the multicast tree structure where all the clients in the tree receive the original stream. At time  $t_0$ , node  $A$  leaves the group either due to congestion or node failure. This causes

the subtrees rooted at nodes  $Y$ , and  $Z$  to suffer service disconnection. This is shown in Figure 4(b).

Nodes  $Y$  and  $Z$  send “rejoin” message to server  $S$  respectively. We describe the rejoin process for node  $Y$ . Node  $Z$  has similar rejoin process. The “rejoin” message includes the rejoin requests for both the original stream and the time-shifted stream. To reconnect the client to the original stream, server  $S$  selects a client which is currently in the multicast tree as the parent node of  $Y$ . In this case, client  $B$  is selected. Node  $Y$  connects with client  $B$  at time  $t_{Y1}$ , and begins to receive the original stream from it. At the same time, the server allocates a free video channel to send out the patching stream to  $Y$ . Assume the server starts to send out the patching stream at  $t_{Y2}$ , and stream  $s_i$  is the designated patching stream, where  $i = \lceil \frac{t_{Y2}-t_0}{d} \rceil$ . Figure 4(c) shows the multicast tree structure at time  $t_1$ , when both nodes  $Y$  and  $Z$  are in patching period. Figure 4(d) shows the tree structure when both patching channels for nodes  $Y$  and  $Z$  are released.

We analyze the influence of this service disconnection to node  $Y$ . The children of  $Y$  suffer the same experience. Node  $Y$  begins to receive the original stream at time  $t_{Y1}$ , it misses video  $[t_0, t_{Y1}]$  due to service disconnection. The patching stream starts at  $t_{Y2}$ . Assume at time  $t_0$ , the client buffer has  $b_0$  time unit of video content.  $Y$  playout the video content from its buffer, when the client is disconnected. If  $b_0 \geq t_{Y2} - t_0$ , then the client still plays video from the buffer when the patching stream begins to feed the client. Thus, the client’s viewing process is not interrupted, and the client viewing delay remain unchanged. Otherwise, the client may suffer a pause of  $t_{Y2} - t_0 - b_0$ . At time  $t_0 + t_{Y2} + t_{Y1}$ , the patching period finishes and the patching stream is released. At this time, the client buffer size is  $b_0 + \max(0, t_{Y2} - t_0 - b_0)$ .

### III. DESIGN DETAILS

In this section, we describe our system in detail, including the time shifting video server design, and the tree management schemes.

#### A. Time Shifting Video Server

A media server is the single source for the live video content, and is the root of the application layer multicast tree. In this section, we describe the time shifting video server design in our system.

Figure 5 shows the structure of our server design. The *Client Request Aggregator* receives four kinds of client requests: *join*, *leave*, *rejoin*, and *patching end*. When the server receives a *join* request, it contacts the *Peer Selector* to find the appropriate parent node<sup>2</sup> to connect with. The *Peer Selector* obtains global topology and resource availability information from the centralized database. In the case of a graceful *leave*, the server signals the children of the leaving client to rejoin other parent nodes. The server also updates the *tree structure and resource availability* information upon a node leave. If a client experiences severe packet loss due to network congestion

<sup>2</sup>Parent node can be either the video server or a client which is currently in the application layer multicast tree.

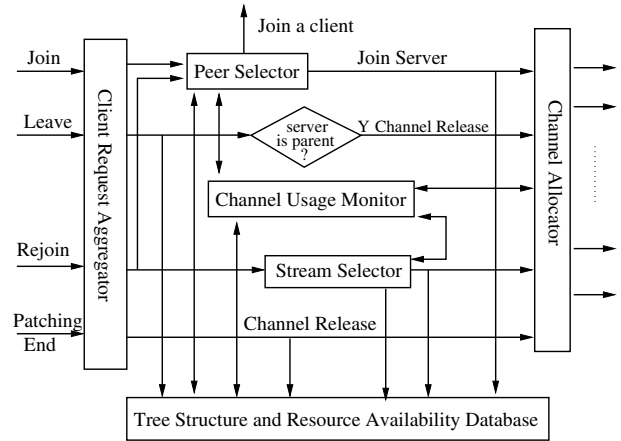


Fig. 5. Video server Design

or ancestor node failure, it sends a *rejoin* message to the server. The server then assigns another parent node to forward the original stream to it. At the same time, the server also contacts the *Channel Usage Monitor*, and assigns a free video channel to transmit the time-shifted video stream to this client. The *Stream Selector* is responsible for assigning the client a stream with appropriate time shifting value. When the missed video portion has been fully received, the client sends a *patching end* message to the server, the server then releases the corresponding patching channel.

#### B. Channel Allocation:

Our video server streams video content in two kinds of channels. The channel for the original stream is called the *live channel*, and the channel for the time-shifted stream is called the *patching channel*. The data rates for the live and the patching channels are the same. Video content that is transmitted on a particular channel is multicast to the clients in the application layer multicast tree associated with this channel. Channel allocation deals with whether the channel is used to transmit the original video stream or the time-shifted video stream. Allocating more channels to the original stream at the server leads to a shorter and wider tree. Reserving more channels for the time-shifted streams means that when the client’s rejoin message reaches the server, there is a higher probability that the server has a free patching channel available. In this way, the client can start to receive the patching stream earlier. In this section, we propose three channel allocation schemes, as described below.

1) *1:1 Allocation*: One way to allocate video channels is that whenever a server allocates one live channel to the client, it also reserves one patching channel for this client. We call this allocation scheme, *1:1 allocation*. This approach has the benefit that whenever there is a node failure in the application layer multicast tree of a live channel, there is a free patching channel available for the clients below the failed node to patch the missed video content.

Figure 6 shows an example of the *1:1* allocation. As shown




Fig. 6. 1:1 channel allocation

in Figure 6(a), node  $A$  in the multicast tree of channel  $L_i$  leaves the group. The subtree of node  $A$  manages to rejoin the patching channel  $P_i$  for the time-shifted video stream, as well as rejoin the multicast tree for the original stream. The problem with this channel allocation scheme is that each patching channel is bound with a live channel. If at time  $t$ , there are more than one patching channel requests from live channel  $L_i$ , only one of them can be served, even though the patching channels for other live channels are left unused.

2) *Static Multiplexing*: To overcome the inefficiency of 1:1 allocation, we propose the *static multiplexing* allocation scheme. In this scheme,  $m$  of the  $C$  video channels are allocated as live channels, while the other  $n = C - m$  channels are allocated as patching channels. There is no fixed binding in this scheme, the patching channel can be used to patch the disconnected clients from any live channel.




Fig. 7. Static multiplex channel allocation

Figure 7 shows how this channel allocation scheme works. Node  $A$  of channel  $L_i$  leaves the group, causing service disconnection for all its children. These nodes receives the patching stream by rejoining one of the free patching channel  $P_j$ . Obviously, this scheme is more efficient than the 1:1 scheme, since the video patching request will be served as long as there is a free patching channel available.

3) *Dynamic Multiplexing with Channel Merging*: In the *static multiplexing* scheme, the value of  $m$  and  $n$  is pre-determined and fixed throughout the video transmission process. A more flexible scheme is to assign the number of live and patching channels dynamically. In this section, we propose a *lifetime-based allocation* scheme. The lifetime of an original stream is the remaining playback time of the video program. The lifetime of the patching stream is the duration of the

patching period. For a video program of length  $T$ , at time  $t_0$ , the lifetime of the original stream is  $T - t_0$ . As to the patching stream, the life time is the time shifting value of this stream; the lifetime for stream  $s_k$  is  $d \times k$ . Our allocation algorithm works as follows: For the patching stream, at time  $t_0$ , the number of patching channels reserved is:

$$|P| = \lceil \frac{\sum_{i=1}^m (i \times d)}{T - t_0 + \sum_{i=1}^m (i \times d)} \times C \rceil$$

For the original stream, at time  $t_0$ , the number of video channels allocated is:

$$|L| = C - |P|$$

The number of patching channels  $|P|$  is monotonically increasing as the video program proceeds. It is possible that when the server needs to allocate more patching channels, there is no free channel available. We introduce a channel merging scheme to deal with this problem. In this solution, when the server's request for more patching channels can not be satisfied, two live channels that have the fewest number of clients in their multicast trees are merged, and the freed channel is used as a patching channel. To merge two live channels, just connect the root node of the channel with fewest members to the highest possible level of the other channel.

### C. Patching Stream Selection

The *Stream selector* in the video server determines which time-shifted stream should be used to patch the missed video portion for a certain client. When a client is disconnected from the group, it records the video playback time  $t_d$ , and sends a node rejoin request to the server. Assume when the server receives this node rejoin request, the live video playback time is  $t_1$ .

If the server has an available channel at this time, it estimates the connection set up time as  $t_1 - t_d$ . This is the time for the rejoin signal to travel from the client to the server. Thus, stream  $s_i$  is selected, while  $i$  is calculated by the formula below:

$$i = \lceil \frac{2 \times (t_1 - t_d)}{d} \rceil$$

If there is no available channel when the rejoin message reaches the server, this rejoin message will be held until a free patching channel is available. For example, at time  $t_2$ . Then the stream  $s_j$  should be selected, while  $j$  is determined by the formula below:

$$j = \lceil \frac{(t_2 - t_d) + (t_1 - t_d)}{d} \rceil$$

### D. Node join algorithm

A "well shaped" application layer multicast tree should be wide and short. A shorter multicast tree means a smaller number of overlay hops from a client to the server, thus a smaller average stretch. Stretch is the ratio of the latency along the overlay to the latency along the direct unicast path [2], [5]. The stretch of the application layer multicast tree is the average stretch value over all the clients in the tree. Furthermore, by

reducing the number of intermediate clients, the chance that a client suffers service disconnection due to ancestor nodes leave is also reduced. In this section, we design a level-based scheme to handle client join. In this scheme, a newly arriving client joins the node whose overlay path distance to the server is shortest, and has enough available bandwidth resource to accommodate a new client. The join algorithm is shown in Figure 8:

```

Join (C, S) {
1:   joined = false;
2:   Push the server IP address  $S$  in current list  $L_c$ ;
3:   while (joined==false and  $L_c \neq \emptyset$ ) {
4:     repeat
5:       Probe the next IP address  $p_n$  in  $L_c$ ;
6:       if ( $p_n$  has enough resource for new client) {
7:          $C$  join the tree and become  $p_n$ 's child;
8:         joined ==true;
9:       }
10:      Push the children of  $p_n$  into  $L_n$ ;
11:     until( $L_c$  is visited or joined == true)
12:      $L_c = L_n$ ;
13:   }
14:   return joined;
15: }

```

Fig. 8. Node Join Algorithm

The patching scheme in our design requires extra bandwidth for each client, we now analyze how this scheme influences the overall structure of the tree. Assume the bandwidth of a video stream is  $b$ , and the average bandwidth of each host is  $k \times b$ . The height of a tree with  $N$  hosts should be  $h_1 \approx \log_k N$ . Because of the use of the video patching scheme, each host has to allocate twice the bandwidth of a video stream: one for the original stream, the other is for the patching stream. Thus, the height of the tree is  $h_2 \approx \log_{\frac{k}{2}} N$ . We now have:

$$\frac{h_2}{h_1} \approx \frac{\log_{\frac{k}{2}} N}{\log_k N} = \frac{\log k}{\log k - 1}$$

To further reduce the height of the tree, we design a *high-bandwidth-first* tree join algorithm. The key idea is to push the client with larger bandwidth up to the higher level of the tree. The multicast tree built by this scheme has the feature that clients in the lower level of the tree do not have more bandwidth capacity than the clients in the upper level of the tree. Figure 9 and Figure 10 shows the *high-bandwidth-first* join algorithm.

#### E. Node Leave

Clients in the application layer multicast tree may leave the group at any time. A leaf node leaving the group does not have much impact on the application layer multicast tree. An internal node leaving the group, on the other hand, will cause service disconnection for its children. There are two

```

BWJoin( $c, S, bw_c$ ) {
01:  joined = false;
02:  Push the server IP address  $S$  in current list  $L_c$ ;
03:  while ( $L_c \neq \emptyset$ ) {
04:    Push the children of members in  $L_c$  into  $L_n$ .
05:    repeat
06:      probe the bandwidth of next host  $p_n$  in  $L_n$ ;
07:      if ( $p_n$ 's bandwidth <  $bw_c$ )
08:        if (Join_Level( $L_c$ ) == true) return(true);
09:      else
10:         $c$  becomes  $p_m$ 's parent's child;
11:         $p_m$  becomes the child of  $c$ ;
12:        return(true);
13:      if ( $p_n$ 's bandwidth ==  $bw_c$ )
14:        if (Join_Level( $L_c$ )==true) return(true);
15:        else
16:          if (Join_Level( $L_n$ )== true) return(true);
17:        until( $L_c$  is visited)
18:        if ( $L_n = \emptyset$ )
19:          if (Join_Level( $L_c$ )==true) return(true);
20:         $L_c = L_n$ ;
21:      }
22:    }
23:  return joined;
24: }

```

Fig. 9. Bandwidth First Node Join Algorithm

```

Join_Level( $L_c$ ) {
1:  repeat
2:    Probe the next host  $p_a$  in  $L_c$ ;
3:    if( $p_a$  has enough bandwidth)
4:       $c$  join the tree and become  $p_a$ 's child;
5:    return true;
6:  until ( $L_c$  is visited)
7:  return(false);
8: }

```

Fig. 10. Level Join Algorithm

kinds of leave: graceful leave and node failure. With a graceful leave, the client first informs the server and its children. Its children then manage to rejoin the tree by connecting to other parent nodes. It leaves the tree after the adaptation of the tree is finished. In the case of a node failure, the client leaves the group without informing any other hosts. The tree recovery operation due to a node failure is a two step process: First, the failed node and affected region<sup>3</sup> detection. Second, disconnected nodes rejoin the tree, this step follows the same operation as the rejoin process under a graceful node leave.

There are two approaches to discover the failed node and the affected region, as shown in Figure 11. First approach

<sup>3</sup>Affected region of a client  $c$  denotes the set of all nodes that are disconnected due to node failure of  $c$ .

is through localized detection. A client that detects a heavy packet loss sends a *hello* message to its parent node. If the parent node is experiencing the same problem, it sends an *echo* message back to the sender, and sends another *hello* message to its parent. If the parent node does not suffer packet loss, then it is the link congestion between the child and the parent causing the problem. This process repeats until a client does not receive *echo* message from its parent, then either this parent node is failed or the link between this client and its parent is disconnected. This approach does not require global topology knowledge, and detection does not exert extra load on the media server. The problem of this approach is that it might be slow in detecting the affected region. The other approach is to use the central video server. The video server maintains the topology of the application layer multicast tree. Each client that suffers service disconnection reports to the server, the server then figures out which node or link is failed, and the corresponding affected region.

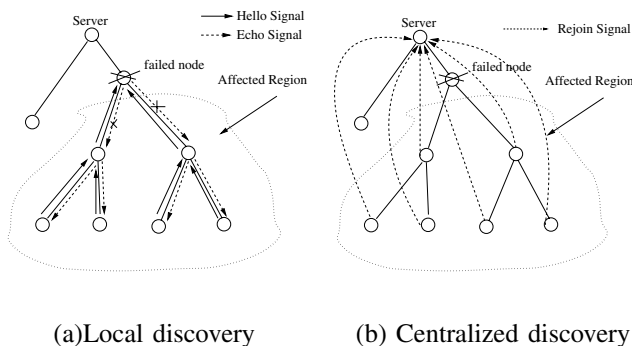


Fig. 11. Failed node discovery process

The node rejoin process works like this: the affected clients first try to elect a central node. A central node is a child of the failed node with enough bandwidth resource to accommodate all its siblings. If there exists such a central node, then the central node rejoins the parent of the leaving node, and all the children of the leaving node rejoin this central node as its children. If no central node can be elected, each child of the leaving node as well as the sub-tree rooted at them rejoins the application layer multicast tree independently.

#### IV. PERFORMANCE EVALUATION

##### A. Simulation Environment Setup

We use the GT-ITM transit-stub model to generate a network topology [4] of about 1400 routers. The average degree of the graph is 2.5, and core routers have a much higher degree than edge routers. The media server and the end hosts are connected with the edge routers. We categorize the bandwidth capacity of the end hosts into three levels. 1) Narrow bandwidth: with  $1.5Mbps$  bandwidth, 70% of the end hosts are in this category. This bandwidth capacity is only enough for receiving video streams. The hosts with such bandwidth capacity can only be leaf nodes in the tree. 2) Medium bandwidth: with  $10Mbps$  bandwidth, 20% of

end hosts belong to this category. 3) High bandwidth: with  $100Mbps$  bandwidth, only 10% of hosts have such high-speed connection.

There is one fixed media server in the topology, it has  $100Mbps$  high speed connection. The time shifting value between stream  $s_i$  and  $s_{i+1}$  is 4 seconds. The server processing capability and I/O capacity is not a bottleneck compared to the server side bandwidth. We further assume links between the routers are never a bottleneck.

##### B. Evaluation of client video reception performance

In this section, we study the client video reception performance in our system. We are interested in the following performance metrics: *client viewing delay*, *maximum buffer usage*, *video continuity*, and *video completeness*. *Client viewing delay* means the delay between the client side playback time and the playback time of the original stream. *Maximum buffer usage* records the maximum buffer usage throughout a client's video reception process. *Video continuity* for a client is evaluated by the duration and frequency of freezing period. *Video completeness* refers to the rate between the received video and the transmitted video.

1) *Effect of video patching*: We first compare the video reception performance of a client with or without video patching scheme. In this experiment, we record a client's viewing behavior from the time it first joins the application layer multicast tree until it leaves the group. We assume infinite client buffer size in this simulation.

Figure 12 shows the simulation result. Figure 12(a) shows the client viewing delay. Without video patching, the client suffers the longest viewing delay. Video patching scheme can greatly reduce the viewing delay. As to the buffer usage, video patching scheme requires some buffer space to store the video packets that is not being played out. The no patching scheme does not use any buffer space, as shown in Figure 12(b). Figure 12(c) shows the video continuity during the viewing process. Without video patching scheme, the client suffers freezing period every time it is disconnected. With video patching, the freezing period is significantly reduced, since the client can still playout the previously buffered video content when the subsequent service disconnection happens. We also find that, if the client starts playout the video after an appropriate initial delay, the client can receive continuous video stream.

2) *Adaptation to different traffic patterns*: We now study the video reception performance of a client under different traffic patterns. In this experiment, the clients join the multicast tree according to a *Poisson* process, their average time in the group varies from 100s to 300s. We manually insert a client into the group, and it does not leave the group throughout the simulation process. The simulation time is one hour, we sample the viewing delay for this client every 100 seconds, and calculate the average buffer usage during this 100 second interval.

Figure 13 shows the simulation result. The shorter the average client lifetime, the longer the client viewing delay.

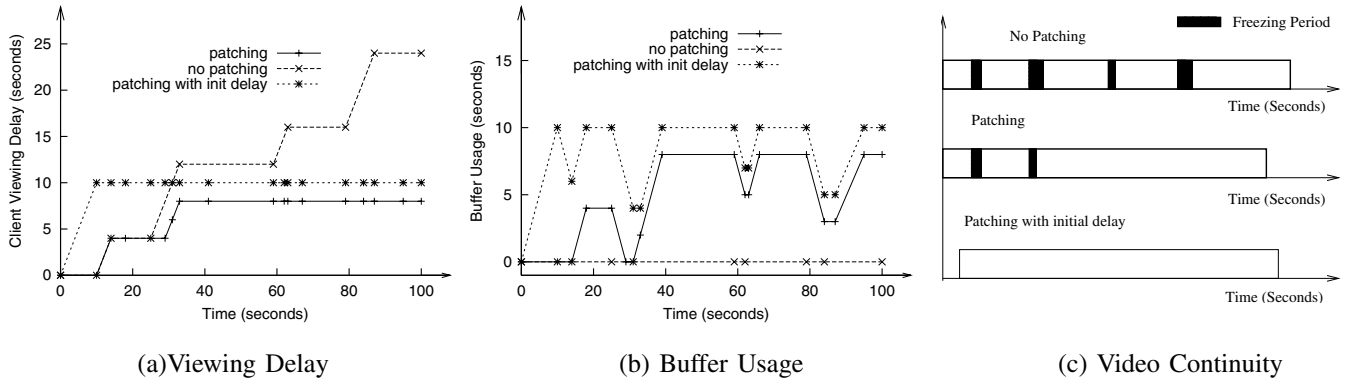


Fig. 12. Video patching on client video reception

Since as the clients leave the group more frequently, it is more possible that the client is disconnected again before the patching period finishes. In this way, the client has to join later and later time-shifted streams each time it rejoins the multicast tree.

The client buffer usage is also increased as the clients join/leave the group in a more frequent manner. Since under this case, the clients have to join the patching stream with larger time shifting value. This causes longer patching period, and accumulates more data in the client buffer.

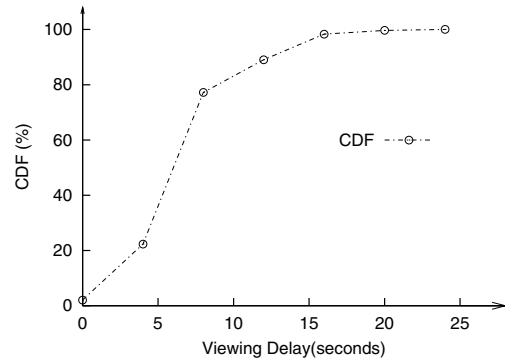


Fig. 14. Viewing delay distribution

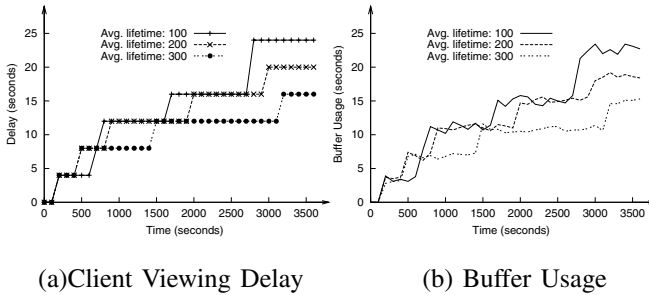


Fig. 13. Video Patching on viewing performance

3) *Client viewing delay distribution*: In this section, we focus on the viewing delay for all the clients in the system. In the simulation experiment, we record the maximum viewing delay of each client, and study their overall distribution. The average lifetime of a client is 100 seconds.

Figure 14 shows the distribution of client viewing delay. The client buffer usage demonstrates very similar distribution. The maximum client viewing delay value is 24 seconds. Most of the clients' viewing delay is between 4 and 16 seconds. The factors that affect the viewing delay are: the client rejoin latency, and the server video channel usage condition. The client rejoin latency is mostly determined by its location in the application layer multicast tree. There are some clients directly connected with the video server, these clients suffer no viewing delay at all. Those clients further away from other clients tend to need longer time to rejoin the application layer multicast tree. Another factor that influences the viewing delay

is the server channel usage condition. If the server does not have a free patching channel available when the client rejoins the tree, the patching stream has to be delayed, as well as the client viewing delay. Clients that suffer longer viewing delay also need more buffer space, since the patching period tends to be longer.

### C. The effect of video patching on tree structure

Video patching requires extra bandwidth on the client side, which reduces the forwarding capability, in terms of fan out, of the clients. Thus, the "width" of the tree will be reduced.

In this section, we study the influence of the patching scheme on the structure of the application layer multicast tree, and the average stretch of the end hosts. In this experiment, the clients join the multicast group in a *Poisson* process, and stay in the group throughout the experiment. Thus, the number of clients in the tree is monotonically increasing. The maximum number of clients in the tree is about 100000 in our simulation. The simulation time is 1 day in this experiment. We compare four tree join algorithms in this experiment: *level-based* algorithm, with and without video patching; *high-bandwidth-first* algorithm, with and without video patching.

Figure 15(a) shows the height of the tree. The level-based tree join algorithm with video patching generates the highest multicast tree. The high-bandwidth-first algorithm promotes the clients with high bandwidth to the upper level of the

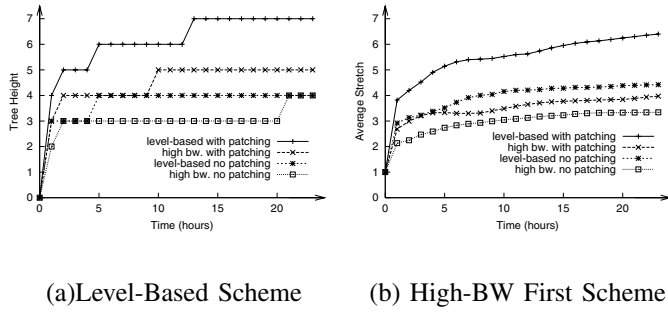


Fig. 15. Influence of video patching on tree structure

multicast tree, thus increases the fan out of the tree in the higher level. The height of the tree for the high-bandwidth-first algorithm is significantly smaller than the level-based scheme. Furthermore, for the high-bandwidth-first algorithm, the height of the tree with video patching does not increase much compared to the scheme without video patching.

Figure 15(b) shows that the level-based tree join algorithm with video patching has the worst stretch performance, and is much worse than the same tree join algorithm without video patching. For the high-bandwidth-first scheme, the shape of the tree is optimized, and the height of the tree is comparatively shorter, the stretch performance with or without video patching is close to each other.

#### D. Server Channel Allocation Scheme

The server side bandwidth resource (video channel) availability determines the shape of the multicast tree, and the starting time of the patching stream. How to effectively assign server channels is important to the overall system performance, and to the satisfaction of client video reception. We have proposed three video channel allocation scheme: 1:1, static multiplexing, and dynamic multiplexing. In this section, we evaluate two aspects of video channel allocation scheme: 1) video channel utilization 2) client queuing delay for the patching channel.

1) *Server channel utilization*: Server channel utilization means the percentage of the number of channels in use to the number of all the video channels. In this experiment, we assume the server can support 100 video channels simultaneously. We run the simulation for 1 hour, and record the average channel utilization value every 100 seconds. Figure 16 shows the simulation results.

For 1:1 channel allocation scheme, the channel utilization value is worst. In this scheme, half of the video channels are allocated as live channel, and each live channel is bound with a patching channel. A patching channel is used only after a service disconnection in the application layer tree of the associated live channel, otherwise, it is free. And the life time of a patching channel is short compared to the live video program time. Furthermore, in 1:1 scheme, even if there are multiple rejoin requests in the application layer multicast tree of a live channel, the associated patching channel can only

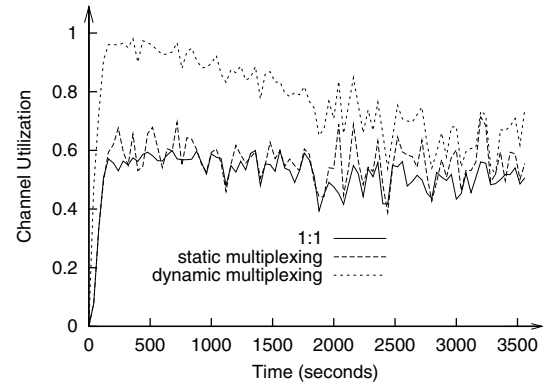


Fig. 16. Video Server Channel Utilization

serve one request at a time, other requests have to be put into a queue until this patching channel is free.

In *static multiplexing* scheme, we use half of all the video channels as live channel, and reserve the other half as patching channel. The channel utilization value is better than the 1:1 scheme. Since multiple rejoin requests from the same live channel can be satisfied as long as there are free patching channels. And the channel utilization has larger variations, since when there are multiple rejoin requests, they are served simultaneously, instead of one at a time in 1:1 scheme.

*Dynamic multiplexing with channel merging* performs best in terms of video channel utilization. In the beginning stage, most of the video channels are used as live channel. The channel utilization value is close to 100%. As time proceeds, the server needs to reserve more channels for video patching. It does this by merging the live channels with fewest clients, and reserve these channels as patching channel. Thus, there is a degradation in channel utilization.

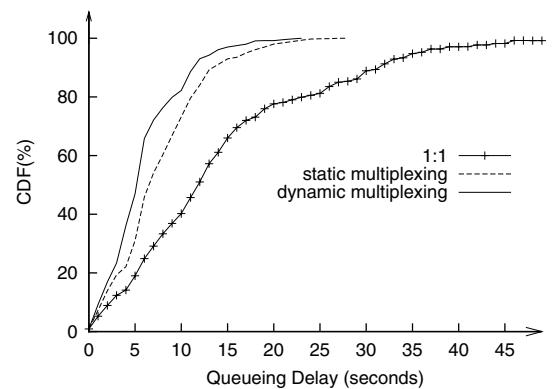


Fig. 17. Client Requests Queuing Delay

2) *Client queuing delay*: Client queuing delay refers to the time when the server begins to receive the video patching request until the time this requests is being served. We compare the client queuing delay under the three video channel allocation schemes. We use the same configuration as last section, Figure 17 shows the simulation result.

We find out that the queueing delay for  $1:1$  scheme is significantly longer than the other two schemes. Since the patching channels are not shared in this scheme. If a node with many children leaves the group, there could be many rejoin requests at the same time. Under  $1:1$  scheme, only one request can be served at a time, while the other requests have to wait until this patching channel is free. For the other two schemes, they use channel multiplexing, so that the rejoin requests can be served as long as there are patching channels available. The queueing delay in these schemes is significantly reduced. For the dynamic multiplexing case, when the patching channel is not enough to handle the rejoin requests, some live channels are merged and used as patching channel. Thus, dynamic multiplexing can further reduce the queueing delay.

### E. Protocol Overhead Analysis

In this section, we consider several aspects of complexity in our system design. We only give qualitative analysis in this paper, for quantitative analysis, please refer to our technical report [12].

- Message processing overhead: in our system design, the server has to process four kinds of messages: *join*, *leave*, *rejoin* and *patching end*. We assume the weight for processing these messages is the same. Our node rejoin message is composed of *join live* and *join patching* messages. The *patching end* message is coupled with a *join patching* message. Assuming there are  $N$  node join message, and  $M$  node rejoin messages throughout the video streaming process, the number of messages should be:  $2 \times N + 3 \times M$ . For those no-shifting, no patching solutions, the number of messages is:  $2 \times N + M$ . For a CoopNet solution with  $m$  descriptions, each *join*, *leave*, *rejoin* message involves operations over  $m$  trees, the number of messages should be:  $(2 \times N + M) \times m$ .

- Tree management overhead: the video server maintains one application layer multicast tree for the original stream, and multiple application layer multicast trees for the patching stream. Note that the patching tree has considerably smaller size and shorter lifetime than the original tree.

- Bandwidth overhead: although the MDC codec used in CoopNet introduces some bandwidth overhead [1], [7], our solution is more bandwidth intensive. Since we need to reserve same amount of bandwidth for each original stream allocated. But the adoption of the MDC approach brings extra overhead such as coding/decoding complexity, synchronization of multiple streams, etc.

## V. CONCLUDING REMARKS

In this paper, we deal with the problem of continuous live video streaming to a group of cooperative clients. Our solution is centered around a time-shifting video server, and a video patching scheme. The time-shifting video server sends multiple video streams with different time shifting values. Clients in the application layer multicast tree can retrieve the missed video content by joining the time-shifted video stream. To avoid *indefinite time-shifting* due to multiple service disconnection during the video reception process, we introduce

the video patching scheme. During the patching period, a client is receiving the time-shifted video stream as well as the original stream. The video content that is not being played out immediately is stored in a client buffer. When a subsequent service disconnection occurs, a client can play the video content from the buffer while rejoining the group. In this way, the client can receive the complete video program even though the forwarding infrastructure is unreliable. Continuous video streaming is achieved if the client starts video playout after some *initial delay*.

Our design has the following features: 1) lossless video reception: by allowing clients rejoin the time-shifted video stream, the client can receive the whole video content from the point it first joined the group. 2) stable video quality: the client receives full quality video throughout the video reception process. 3) continuous video streaming: continuous video streaming can be achieved by sacrificing initial video access delay. 4) Compared to CoopNet's MDC-based system, our system has the advantage that it can use standard-based single description encoded streams. 5) Moderate complexity: the overhead of message processing and tree management is at the same level with a no-shifting, no-patching solution.

## REFERENCES

- [1] J. Apostolopoulos. Reliable Video Communication over Lossy Packet Networks using Multiple State Encoding and Path Diversity . In *Proceedings of VCIP 2001*.
- [2] S. Banerjee, B. Bhattacharjee, C. Kommareddy Scalable Application Layer Multicast In *Proceedings of ACM SIGCOMM*, August 2001.
- [3] S. Banerjee, C. Kommareddy, K. Kar, B. Bhattacharjee, S.Khuller Construction of an Efficient Overlay Multicast Infrastructure for Real-time Applications To appear in *Proceedings of Infocom*, 2003
- [4] K. Calvert, M. Doar, and E. Zegura. Modeling Internet topology. *IEEE Communications Magazine*, June 1997.
- [5] Y.H. Chu, S.G. Rao and H. Zhang A Case For End System Multicast In *Proceedings of ACM SIGMETRICS*, June 2000.
- [6] Y.H. Chu, S.G. Rao and H. Zhang Enabling Conferencing Applications on the Internet using an Overlay Multicast Architecture In *Proceedings of ACM SIGCOMM*, August 2001.
- [7] V.K. Goyal, J. Kocevici, R. Armean, and M. Vetterli. Multiple Description Transform Coding of Images In *Proceedings of ICIP*, 1998
- [8] D. Hrishikesh; B. Mayank; G. Hector Streaming Live Media over a Peer-to-Peer Network. Technical Report, Stanford, 2001
- [9] K.A. Hua, Y. Cai, S. Sheu. Patching: a multicast technique for true video-on-demand services . In *Proceedings of the sixth ACM international conference on Multimedia*, 1998.
- [10] Duc A. Tran, Kien Hua, Tai Do ZIGZAG: An Efficient Peer-to-Peer Scheme for Media Streaming In *Proceedings of InfoCom 2003*.
- [11] M. S. Kim, S. S. Lam, D.Y. Lee Optimal Distribution Tree for Internet Streaming Media. *Technical Report* , U.T. Austin, 2002
- [12] M. Guo, M.H. Ammar Scalable live video streaming to cooperative clients using time shifting and video patching. *Technical Report GIT-CC-03-40*, College of Computing, Georgia Tech, 2003.
- [13] V. N. Padmanabhan, H. J. Wang, P. A. Chou, K. Sripanidkulchai. Distributing Streaming Media Content Using Cooperative Networking. In *Proceedings of NOSSDAV 2002*.
- [14] V. N. Padmanabhan, H. J. Wang, P. A. Chou. Resilient Peer-to-Peer Streaming In *Proceedings of ICNP*, November, 2003.
- [15] D. A. Tran, K. A. Hua, T. T. Do Peer-to-Peer Streaming Using A Novel Hierarchical Clustering Approach To appear at *IEEE JSAC Special Issue on Advances in Service Overlay Networks*, 2003
- [16] D.Y. Xu, M. Hefeeda, S. Hambrusch, B. Bhargava On Peer-to-Peer Media Streaming. In *Proceedings of ICDCS*, 2002
- [17] L. Zou, E.W. Zegura, M.H. Ammar The Effect of Peer Selection and Buffering Strategies on the Performance of Peer-to-Peer File Sharing Systems. In *Proceedings of MASCOTS 2002*, October 2002.