A Performance Comparison of Multiple Description Video Streaming in Peer-to-Peer and Content Delivery Networks

Shoaib Khan, Rüdiger Schollmeier, Eckehard Steinbach
Institute of Communication Networks
Technische Universität München
{khan, hrs, steinb}@lkn.ei.tum.de

Abstract

In this paper we examine the performance of Peer-to-Peer (P2P) media streaming using Multiple Description Coding (MDC) and compare it with that of a Content Delivery Network (CDN). In both approaches multiple servers simultaneously serve one requesting client with complementary descriptions. This approach improves reliability and decreases the data rate a server has to provide. We have implemented both approaches in the ns-2 network simulator. The experimental results indicate that the user perceived video quality of MDC-based streaming in a P2P network can be significantly better than in a CDN, despite the high degree of unreliability of the P2P network.

1. Introduction

Video streaming over the Internet has become increasingly popular in recent years due to the rapid rise in network access speed of the end-users. Media streaming systems are distinct from file-sharing systems, in which a client has to download the entire file before using it. In a media streaming session, the receiver can already consume the file while downloading. Video streaming applications are highly susceptible to packet delay and packet loss. A packet arriving after its scheduled playback time is useless and considered lost. A lost or corrupted frame creates an avalanche effect in the decoding process, as the decoding of subsequent frames is impaired by spatio-temporal error propagation. Therefore, in the case of loss, decoding is often completely stopped until the next I-frame arrives. In Internet video streaming, variations in transmission quality (throughput, delay etc.) are smoothed using a receiver buffer. The size of the buffer corresponds to the user-perceived initial delay of the application.

Today's video streaming systems are mostly based on the client server model of Content Delivery Networks (CDN) which leads to a number of problems. The most important ones are:

- Flash Crowds: Even the large streaming servers are not able to feed more than a few hundred streaming sessions simultaneously [12]. P2P based media streaming solutions can alleviate the effects of flash crowds more effectively [1].
- Single Point of Failure: Like any client-server model, the server is the single point of failure.
- Bandwidth cost: It can be a significant problem to the content provider. In contrast, these costs are shared by every participant in the P2P streaming network.
- Selection of the optimal server is difficult in CDNs. In P2P networks, a node located in the proximity of the requesting peer might serve the content, as the content is available at the edge of the network.

P2P networks offer characteristics and possibilities which can not be provided by CDNs. As we show in this work, the performance of media streaming can be better in a P2P network, although the probability that one stream breaks is higher [7], [8]. The reason for this is that the replication rate of the video streams in a P2P network is typically significantly higher than in a CDN, due to the large number of participating hosts. In Gnutella, for example, every peer shares an average of 500 files [13] and many peers host the same file.

Multiple Description Coding is a source coding technique which encodes a signal into a number of separate bitstreams. Each individual bitstream is called a description. The descriptions are sent through different network paths to a destination. The receiver can make a useful reproduction of the signal when at least one of the descriptions is received. The quality of the reproduced source is proportional to the number of descriptions received. This property makes MDC highly suitable for lossy packet networks where there is no prioritization among the packets. For a general overview on Multiple Description Coding refer to [2]. It has been demonstrated in [6], [10], that using MDC in combination with packet path diversity significantly improves the robustness of a real-time video application. We can therefore use a small receiver buffer which translates to a low initial startup delay of the streaming session. In a similar spirit, in this work, an MDC encoded video is streamed simultaneously from multiple geographically distributed peers. Using MDC in a P2P streaming scenario is illustrated in Fig. 1. Peer p1 wants to receive video file D which is available in the MDC format on p3, p5 and p6. In this example the video is encoded using two descriptions D1 and D2. Peers p3 and p6 are chosen based on, e.g., the distance in hops. Thus they can simultaneously serve the video file D, each one providing a complementary description. If any of the descriptions is affected by packet loss or excessive delay, the receiver can still decode and display video D but at the expense of a degradation of the quality, as the descriptions are independently decodable. As this example shows, MDC-based media streaming is well suited for situations where the quality and availability of connections vary significantly over time. In this work we examine the performance of MDC-based video streaming in P2P networks and additionally compare it to MDC-based video streaming in traditional CDNs.

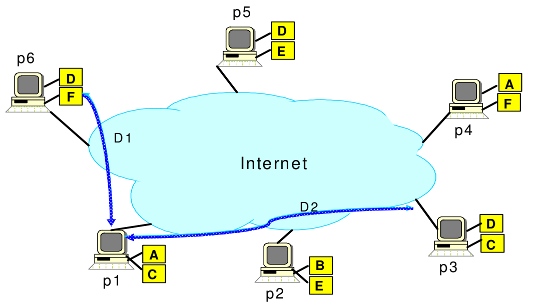

Figure 1. MDC-based video streaming over a P2P network.

Peer-to-peer based media streaming approaches using multiple serving hosts have been proposed in [3] and [4]. CoopNet [14] proposes to build multiple description trees spanning the source and all the receivers, each tree transmitting a separate description of the media signal. CoopNet puts a heavy control overhead on the source since the source must maintain full knowledge of all distribution trees. In [6] MDC-based distributed video streaming has been proposed for content delivery networks. Our approach is inspired by this work and we use the same MDC technique for our P2P based media streaming.

The rest of the paper is organized as follows. Section 2 presents our simulation model for the performance measurements. The performance study of our system is described in Section 3, along with simulation results.

2. Simulation Model

In P2P networks, peer and content availability pose a challenging problem. Availability of a peer in a P2P network is quite unpredictable, depending primarily on human presence. In our experiments we model peer availability as a 2 state process, having the states ON and OFF, where the OFF state is non-returning (see Fig. 2). The median session duration of a peer in a Gnutella network is found to be about 60 minutes [8]. In our work a serving peer is assumed to be in the ON state when a streaming request arrives. Once the peer availability model transitions to the OFF state this peer becomes unavailable for the current session. In our experiments we select the transition probabilities such that the average serving time of a peer equals 30 minutes.

In a CDN, the server placement problem addresses how to optimally place a number of servers in order to maximize the quality at the end user. For a particular number of servers, we

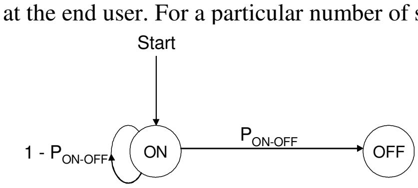

Figure 2. Peer availability model

place the servers randomly in the network and measure the average round-trip-time from each user to the servers. We perform this random placement 10 times and choose in our experiments the one yielding the smallest average round-trip-time.

The server selection problem addresses how to optimally choose a pair of servers to get complementary descriptions in order to maximize the perceived quality at the receiver. In [6] Apostolopoulos et al. propose a path diversity model which requires knowledge about the network topology of the CDN including knowledge of joint and disjoint links, and loss characteristics for each link. In our experiments for the CDN case we simply choose the closest two servers, in terms of roundtrip time, for each client request. For the P2P case, we choose the closest two serving peers offering the requested content.

Another performance parameter is the optimal distribution of the MDC streams among a set of servers. In this paper we assume that all the CDN servers hold both descriptions. This simplifies the server selection problem by merely choosing the two closest servers, as mentioned above.

In the P2P case, we model the distribution of content among the peers by randomly selecting peers to share a particular media file. We vary the percentage of peers having the file from 0.5% to 100%. To simulate various degrees of network load, we create random TCP connections originating from randomly picked nodes with randomly picked destination nodes.

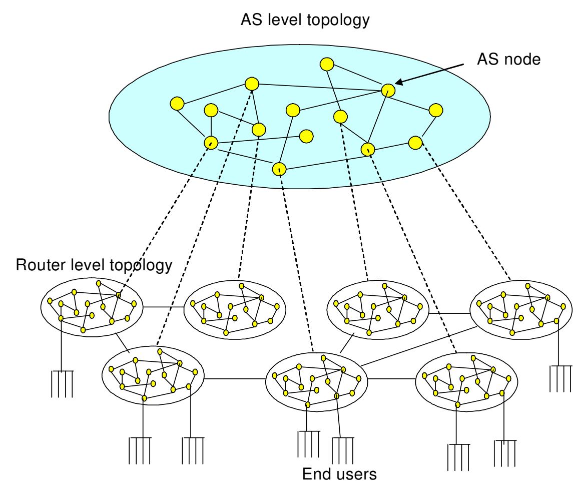

Figure 3. Abstract view of the topology used for simulation.

3. Simulation Results

We have implemented both the P2P and CDN approaches within the Network Simulator ns-2 [5]. A random Internet-like topology is generated using the GT-ITM topology generation tool as shown in figure 3. The topology has 3 levels which approximately resembles the Internet. The highest level (Autonomous System or AS level) represents the ISPs. Next, for each node in the AS-level topology we generate a router-level topology. We use an edge connection mechanism to interconnect router-level topologies as dictated by the connectivity of the AS-level topology. We then add a next level which corresponds to the end-user level. We generate a topology of 100 end nodes, each node corresponding to 10 peers. Thus a total of 1000 peers are simulated. The end user network access speed is 2Mbit/s. The nodes on the router level are connected with 20 Mbit/s links. Peers connected to the same router are not allowed to serve each other. At least one router has to be in between. A video file of 120 minutes duration (typical length of a Hollywood movie), having a data-rate of 100 kbit/s is selected for all simulations. As peers usually have asymmetric data rates (downstream rate being a number of times greater than the upstream rate), we assume that a serving peer can serve only one client at a time, supplying with a rate of 50 kbit/s. Thus, a receiving peer receives two descriptions from two supplying peers with an aggregate rate of 100 kbit/s. Each packet contains 1000 bytes. In both the CDN and P2P based systems, there is one new request every 30 seconds, originating from a randomly picked node. The total simulation time is 10 hours which corresponds to about 1200 streaming requests.

In the P2P network, the file is streamed from the two closest available peer nodes with complementary descriptions, whereas in the CDN, the two descriptions are served by the two closest CDN servers. It is assumed that a peer can serve only one request at one time, while a CDN server can serve a maximum of 200 streams simultaneously. This corresponds to a maximum serving rate of 10 Mbit/s per server.

Fig. 4 shows the results obtained through simulations. Performance parameters such as the rate of packet loss, number of non-decodable video frames, and the average response time, i.e., the time to receive the first video packet after the request has been sent, are compared for P2P and CDN networks. For the count of non-decodable frames, it is assumed that the descriptions contain an Intra frame once every second, and in case of a packet loss for the P-frames, all the subsequent frames become non-decodable, until the next I-frame is received. Because of MDC coding, the receiver can still display the video sequence with a reduced frame rate, unless both descriptions are corrupted simultaneously. This scenario is shown in Fig. 5, where description s1 contains a packet loss, but s2 is received error-free. The receiver can display the video at ½ the original frame rate until the next I-frame is received in s1.

Regarding the average response time, we can see in the first row of Fig. 4 that media streaming in a P2P network with good content replication exhibits a significantly lower initial delay compared to the CDN network.

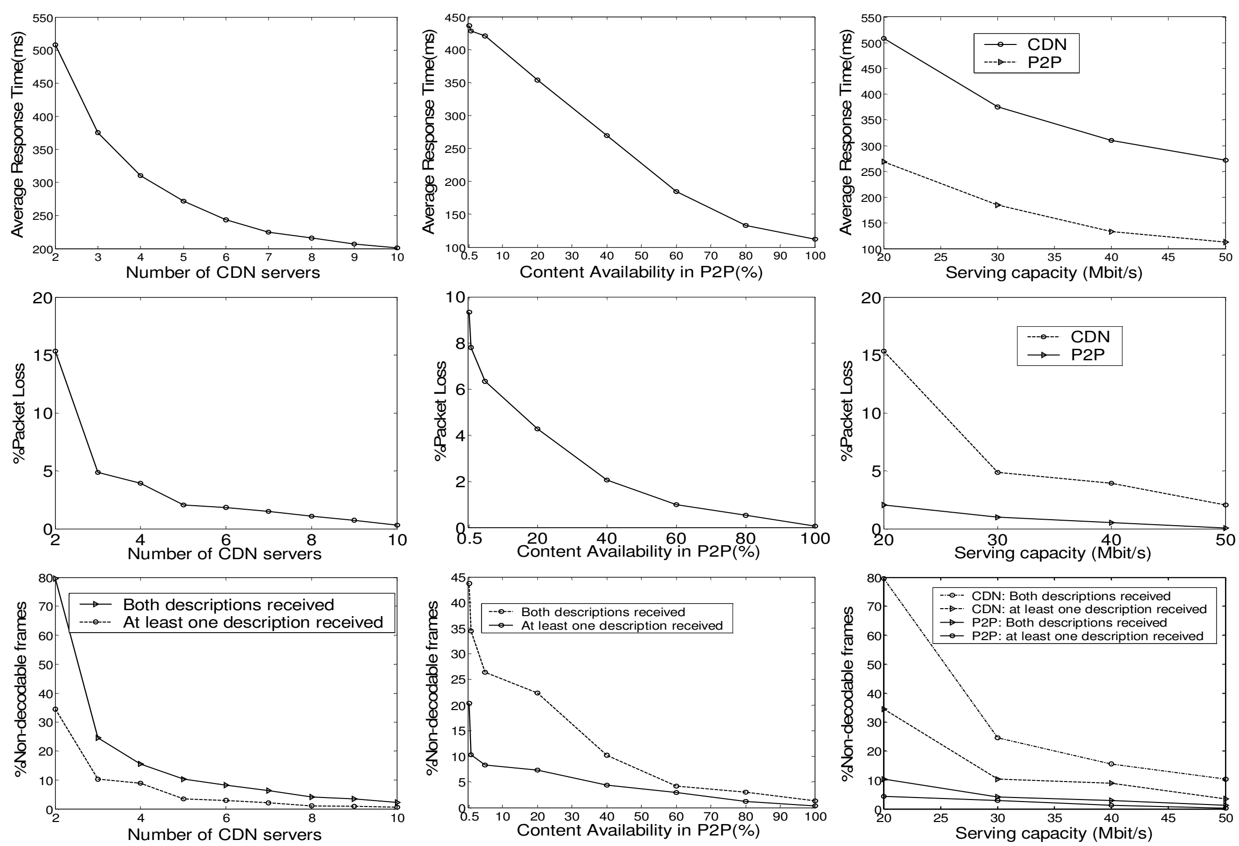

Figure 4. Performance of media streaming in P2P and CDN networks using MDC: (Row 1) average response times for P2P and CDN. (Row 2) Packet loss rate varies with varying number of CDN servers and content availability in P2P network. (Row 3) Number of decodable frames increases with increasing number of servers and availability. Serving capacity denotes the combined capacity of outgoing data rate of all the servers (peers sharing the video) available. This set of simulations was performed for moderate network load.

Additionally we can state that if more than 60% of the peers share the video, P2P based media streaming achieves lower initial delays than a CDN network with 10 servers.

To make a fair comparison of the CDN and the P2P approach, we employ as a basis the serving capacity of each network. The serving capacity is the aggregate available serving data rate of all CDN servers (for CDN case) or peers (for P2P case). Thus a serving capacity of e.g. 20 Mbit/s can either be provided by 400 peers (corresponds to 40% content availability in our simulation) or by 2 CDN servers. We can clearly see (Fig. 4, row 1) that the P2P approach always achieves a significantly smaller initial delay than the CDN approach. The reason for this is that the serving peers are typically much closer, in terms of the number of links, than the CDN servers. In a P2P network the requested content is automatically located where it is demanded, which is at the edge of the network. In a CDN, we need to have a large number of servers so that the content can also be located nearer to the user. This would result in very high costs, which do not occur in a P2P network.

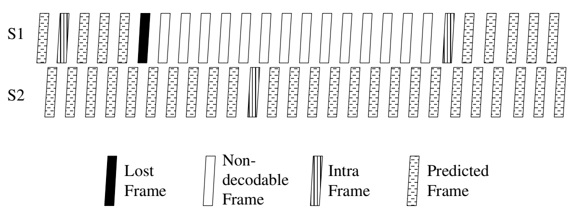

Figure 5. Impact of packet loss for MDC-based video streaming. Only stream s2 can be decoded completely. Description s1 is affected by packet loss which leads to locally reduced frame rate of the reconstructed video.

The phenomenon that in a CDN the content has to be retrieved from the core of the network, whereas it is available at the edge of the network in a P2P network, can also be observed in the packet loss graph of Fig. 4 (Row 2). In our P2P network, the percentage of lost packets is significantly lower than for the CDN. As the packet loss particularly impacts the number of non decodable frames, we can also observe in row 3 of Fig. 4 that the percentage of non decodable frames in the P2P case is significantly smaller than for a CDN at same serving capacity.

As the initial delay (average response time) and the percentage of non decodable frames are two of the major criterions for the user perceived quality, we can state that P2P based media delivery can considerably outperform CDN-based approaches. Even for low content availability (0.5%), lower residual loss rates, lower response time and fewer non-decodable frames can be achieved in comparison to a CDN with two servers, where we assume a 100% availability (see Fig. 4).

4. Conclusion

In this work we compare MDC-based video streaming in CDN and P2P networks. In our performance analysis we could show that a P2P based approach has a significantly better performance than a CDN approach. In a P2P network, the video can be provided with a smaller playout delay, a significantly smaller packet loss (2% instead of 15% in the CDN) and thus also a significantly smaller number of nondecodable frames (10% instead of 80% for a CDN). To guarantee comparability we conducted all simulations with the same total serving capacity. Besides these advantages, streaming via P2P networks does also not rely on additional servers in the network. The main reasons for this good performance of the P2P network are that the content can be accessed in the local proximity of each node, and the higher replication rate.

References

[1] T. Stading, P. Maniatis, M. Baker, “Peer-to-Peer Caching Schemes to Address Flash Crowds,” Proceedings of the 1st International Peer To Peer Systems (IPTPS ’02) Workshop, pp. 203 – 213, Mar. 2002.

[2] V. K. Goyal, “Multiple Description Coding: Compression Meets the Network,” IEEE Signal Processing Magazine, vol. 18, no. 5, pp. 74 – 94, Sept. 2001.

[3] T. Nguyen, A. Zakhor, “Distributed Video Streaming,” Proceedings of MMCN 2002, pp. 186 – 195, Jan. 2002.

[4] M. Hefeeda, B. Bhargava, “On-Demand Media Streaming Over the Internet,” CERIAS TR 2002-20, Purdue University, June 2002.

[5] The Network Simulator, http://www.isi.edu/nsnam/ns/

[6] J. Apostolopoulos, T. Wong, W. Tan, S. Wee, “On Multiple Description Streaming with Content Delivery Networks,” Proceedings of IEEE INFOCOM, June 2002.

[7] R. Baghwan, S. Savage, G, M. Voelker, “Understanding Availability,” Proceedings of the 2nd International Workshop on Peer-to-Peer Systems (IPTPS ’03), Feb. 2003.

[8] S. Saroiu, P. K. Gummadi, S. D. Gribble, “A measurement study of peer-to-peer file sharing systems,” Proceedings of Multimedia Computing and Networking (MMCN), Jan. 2002.

[9] E. Adar, B. Huberman. “Free Riding on Gnutella,” First Monday, vol. 5, no. 10, Oct. 2000.

[10] J. G. Apostolopoulos, “Reliable video communication over lossy packet networks using multiple state encoding and path diversity,” Proceedings of VCIP’01, pp. 392 – 409, Jan. 2001.

[11] M. S. Borella, D. Swider, S. Uludag, G. B. Brewster, “Internet Packet Loss: Measurement and Implications for End-to-End QoS,” Proceedings of International Conference on Parallel Processing, pp. 3 – 13, Aug. 1998.

[12] nCUBE Corporation homepage. www.ncube.com.

[13] J. Chu, K. Labonte, B. N. Levine, “Availability and locality measurements of peer-to-peer file systems,” in ITCom: Scalability and Traffic Control in IP Networks, Proceedings of SPIE, vol. 4868, Jul. 2002.

[14] Padmanabhan, H. Wang, P. Chou, and K. Sripanidkulchai, “Distributing streaming media content using cooperative networking,” In Proc of NOSSDAV’02, Miami Beach, FL, USA, May 2002.