

Caching Video Objects: Layers vs Versions ?

Felix Hartanto* Jussi Kangasharju† Martin Reisslein‡ Keith W. Ross†

*felix@ie.cuhk.edu.hk
Dept. of Information Engineering
Chinese University of Hong Kong
Shatin, N.T.
Hong Kong

†{kangasha, ross}@eurecom.fr
Institut Eurecom
2229, route des Cretes
06904 Sophia Antipolis
France

‡reisslein@asu.edu
Telecommunications Research Center
Dept. of Electrical Engineering
Arizona State University
Tempe, AZ 85287–7206

Abstract—Because Internet access rates are highly heterogeneous, many video content providers today make available different versions of the videos, with each version encoded at a different rate. Multiple video versions, however, require more server storage and may also dramatically impact cache performance in a traditional cache or in a CDN server. An alternative to versions is layered encoding, which can also provide multiple quality levels. Layered encoding requires less server storage capacity and may be more suitable for caching; but it typically increases transmission bandwidth due to encoding overhead. In this paper we compare video streaming of multiple versions with that of multiple layers in a caching environment. We examine caching and distribution strategies that use both versions and layers. We consider two cases: the request distribution for the videos is known *a priori*; and adaptive caching, for which the request distribution is unknown. Our analytical and simulation results indicate that mixed distribution/caching strategies provide the best overall performance.

Keywords—Proxy Caching, Streaming Video, Layered Video, Multi-version Video.

I. INTRODUCTION

MANY analysts expect streaming stored video to be the dominant traffic type in the Internet in the upcoming years, dwarfing the bandwidth usage of other Internet applications. Driving this demand is the current deployment of residential broadband access technologies, such as cable modem and xDSL technologies.

Given that the Internet will soon be transporting vast quantities of video traffic, a major concern becomes the efficient distribution of the video data. As with Web objects, video data can be transported to the client in many different ways, including (*i*) directly from origin server to client; (*ii*) through intermediate ISP caches; and (*iii*) through content distribution networks (CDNs) such as the Akamai network.

In designing new strategies for distributing stored video over the Internet, we also must take into account that access to the Internet is highly heterogeneous [1], [2]. Today, Internet access includes 56 kbps modem connections, 128

kbps ISDN connections, shared-bandwidth cable modem connections, xDSL connections with downstream rates in 100 kbps to 6 Mbps range, and high-speed switched Ethernet connections at 10–100 Mbps. Because Internet access is heterogeneous, video content providers typically provide multiple quality levels, with each quality level having a different encoding rate.

Multiple quality levels can be created by encoding video into multiple versions, each version encoded at a different rate. However, multiple versions of the same video can cause large increases in the amount of storage. For example, for storing 1000 videos, each video having an average length of one hour and having an average encoded bit rate of 4 Mbps, the required storage is 1800 Gbytes of storage. If for each video there is a second lower-quality version at half the bit rate of the high-quality version, then we need an additional 900 Gbytes of storage. When there are many versions, the additional required bandwidth can be yet much more.

Layered encoding (also known as hierarchical encoding) can also be used to create multiple quality levels. The storage requirements at a server for maintaining multiple layers is typically much less than maintaining the same number of versions. However, creating video layers generates additional bandwidth overhead [3], [4]. In particular, for the same quality level, layered encoding typically requires more transmission bandwidth than does a video version.

Given the presence of a caching and/or content distribution network infrastructure, and the need for multiple video quality levels, in this paper we compare distributing video versions to distributing video layers. We also examine mixed strategies consisting of both versions and layers.

Specifically, we consider a model in which all video content is encoded into two versions: a low-quality version and a high-quality version. All videos are also hierarchical encoded into a base layer and an enhancement layer. A proxy, representing an institutional cache or a server in a CDN, sits between the origin servers and the

clients. Bandwidth between the proxy and the clients is assumed to be abundant. However, bandwidth between the origin server and the proxy is a constrained resource, as well as is the storage capacity at the proxy. When the proxy receives a request for a video at a specific quality level, the proxy will directly satisfy the request if it has cached the appropriate version or layers; otherwise, if sufficient bandwidth is available between origin server and proxy, the origin server will stream the needed version or layers to the client through the proxy.

We first consider the case when the request distribution for the videos is known. We consider three natural distribution strategies and develop an analytical performance methodology. We then consider the case when the request distribution is unknown. We propose three natural adaptive caching studies and use simulation to compare their performance. Broadly speaking, we find that mixed strategies that use both versions and layers provide the most robust performance. Our model and methodology brings out a number of subtle issues that shed important insights on the distribution of multi-quality in the Internet.

This paper is organized as follows. We end this section with an overview of related work. In Section 2, we present our model and establish some basic properties of optimal caching strategies. In Section 3 we consider the case when the request distribution is known. We develop an analytical methodology, which we use to study the performance of three natural caching strategies. In Section 4, we consider adaptive caching and again study three natural distribution schemes. In the concluding Section 5, we summarize our findings.

A. Related Work

Decuetos *et al* [5] also compared streaming of video versions to streaming of video layers. In particular, in a TCP-friendly context, they proposed prefetching and quality-level switching schemes for both pure versions and pure layers. The paper [5] focused on time-dependent streaming of a single video from origin server to client; it did not take into account an intermediate cache sitting between origin servers and clients.

Kangasharju *et al* [7] considered caching strategies for layered video. In particular, they formulated the problem as optimization problem, showed that the optimization problem was intractable, and proposed and studied several natural heuristics. The paper [7] did not take into account multiple versions, and therefore did not compare caching layers, caching versions, and mixed strategies.

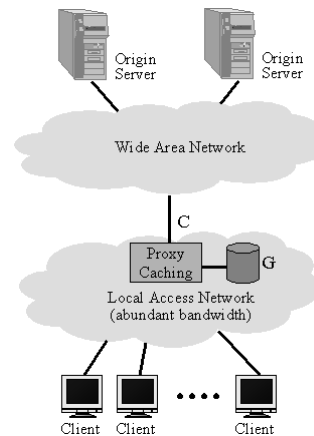


Fig. 1. Architecture for caching and streaming of adaptive video.

II. MODEL AND NOTATION

Fig. 1 illustrates our architecture for video caching. Suppose there are M videos available; and all of them are stored on the origin servers. Popular videos are cached in a proxy server, which is located close to its client community.

A. Proxy Server

The clients direct their requests to the proxy server; if the requested video (defined by type and quality) is in the proxy, then the video is streamed from the proxy to the client; if it is not in the proxy, the video is streamed from the origin server to the proxy, and then from the proxy to the client.

The proxy server is connected to the origin servers via a wide area network (e.g., the Internet). We model the bandwidth available for streaming from the origin servers to the proxy server as a bottleneck link of fixed capacity C (bit/sec). The proxy is connected to the clients via a local access network, which could be a LAN running over Ethernet, or a residential access network using xDSL or HFC technologies. For the purposes of this study, we assume that there is abundant bandwidth for streaming from the proxy to the clients. We model the proxy server as having a storage capacity of G (bytes) and having infinite storage bandwidth (for reading from storage). Our focus in this study is on caching strategies that cache complete layers or versions of videos in the proxy. Our goal is to cache video layers or versions so as to maximize the number of supported streams.

B. Versions

Real Networks [9] and other video streaming technology companies today allow content providers to encode

video into multiple quality versions. Video versions allow service and content providers to offer flexible streaming services to clients with vastly different access bandwidths and decoding capabilities. Clients with low-speed access will only be interested in the low-quality stream. Clients with LAN, cable-modem or ADSL access will be interested in high-quality streams.

Many content providers today store multiple versions of the same video on the origin server and stream the video version that is most appropriate on a user-to-user basis. This approach allows for flexible pricing structures. A content or service provider may offer the low-quality version for a standard charge and charge a premium for the high-quality version. Throughout this paper we shall assume that two quality levels are available for each video.

Although the approach of multiple versions offers greater service and pricing flexibility, it has major drawbacks. First, it requires more storage at the origin server than does the approach that makes only one quality level available. Second, if one quality version is cached in a proxy server, and there is a request for a different quality version, then the version must be fully streamed from the origin server, i.e., the cached version is of no use. And third, if both quality levels are cached in the proxy, then more storage is necessary than when only one version is used.

C. Layers

An alternative to using versions is to use layered (also known as hierarchical encoded) video. With layered encoding, each video object is encoded into a base layer and one or more enhancement layers. The base layer contains the most essential information, and the enhancement layers provide quality enhancements. A particular enhancement layer can be decoded only if all lower quality layers are available. Throughout this paper, we will assume that each video has been coded into two layers, a base layer and a single enhancement layer.

The storage requirements for the base and enhancement layer together are typically less than the requirements for the low-quality and high-quality versions together, for both the origin server and the proxy. Furthermore, if the base layer is cached in the proxy, and a client requests a high-quality version, then only the enhancement layer needs to be streamed from the origin server. Nevertheless, layered encoding has one major drawback, namely, encoding overhead. Typically, for the same high-quality level, the total rate of the base and enhancement layer combined is greater than the rate of the high-quality version. Also, for the same low-quality level, the rate of base layer is often greater than that of the low-quality version. This over-

head impacts both transmission and storage resources.

In summary, in this study each video can be encoded into either versions or layers. For versions, we suppose that there are two possible versions, namely, a high-quality version and a low-quality version. For layers, we suppose that the video is encoded into two layers, namely, a base layer and an enhancement layer. Thus, each video has four objects associated with it: a low-quality version, a high-quality version, a base layer, and an enhancement layer. We denote these four objects by l , h , b , and e , respectively.

If $T(m)$ is the length of video m , $m = 1, \dots, M$, in seconds and $r(m)$ is the encoding rate for one of the versions or layers, then the corresponding storage requirement for the object is $S(m) = T(m) \cdot r(m)$. Table I summarizes the notation we will use for the two versions and the two layers. We naturally assume that the rate of the high-quality version is greater than the rate of the low-quality version, i.e., $r_h(m) > r_l(m)$.

In order to compare the caching of layers and versions, we suppose throughout that the encodings are such that the visual quality of the base layer is the same as the visual quality of the low-quality version; and the video quality of the base and enhancement layer combined is the same as the high-quality version. However, due to encoding overhead to create layers, we do not assume that the layers and versions have the same encoding rates. Instead, we make the following three natural *Rate Assumptions* which are based on video encoding experiments [3], [4]:

1. Due to the overhead of layered encoding, the base layer has at least the same rate as the low-quality version, i.e., $r_b(m) = r_l(m) \cdot [1 + O_l(m)]$ where $O_l(m) \geq 0$ is the low-quality coding overhead.
2. Again due to the overhead of layered encoding, the base and enhancement layers together have at least the same rate as the high-quality version, i.e., $r_b(m) + r_e(m) = r_h(m) \cdot [1 + O_h(m)]$ where $O_h(m) \geq 0$ is the high-quality coding overhead.
3. The base and enhancement layers together have smaller rate than the two versions, i.e., $r_b(m) + r_e(m) < r_l(m) + r_h(m)$.

For any video, the proxy can contain objects made from versions and/or layers. However, we assume the *decoding constraint*, namely, that the proxy never caches the enhancement layer if the base layer is not cached. When a request arrives to the proxy for some low-quality video, the proxy can satisfy the request if it is currently storing either the low-quality version or the base layer of the video. Otherwise, the proxy must obtain either the low-quality version or the base layer from the origin server and relay the object to the requesting client. When a request arrives to the proxy for some high-quality video, the proxy can

	Base layer	Enh. layer	Low quality	High quality
Encoding rate	$r_b(m)$	$r_e(m)$	$r_l(m)$	$r_h(m)$
Size	$S_b(m)$	$S_e(m)$	$S_l(m)$	$S_h(m)$

TABLE I

NOTATION: RATES AND STORAGE REQUIREMENTS OF LAYERS AND VERSIONS OF VIDEO m , $m = 1, \dots, M$.

satisfy the request if it is currently storing either the high-quality version or if it is storing both the base and enhancement layers of the video. Otherwise, it must retrieve an object from the network to satisfy the request. If the proxy has stored the base layer, then the proxy can retrieve either the enhancement layer or the high-quality version.

D. Basic Properties

For a given video, there are four cachable objects: the low-quality version (l), the high-quality version (h), the base layer (b) and the enhancement layer (e). Thus for any given video, there are $2^4 = 16$ different combinations of objects that can be put in the cache, including putting no object in the cache. This is a daunting number of combinations to analyze. Fortunately, without loss of generality, we may restrict ourselves to only five of the combinations:

Theorem 1: There is an optimal caching configuration such that for each video one of the following five object combinations is used: \emptyset , $\{l\}$, $\{h\}$, $\{b\}$, or $\{b, e\}$. In other words, for each given video we either cache just the low-quality version, just the high-quality version, just the base layer, the base and enhancement layers together, or no objects at all.

Proof: Because of the decoding constraint for layered video, we can rule out all combinations that include e but not b .

Now consider $\{b, h\}$. Note that Rate Assumptions 3 and 1 together imply that $r_h(m) > r_e(m)$. Hence $r_b(m) + r_h(m) > r_b(m) + r_e(m)$. It follows from this last expression that we can replace the combination $\{b, h\}$ with $\{b, e\}$ and use less storage while still satisfying all requests at the proxy for the video. Thus we can rule out $\{b, h\}$.

Now consider $\{b, l\}$, $\{b, l, e\}$, $\{b, l, h\}$, $\{b, l, h, e\}$. By caching the base layer, we satisfy all low-quality requests and we partially satisfy higher quality requests (only need to get enhancement layer from network). If we additionally cache the low-quality version, we take up more storage and we do not satisfy more requests for low-quality video. Combining this observation with $r_h(m) > r_e(m)$ implies that if we cache the base layer, then there is no need to also cache the low-quality layer.

Thus we can rule out all these four cases.

Now consider $\{l, h\}$. This combination will satisfy all requests at the proxy. However, the combination $\{b, e\}$ also satisfies all requests and, by Rate Assumption 3, takes less storage. Thus, we can rule out $\{l, h\}$.

Finally, we can also rule out $\{b, e, h\}$ since the combination $\{b, e\}$ also satisfies all requests but takes less storage. ■

As a corollary to the above theorem, for any given video we use either versions or layers but not both.

Motivated by the above theorem, in the following sections we will propose and examine some strategies for caching layer and version objects. But it is also useful to make a few additional *Observations* about extreme cases:

1. For a given video if all (or “nearly all”) requests are for the low-quality version (and none or “nearly none” are for the high-quality version), then we would either cache the low-quality version or cache no objects for that video, i.e., for object combination we would use either $\{l\}$ or \emptyset .
2. Similarly, if for a given video if all (or “nearly all”) requests are for the high-quality version, we would use either $\{h\}$ or \emptyset .
3. If there is no overhead for layered encoding, that is, if $O_l(m) = O_h(m) = 0$, then for video m we would only use layers; in particular, we would use either \emptyset , $\{b\}$ or $\{b, e\}$.

However, when (i) there is layering overhead, and (ii) request rates for low- and high-quality versions are both significant, then it is not obvious whether we should use versions or layers; furthermore, for some videos it may be preferable to use versions whereas for others it may be preferable to use layers.

III. KNOWN REQUEST DISTRIBUTION

We start by modeling the steady-state cache performance using a static caching model. With this model, we assume that the request pattern is known *a priori* and does not change dynamically. Suppose that there are M videos. Suppose that requests for video streams arrive according to a Poisson process with rate λ (requests/hour). Let j denote the requested quality level with $j = 0$ indicating a request for a low quality video, and $j = 1$ indicating a request for a

high quality video. Let $p(j, m)$, $j = 0, 1$; $m = 1, \dots, M$, denote the probability that a given request is for the j -quality stream of video m . As a proper mass distribution the $p(j, m)$'s satisfy $\sum_{m=1}^M \sum_{j=0}^1 p(j, m) = 1$.

The corollary to Theorem 1 suggests three caching strategies, namely:

1. Pure version caching, where we cache only video versions.
2. Pure layer caching, where we cache only video layers.
3. Mixed caching, where we cache layers for some videos and versions for others.

For all three caching strategies we first order the request probabilities $p(j, m)$, $j = 0, 1$; $m = 1, \dots, M$ in decreasing order. We then fill the cache by considering the objects (j, m) that are the most requested. First, we put the object (j, m) with the largest request probability $p(j, m)$ into the cache. Next, we cache the object (j, m) with the next largest probability $p(j, m)$, and so on. If at some point (as the cache fills up) the object needed to satisfy the request with the next largest request probability does not fit into the remaining cache space, we skip this object and try to cache the objects with the next largest request probabilities.

With *pure version caching* we cache only versions of the videos. We cache the high quality version of video m if the next largest probability $p(j, m)$ is for the high quality stream of video m (i.e., $j = 1$). If the next largest probability is for the low quality stream of video m , then we cache the low quality version of video m . Note that with pure version caching we may end up caching both high and low quality versions of the same video (which we know from Theorem 1 is sub-optimal).

With *pure layer caching* we cache only video layers. If the next largest request probability $p(j, m)$ is for low quality stream of video m (i.e., $j = 0$), then we cache the base layer of video m . On the other hand, if the next largest probability is for the high quality stream of video m (i.e., $j = 1$), then we cache both base and enhancement layer of video m . If the base layer has already been cached, i.e., if $p(0, m) > p(1, m)$, then we need to cache the enhancement layer only. Due to decoding constraint, we never cache the enhancement layer of a given video without caching the corresponding base layer.

With *mixed caching* we cache the high quality version of video m if the next largest $p(j, m)$ is for the high quality stream of video m and no other object of the video has been cached. On the other hand, if the next largest probability is for the low quality stream of video m and no other object of the video has been cached, then we (i) cache the low version of video m if $r_b(m) > r_l(m)$, and (ii) cache the base layer of video m if $r_b(m) = r_l(m)$. However, if

we have already cached the low (or high) quality version of a given video and the next largest probability is for a different quality of the video, then we replace the low (or high) quality version of the video with the base and enhancement layer of the video.

A. Video Caching Model

In this section we develop an analytical model for the caching and streaming of video layers and versions. We derive expressions for the blocking probability of a client request and the long run rate at which client requests are satisfied. To keep track of the objects in the cache we introduce a vector of cache indicators $\mathbf{c} = (c_1, c_2, \dots, c_M)$, with $c_m = \{\emptyset\}, \{l\}, \{h\}, \{l, h\}, \{b\}$, or $\{b, e\}$, for $m = 1, \dots, M$. c_m indicates whether no object, the low-quality version, the high-quality version, both the low- and high-quality version, the base layer, or the base layer together with the enhancement layer is cached for video m . (We allow for $c_m = \{l, h\}$ to accommodate pure version caching in our model; note, however, that by Theorem 1 it is sub-optimal to cache both the low- and high-quality version for a given video m .) In our model we focus on the bottleneck link of capacity C , that connects the proxy server to the origin servers. We model this link as a stochastic knapsack [8]. Let $b_{c_m}(j, m)$, $j = 0, 1$, $m = 1, \dots, M$, denote the link capacity required for satisfying a request for a j -quality stream of video m , given that the object(s) c_m are cached for video m . Table II gives the $b_{c_m}(j, m)$'s for all possible combinations of c_m and j . We assume that the lower rate versions are streamed over the bottleneck link whenever a request cannot be satisfied by the cache; except in the case where the base layer is cached and the high-quality stream is requested, in that case we stream the enhancement layer. Without loss of generality we assume that C and all $b_{c_m}(j, m)$'s are positive integers. Let $\mathbf{b}_c = (b_{c_m}(j, m))$, $j = 0, 1$, $m = 1, \dots, M$, be the vector of the bandwidth requirements of the requests. Note that this vector has $2M$ elements. Throughout we assume that the client watches the entire stream without interruption, thus the bandwidth $b_{c_m}(j, m)$ is occupied for $T(m)$ seconds. Let $\mathbf{n} = (n(j, m))$, $j = 0, 1$, $m = 1, \dots, M$, be the vector of the numbers of ongoing j -quality streams of video m . The $n(j, m)$'s are non-negative integers. Let $\mathcal{S}_c = \{\mathbf{n} : \mathbf{b}_c \cdot \mathbf{n} \leq C\}$ be the state space of the stochastic knapsack model of the bottleneck link, where $\mathbf{b}_c \cdot \mathbf{n} = \sum_{m=1}^M \sum_{j=0}^1 b_{c_m}(j, m) \cdot n(j, m)$. Furthermore, let $\mathcal{S}_c(j, m)$ be the subset of states in which the knapsack (i.e., the bottleneck link) admits a stream with the bandwidth requirement $b_{c_m}(j, m)$. We have $\mathcal{S}_c(j, m) = \{\mathbf{n} \in \mathcal{S}_c : \mathbf{b}_c \cdot \mathbf{n} \leq C - b_{c_m}(j, m)\}$. The blocking probabilities

$b_{c_m}(j, m)$	$c_m = \{\emptyset\}$	$c_m = \{l\}$	$c_m = \{h\}$	$c_m = \{l, h\}$	$c_m = \{b\}$	$c_m = \{b, e\}$
$j = 0$	$r_l(m)$	0	$r_l(m)$	0	0	0
$j = 1$	$r_h(m)$	$r_h(m)$	0	0	$r_e(m)$	0

TABLE II

BANDWIDTH REQUIREMENT FOR STREAMING j -QUALITY STREAM OF VIDEO m GIVEN CACHE CONFIGURATION $b_{c_m}(j, m)$.

can be explicitly expressed as

$$B_{\mathbf{c}}(j, m) = 1 - \frac{\sum_{\mathbf{n} \in \mathcal{S}_{\mathbf{c}}(j, m)} \prod_{m=1}^M \prod_{j=0}^1 (\rho(j, m))^{n(j, m)} / (n(j, m))!}{\sum_{\mathbf{n} \in \mathcal{S}_{\mathbf{c}}} \prod_{m=1}^M \prod_{j=0}^1 (\rho(j, m))^{n(j, m)} / (n(j, m))!},$$

where $\rho(j, m) = \lambda p(j, m) T(m)$ is the load offered by requests for j -quality streams of video m . These blocking probabilities can be efficiently calculated using the recursive Kaufman–Roberts algorithm [8, p. 23]. The expected blocking probability of a client's request is given by

$$B(\mathbf{c}) = \sum_{m=1}^M \sum_{j=0}^1 p(j, m) B_{\mathbf{c}}(j, m).$$

The long run throughput, i.e., the long run rate at which client requests are satisfied is given by

$$TH(\mathbf{c}) = \lambda \cdot \sum_{m=1}^M \sum_{j=0}^1 p(j, m) (1 - B_{\mathbf{c}}(j, m)).$$

We define the normalized throughput $TH_n(\mathbf{c})$ as the ratio of the rate of satisfied requests to the total request arrival rate, i.e., $TH_n(\mathbf{c}) = TH(\mathbf{c})/\lambda$.

B. Numerical Results

We assume that there are $M = 1,000$ different videos. For a given video m we generate the version and layer rates as follows. The rate of the high quality version $r_h(m)$ is drawn randomly from a uniform distribution between 2 and 6 Mbps with a granularity of 0.1 Mbps and an average of 4 Mbps. The rate of the low quality version $r_l(m)$ is uniformly drawn between $0.5 \cdot r_h(m)$ and $0.7 \cdot r_h(m)$ with an average of $0.6 \cdot r_h(m)$. The length of the video $T(m)$ is drawn from an exponential distribution with an average length of one hour.

We assume that the aggregate rate for the layered video has an overhead $O_h(m)$ over the high quality version, i.e., $r_b(m) + r_e(m) = [1 + O_h(m)] \cdot r_h(m)$. We consider two cases: (i) $r_b(m) = r_l(m)$, and (ii) $r_b(m) > r_l(m)$, in this case we vary $r_b(m)$ between $r_l(m)$ and $[1 + O_h(m)] \cdot r_l(m)$. With $r_b(m)$ fixed, the rate of the enhancement layer

$r_e(m)$ is then computed as $r_e(m) = [1 + O_h(m)] \cdot r_h(m) - r_b(m)$.

Client requests arrive according to a Poisson process. The average request arrival rate is $\lambda = 270$ requests/hour, chosen to give a blocking probability of 2% when $q = 1.0$. The $p(j, m)$'s are determined as follows. Let p_m , $m = 1, \dots, M$, denote the probability that a given client request is for video m (irrespective of whether the request is for the low quality stream or the high quality stream of the video). We draw the p_m 's from a Zipf distribution with parameter $\zeta = 1$. Let q denote the probability that the request for a given video is for the low quality stream of the video. We fix q as a system parameter in our numerical analysis. We set $p(0, m) = q \cdot p_m$ and $p(1, m) = (1 - q) \cdot p_m$.

The cache size is set to $G = 200$ Gbytes and the link capacity is $C = 150$ Mbps. For a given realization of the layer and version rates ($r_l(m)$, $r_h(m)$, $r_b(m)$, $r_e(m)$) as well as video lengths $T(m)$, $m = 1, \dots, M$, we apply the three outlined caching strategies to obtain the cache indicators c_m , $m = 1, \dots, M$. With these cache indicators we calculate the normalized throughput using the stochastic knapsack analysis introduced in the previous section. We run many independent replications of this procedure to obtain confidence intervals for the normalized throughput. For every independent replication we draw a new independent set of layer and version rates and video lengths. We repeat this procedure until the 95% confidence interval of the normalized throughput is less than 1% of the corresponding sample mean.

In Fig. 2 we plot the normalized throughput as a function of the probability of a low quality request q . The results show that if no overhead is incurred in generating layered videos (i.e., $O_h = 0$), then pure layer caching is the best strategy as suggested by Observation 3 above. Caching layers is also favorable when the requests are non-homogeneous ($0.1 < q < 1$) and the overhead is low. We see that the throughput for pure layer caching increases monotonically as more requests are for low quality videos and decreases with increasing overhead. The throughput for pure layer caching is strongly affected when the base layer includes overhead (i.e., $r_b > r_l$). This can be considered as the worst case and therefore, we always assume

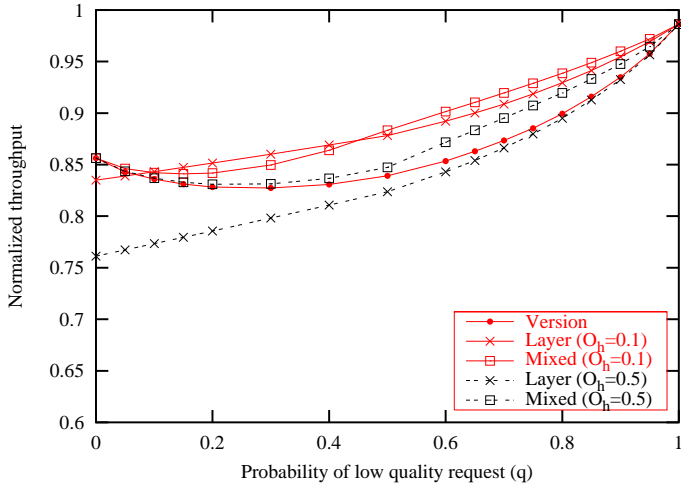
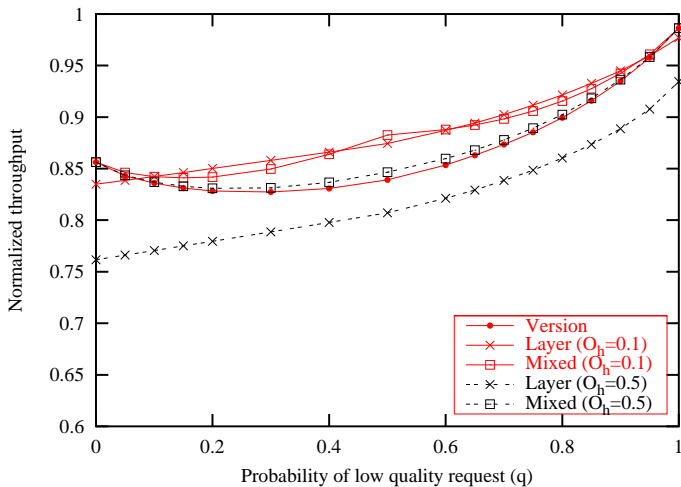
(a) $r_b = r_l$ (b) $r_b > r_l$

Fig. 2. Static caching scenario with varying probability of low quality requests.

$r_b > r_l$ in future plots.

Pure version caching is only favorable in case of homogeneous request quality, i.e., all requests are either for low quality ($q = 0$) or for high quality streams ($q = 1$). The largest throughput is achieved if all requests are for low quality streams. This is expected because in this scenario more videos are cached and hence the cache hit rate is higher compared to a scenario where all requests are for high quality streams. The throughput is lowest when the requests are non-homogeneous as sometimes we need to cache both the low- and the high-quality version.

The results indicate that mixed caching strikes a good balance between pure layer caching and pure version

caching for all cases and offers the best overall performance. It performs as well as pure layer caching when the overhead is zero and as well as pure version caching when $O_h = 0.5$. Since the smallest r_l is $0.5 \cdot r_h$, $O_h = 0.5$ is the largest overhead incurred in creating layered video while meeting Rate Assumption 3.

Fig. 3 gives the normalized throughput as a function of the overhead O_h of layered encoding. We can clearly see that mixed caching gives better performance than pure version caching and pure layer caching for the range of overhead. Its performance is less sensitive to the overhead than pure layer caching.

The superiority of mixed caching is independent of the cache size and the link capacity. In Fig. 4 we plot the normalized throughput as a function of the cache size G and the link capacity C . The cache size is chosen between $G = 45$ and 900 Gbytes or between 2.5% and 50% of the total video data. Given the average video length T_{avg} (in seconds), the average rate of a video r_{avg} (in bit/sec), and the client request rate λ (in requests/second), we would need on the average $C = T_{avg} \cdot r_{avg} \cdot \lambda$ Mbps of bandwidth to stream all the requested videos. We varied the link capacity between $C = 10$ and 160 Mbps or between 1% and 16% of the total requested video bit rate.

Both figures show that in all cases mixed caching offers the best overall performance. It shows that mixed caching gives similar performance to pure layer caching for small overhead and similar performance to the pure version caching for $O_h = 0.5$.

In summary, the results with static caching model show that a mixed caching strategy can strike a good balance between pure layer caching and pure version caching.

IV. ADAPTIVE CACHING

With the static caching model, the request distribution is assumed to be known beforehand. However, in practice, the actual request distribution may not be known. When the distribution is unknown, we need to make caching and replacement decisions on-the-fly. Moreover, in most video distribution systems, new videos are being continually released. As the video popularities change, providers replace the least popular videos in their systems with new videos. In this section we will consider adaptively caching and replacing videos when the request distribution is unknown and new videos are being continuously released. We will compare the performance differences of static caching (with known distributions) and adaptive caching, and identify the factors causing the differences. We will also investigate whether the basic observations of Section 2 still apply.

In order to allow for direct comparison with the static

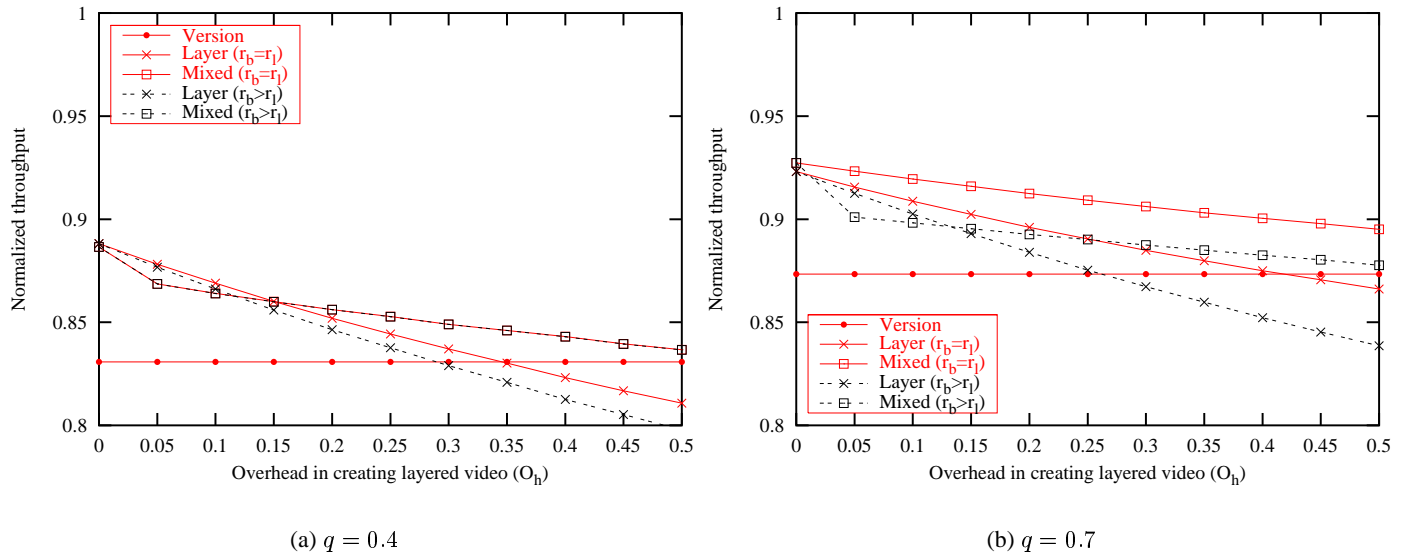


Fig. 3. Static caching with varying amount of overhead of layered encoding.

caching scenario, we model the dynamic request distribution as follows. We start with the same 1000 videos and generate requests in a similar fashion as in the static caching scenario. However, in this adaptive model, we assume that a fresh set of videos is made available periodically and the least popular videos are replaced by this new set of videos. We assume that 1 to 50 new videos are released every week and that the exact number of new videos is uniformly distributed between 1 and 50. The characteristics of these new videos follow the same distribution as used in Section III-B. Once the least popular videos are replaced by these new videos, the popularity of all videos in the system are re-shuffled and requests are generated based on the new popularity distribution. Upon re-shuffling, we also evict all currently cached objects from the videos that have been replaced. If a stream is currently using the objects, then we remove the objects as soon as the ongoing stream finishes.

Now, we explain the caching strategies. We start with an empty cache and cache the layer or version of a video as it is requested and streamed to the client. If the cache is full, then we replace the video in the cache following a *least recently used (LRU)* replacement strategy. We replace videos in the cache until enough space is obtained. In all strategies, we do not replace a video object from the cache if the object is currently being used for streaming the video. In the following we describe the three caching strategies identified in Section III.

With *pure version caching*, we cache the high-quality version if the high quality video is requested and that version is not in the cache regardless of whether we have the

low version in the cache or not. Likewise, we cache the low-quality version if the low quality video is requested and that version is not in the cache. Again, we do it regardless of whether we have the high-quality version in the cache or not. Therefore, we can have both high- and low-quality versions in the cache (which we know from Theorem 1 is suboptimal).

With *pure layer caching*, we cache both base and enhancement layer if the high quality video is requested and the video is not in the cache. If we already have the base layer, then we only stream the enhancement layer from the origin server and cache it. We stream the base layer from the origin server and cache it if the request is for low quality video and the base layer is not cached. During replacement, we remove the enhancement layer before the base layer.

With *mixed caching*, we have a similar objective as in the static caching model. We basically want to reduce the resource usage by mixing layers and versions in the cache. Here, we consider two simple heuristics to illustrate our findings.

The *first heuristic* corresponds exactly to mixed caching in the static model. Its objective is to replace the caching of both high and low version of a video with the layers of the video since they use less resources as $r_b(m) + r_e(m) < r_h(m) + r_l(m)$. The caching proceeds as follows. For the first request we stream and cache the version of the video. So, if the request is for high-quality video, we stream and cache the high-quality version of the video. If the request is for the low-quality video, then we stream and cache the low-quality version if $r_b(m) > r_l(m)$, or the base layer

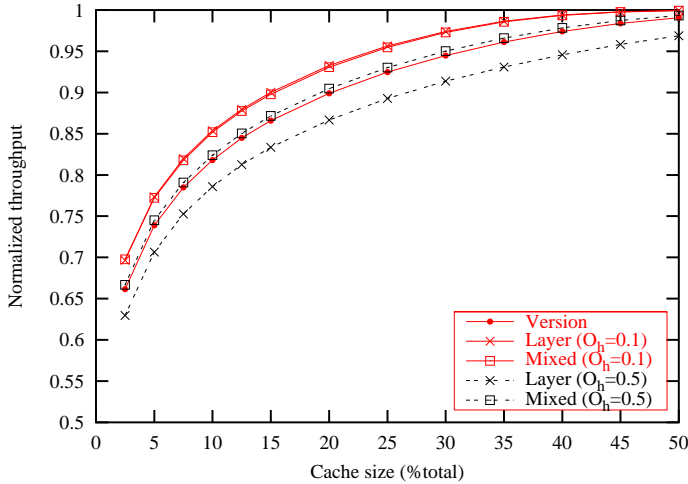
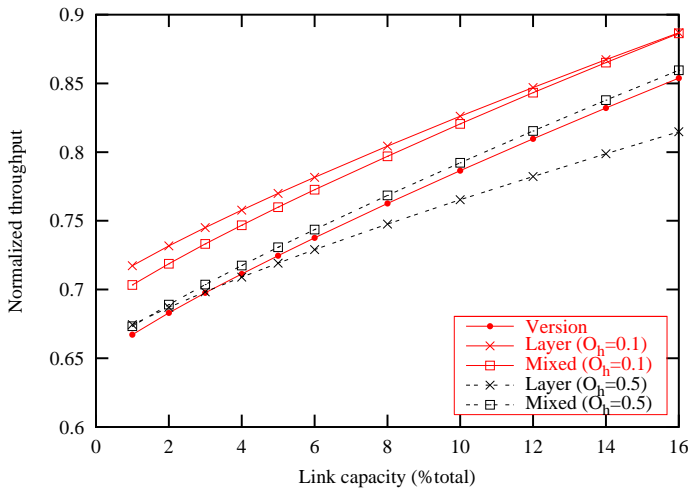
(a) varying cache size G (b) varying link capacity C

Fig. 4. Static caching with varying cache size G and link capacity C ($q = 0.4$).

if $r_b(m) = r_l(m)$. If there is a second request for different quality level of the same video, then we try to satisfy the request with layers and remove the version from the cache. Otherwise, we proceed with pure version caching. In replacing the version by layers, we reject the request and keep the version if we do not have enough link capacity to stream the layers. Moreover, since we do not want to interrupt ongoing streams, we cannot remove the version if it is being used, but we still cache the layers. Once we have cached the layers, the version will be removed as soon as the ongoing streams using that version are terminated. If both the base and enhancement layers of a video are removed from the cache (by LRU replacement), then

we will start again with the streaming and caching of versions for the next request. The motivation is that if a video object can be removed from the cache, then the video object is probably not very popular. So, it is better to start again with versions.

The *second heuristic* is similar to the first heuristic, except that for the first request we stream the version of the video but we do *not* cache it. So, if the request is for high quality video, then we stream the high-quality version of the video but do not cache it, and if it is for low-quality video, then we stream the low-quality version of the video but do not cache it. If there is a second request for the same video then we cache layers of the video. If the second request is for high-quality video then we stream and cache both base and enhancement layers. If it is for low-quality video, then we stream and cache the base layer only. In this way, we stream versions but never cache them. Instead, we cache only layers. The motivation of this heuristic is to avoid caching objects for videos which are requested once only. Moreover, for videos which are requested more than once, caching video layers can serve requests of different quality while using less resources. This caching strategy requires the proxy to keep track of videos that have been previously streamed. If all layers of a video are removed from the cache, then we will start again, streaming but not caching versions, and caching layers upon second requests.

A. Numerical Results

We now present simulation results for adaptive caching. We use the same distributions for the layer and version rates as well as the video lengths as were used in Section III-B. While we evaluated the normalized throughput with the stochastic knapsack analysis in Section III-B, we now obtain the normalized throughput from simulations of the cache operation. We use sequential simulation [?] to stop a simulation run automatically once the 95% confidence interval is reached or it has run for 10^8 seconds. We then repeat the simulation by using different seeds. This ensures a different mixture of videos and hence cache composition. The final results are obtained by averaging the values from all runs. The simulation run are repeated until the final results with 95% confidence intervals across different video mixtures are reached.

Fig. 5 gives the normalized throughput as a function of the probability of a low quality request q . The figure shows that pure version caching is only favorable in case of homogeneous requests. For heterogeneous requests, pure layer caching offers better performance than pure version caching, especially when the layering overhead is low and no overhead is incurred in creating the base layer. As with

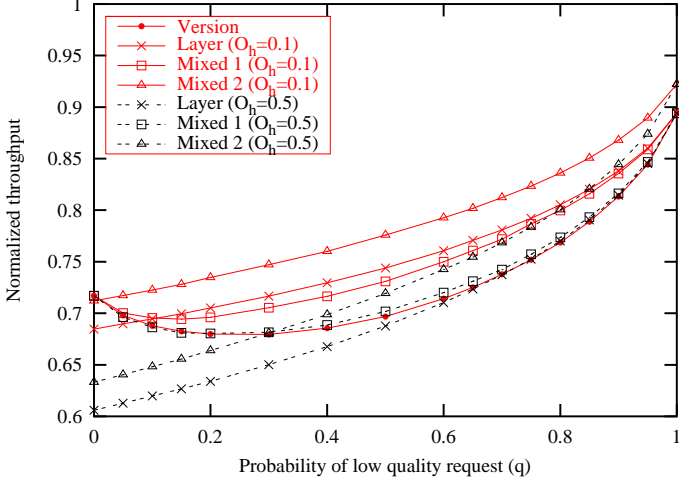
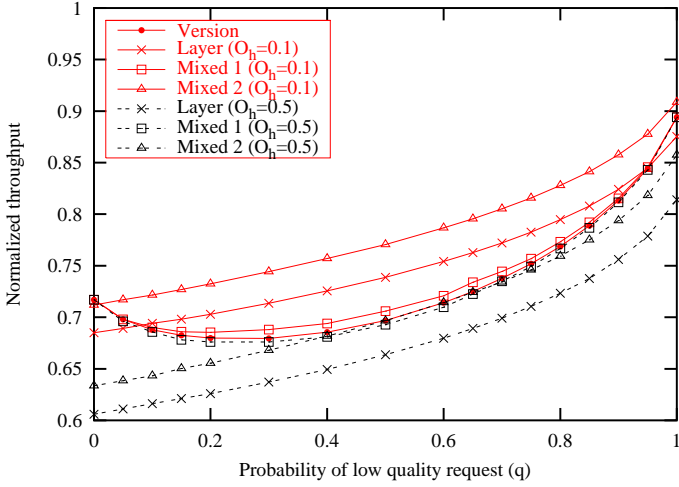
(a) $r_b = r_l$ (b) $r_b > r_l$

Fig. 5. Adaptive caching for varying probability of low quality request.

the static model, we see that mixed caching — using both heuristics 1 and 2 — provides a good balance between pure layer and pure version caching. It performs better than pure layer caching for small overhead and as well as pure version caching for large overhead. We also observe that heuristic 2 gives excellent results for a small layered encoding overhead. Note that heuristic 2 can be considered a variation of pure layer caching where we require to see two requests before caching layers of a video. Throughout, heuristic 2 performs much better than pure layer caching. This demonstrates the importance of weeding out the one-timer requests.

Comparing the plots with Fig. 2 we notice that in gen-

eral the throughput for adaptive caching is smaller than the throughput for static caching. This is mainly because the request pattern is not known *a priori* in the adaptive caching model. In adaptive caching, videos are (1) cached as requests arrive, and (2) evicted from the cache when there is not enough space for new video objects. Thus the order of the request arrivals has a strong impact on the cache composition, whereas the cache composition is exclusively based on the stream popularities in the static caching model. The difference in performance between static caching and adaptive caching widens as the average request arrival rate λ increases, as is illustrated in Fig. 6. This can be explained as follows. Consider a cache with a

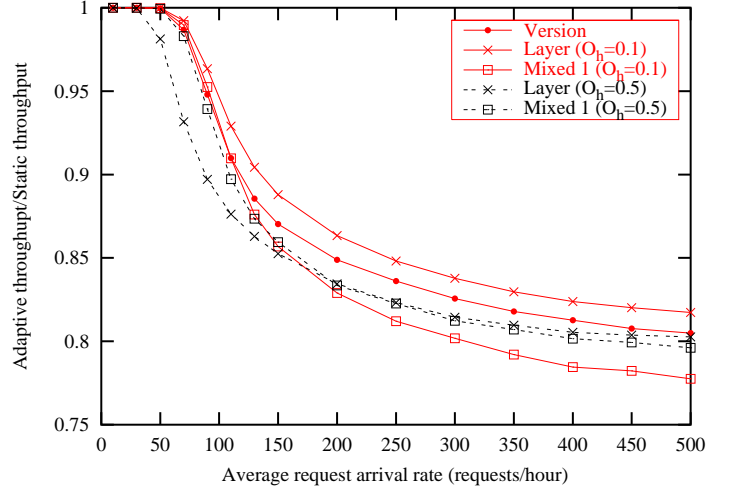


Fig. 6. Differences in throughput between adaptive and static caching ($q = 0.4$).

large request arrival rate. and suppose that a “mistake” has been made by caching a moderately popular object. With a large request arrival rate even a moderately popular object could receive enough requests to have continuously one or more ongoing streams. These ongoing streams, however, keep the object in the cache and prevent more popular objects (which would have been cached in the known request distribution scenario) from entering the cache.

Fig. 7 gives the normalized throughput as a function of the amount of overhead incurred in layered encoding. We observe that heuristic 2 offers the best overall performance. However, similar to pure layer caching, it is highly sensitive to the overhead. On the other hand, heuristic 1 behaves similar to pure version caching and is hence less sensitive to the overhead.

The effects of varying the cache size and link capacity on the normalized throughput are shown in Fig. 8. Comparing Figs. 8(a) and 8(b) with Figs. 4(a) and 4(b) we see that the normalized throughput does not grow as fast as

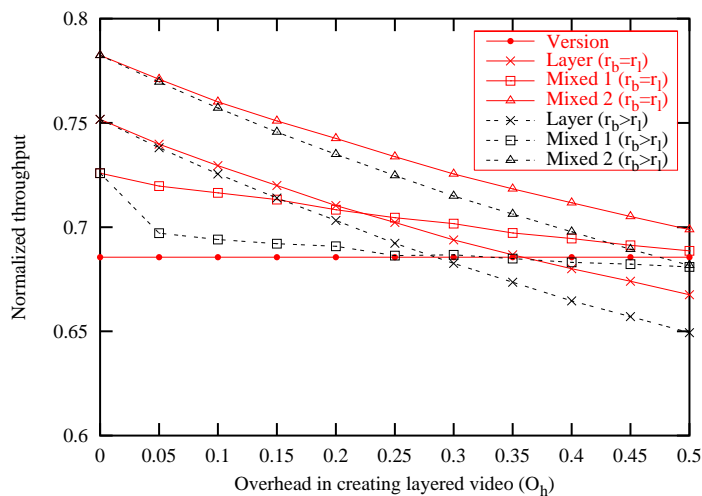


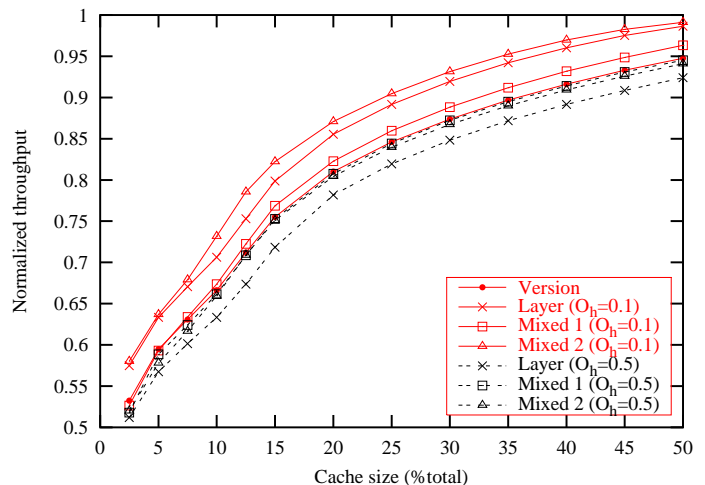
Fig. 7. Adaptive caching with varying amount of overhead for layered encoding

in the static caching model for small cache sizes and small link capacities. This is again due to the fact that without *a priori* knowledge of the request distribution the order of the request arrivals has a strong impact on the cache composition. Also, moderately popular objects tend to keep the few extremely popular objects from being cached. Note again that by weeding out one-timer requests, heuristic 2 achieves a higher throughput than the other strategies.

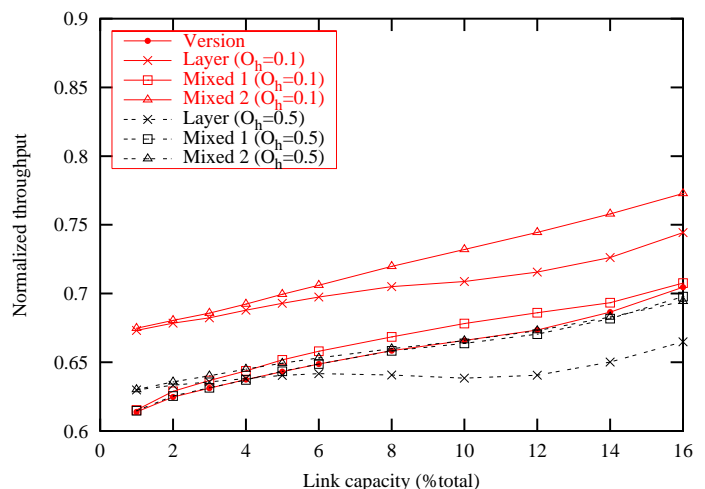
V. CONCLUSION

In this paper we have studied pure versions, pure layers, and mixed distribution strategies. We found that mixed distribution strikes a good balance to offer the best overall performance. Our study leads to the following guidelines for distributing multi-quality video in the Internet:

1. Caches and CDN servers should be partially pre-filled with the most popular videos. If there are requests for both quality levels of a popular video, than the server should cache both the base and the enhancement layer of the video (rather than use versions). It is important to pre-fill the cache with the popular videos; otherwise, continuously streaming moderately-popular videos may prevent popular videos from getting stored in the cache.
2. For a first-time request of a video with unknown popularity, the origin server should stream the requested quality level as a *version*, and the proxy should not cache the version. If the video experiences multiple requests, then layers should be streamed and stored in the cache.
3. Although we should use versions to stream first-time requests from origin server to client, we should not cache versions (unless all the requests for a specific video are for one quality level).



(a) varying cache size



(b) varying link capacity

Fig. 8. Adaptive caching with varying cache size and link capacity ($q = 0.4$).

The focus of this study was on two quality levels. We expect many of the conclusions to be even more pronounced for three or more quality levels.

REFERENCES

- [1] T. Abdelzaher and N. Bhatti, "Web server qos management by adaptive content delivery," in *Proc. of International Workshop on QoS*, May 1999.
- [2] W. Ma, I. Bedner, G. Chang, A. Kuchinsky, and H.J. Zhang, "A framework for adaptive content delivery in heterogeneous network environments," in *Proc. of MMCN 2000*, San Jose, CA, January 2000.
- [3] K. Chandra and A. Reibman, "Modeling one- and two-layer variable bit rate video," *IEEE/ACM Trans. on Networking*, vol. 7, no. 3, pp. 398–413, June 1999.

- [4] J. Kimura, F. Tobagi, J. Pulido, and P. Emstad, "Perceived quality and bandwidth characterization of layered mpeg-2 video encoding," in *SPIE International Symposium on Voice, Video and Data Communications*, Boston, MA, September 1999.
- [5] P. De Cuetos, D. Saporilla, and K. Ross, "Adaptive streaming of stored video in a tcp-friendly context: Multiple versions or multiple layers?," in *Proc. of International Packet Video Workshop*, Kyongju, Korea, April 2001.
- [6] T. Kim and M.H. Ammar, "A comparison of layering and stream replication video multicast schemes," in *Proc. of NOSSDAV 2001*, Port Jefferson, NY, Jun 2001.
- [7] J. Kangasharju, F. Hartanto, M. Reisslein, and K. Ross, "Distributing layered encoded video through caches," in *Proc. of IEEE INFOCOM*, Anchorage, AL, April 2001.
- [8] K.W. Ross, *Multiservice Loss Models for Broadband Telecommunication Networks*, Springer-Verlag, 1995.
- [9] Real Networks. www.realnetworks.com.