Models for Optimized Caching in Systems with Heterogeneous Client Populations

Derek L. Eager
Michael C. Ferris
Mary K. Vernon

Technical Report #1402

August 1999

Models for Optimized Caching in Systems with Heterogeneous Client Populations *

Derek L. Eager | Michael C. Ferris | Mary K. Vernon

University of Saskatchewan
Saskatoon, SK Canada S7N 5A9
eager@cs.usask.ca

University of Wisconsin – Madison
Madison, WI 53706
{ferris, vernon}@cs.wisc.edu

Computer Sciences Technical Report No. 1402, University of Wisconsin – Madison, August 1999.

Abstract

A recent paper develops and applies two new models for determining optimal proxy cache content in systems supporting on-demand access to large, widely shared data, such as popular video objects. The new models permit the study of specific types of heterogeneous systems in which the proxy servers have differing client populations and server capabilities. In one model, the clients of a given proxy server have either a higher or lower total request rate for the data files, and the server in that region may have different storage and bandwidth capabilities than the other proxy servers. In the other model, each group of proxy servers has a *preferred group* of files that the clients of those proxy servers request most frequently. This technical report provides details regarding the equations that the models employ.

1 Introduction

Supporting on-demand access to large widely shared data, such as popular video objects, requires effective use of (regional) proxy servers that store some of the data close to the clients. In [7], we presented two models for determining optimal proxy cache contents in heterogeneous systems in which the proxy servers have differing client populations and server capabilities. This technical report provides the full set of equations used in the models, including those omitted from [7] due to space constraints.

Section 2 provides an overview of the partitioned dynamic skyscraper delivery technique that the models assume, and summarizes an optimization model for homogeneous systems that was developed previously [6]. Section 3 presents the equations for a model used to study systems in which the proxy servers have differing client request rates and server capabilities, while Section 4 gives the equations for a model designed to study the impact of heterogeneous object selection frequencies. Experiments using the models, and the insights and design principles for caching in heterogeneous systems gained from the experiments, are described in [7].

*This research was partially supported by the Natural Sciences and Engineering Research Council of Canada under Grant OGP-0000264, by the Air Force Office of Scientific Research under Grant F49620-98-1-0417, and by the National Science Foundation under Grant CCR-9975044.

2 Background

2.1 Partitioned Dynamic Skyscraper

Partitioned dynamic skyscraper delivery [6] is one of a number of recently proposed segmented multicast delivery techniques. These techniques achieve bandwidth savings by dividing each hot object into fixed increasing-sized segments, and employing transmission schedules in which the smaller initial segments are multicast more frequently than the remaining larger segments. Most of the segmentation-based delivery techniques employ static transmission schedules and are thus only applicable to objects that are steadily extremely popular. An exception is the dynamic skyscraper technique [5], which dynamically assigns server bandwidth, in the form of multicast "transmission clusters", for segmented delivery of objects in response to client requests.

The dynamic skyscraper technique was recently extended [6] to support delivery in environments that have regional (or proxy) caching. This *partitioned skyscraper* technique partitions the segments of each object into two sets, one composed of the first k segments (the "leading segment set") and one composed of the remaining $K-k$ segments (the "trailing segment set"). Separate transmission clusters are dynamically scheduled for delivery of the leading segment sets and the trailing segment sets. Scheduled delivery of the two sets is coordinated so that, for each client, jitter between playback of each set is avoided. For each object, a proxy server may cache just the leading segment set, both sets of segments, or neither set. A full description of the partitioned dynamic skyscraper system is provided in [6].

2.2 An Optimization Model for Homogeneous Systems

In previous work [6], a model was developed for determining optimal regional/proxy cache content, for systems with (a) identical regional/proxy server bandwidth and storage capabilities, (b) statistically homogeneous regional client workloads, (c) partitioned dynamic skyscraper delivery, and (d) equal-sized objects. For each object, the proxy servers can store either none, all, or just the leading segment set of the object. The model considers client cost sharing for the multicast delivery, as well as the relative cost of remote and proxy server resources, in determining the optimal data to store at the proxy servers. This model is reviewed here, as it is the starting point for the more complex models of systems with heterogeneous proxy server capabilities and workloads that are described in the next section.

To determine delivery cost, the homogeneous model uses simple analytic estimates of the number of remote server channels (C_{remote}) and proxy server channels ($C_{regional}$) needed to support a given client workload as a function of the object segment sets that are stored at the proxy. Results given in [6] show that these server bandwidth estimates are very close to the knee of the curve of mean client waiting time versus the inverse of the number of server channels, and that the knee of the curve is typically quite sharp for systems that use multicast delivery. Furthermore, since the system can provide immediate service to client requests, the average client waiting time is typically very small near the (sharp) knee of the curve.

Table 1 defines the model input and output parameters. Note that the last six input parameters specify the particular configuration of the partitioned dynamic skyscraper delivery system. The key system constraints in the optimization are the maximum proxy server bandwidth ($N_{channels}$), and storage capacity ($N_{segments}$). The key model outputs are the θ_i values that specify whether object i should be fully or partially cached, or not cached, at the proxy servers.

Given that β is the ratio of costs for proxy server channels and remote server channels, and given that P is the number of proxy servers in the system, the optimization problem for our homogeneous

Table 1: **Parameters of the Homogeneous System Optimization Model**

Input	Parameter Definition
n	number of objects
$N_{channels}$	maximum number of channels at each proxy server
$N_{segments}$	maximum storage capacity (measured in number of unit-segments) at each proxy server
P	number of proxy servers
β	the cost ratio of a proxy server channel and a remote server channel
λ_i	total arrival rate of requests for object i (from all regions)
k	number of segments in the leading segment set
K	total number of segments in the leading and trailing segment sets
s_j	size of the j'th segment (relative to the size of the first)
T_1	duration of a transmission of segment s_1
w	the largest segment size in the leading segment set
W	the largest segment size in the trailing segment set
Output	**Parameter Definition**
C_{remote}	number of channels needed at the remote server
$C_{regional}$	number of channels needed at the proxy server
$D_{regional}$	storage needed at each proxy server (measured in units of s_1)
$X_{i,l,x}, X_{i,t,x}$ $x \in \{R, p, r\}$	maximum rate at which new transmission clusters can be allocated for multicasts of the leading/trailing (l/t) segment set for object i, distinguished by whether the object is stored only at the remote server ($x = R$), partially (leading segment set) at the regional server ($x = p$), or fully at the regional server ($x = r$)
θ_i^R	equals 1 if object i is stored only at the remote server; 0 otherwise
θ_i^p	equals 1 if only the leading segment set of object i is cached regionally; 0 otherwise
θ_i^r	equals 1 if object i is entirely cached regionally; 0 otherwise

system model is defined as follows:

$$\min_{\theta} \quad C_{remote}(\theta) + P\beta C_{regional}(\theta)$$

$$\begin{aligned} \text{subject to} \quad & C_{regional}(\theta) \le N_{channels} \\ & D_{regional}(\theta) \le N_{segments} \\ & \theta_i^R + \theta_i^p + \theta_i^r = 1, \quad i = 1, 2, \ldots, n \\ & \theta_i^R, \theta_i^p, \theta_i^r \in \{0, 1\}, \quad i = 1, 2, \ldots, n \end{aligned}$$

In the above we have used the symbol θ to represent the vector with components θ_i^R, θ_i^p, θ_i^r, $i = 1, 2, 3, ..., n$. Note that the expression to be minimized is the total delivery cost for all objects to clients in all regions. However, dividing this expression by P gives the cost for delivery to an individual region, whose clients collectively pay for $\frac{1}{P}$ of the remote delivery cost, since the regional client populations are statistically the same.

Making use of the fact that transmission clusters for leading and trailing segment sets are of duration wT_1 and WT_1, and include k and $K - k$ channels, respectively [6], the number of channels required for remote delivery is given by the following equation:

$$C_{remote}(\theta) = \sum_{i=1}^{n} \theta_i^R X_{i,l,R} kwT_1 + (\theta_i^R X_{i,t,R} + \theta_i^p X_{i,t,p})(K-k)WT_1$$

That is, the number of channels needed for remote server multicasts is the sum over all objects of the number of channels needed to deliver leading segment sets (i.e., if the object is stored only at the remote server, the maximum rate at which new transmission clusters can be allocated for the leading segment set ($X_{i,l,R}$) times the number of channels in the cluster (k) times the duration of the cluster), plus the number of channels needed to deliver trailing segment sets (i.e, if the object is delivered fully or partially by the remote server, the maximum rate at which new transmission clusters can be allocated for the trailing segment set ($X_{i,t}$) times the number of channels in the cluster ($K-k$) times the duration of the cluster).

Similarly, the number of channels required at each proxy server, and the storage required at each proxy server, are given by

$$C_{regional}(\theta) = \sum_{i=1}^{n} (\theta_i^r X_{i,l,r} + \theta_i^p X_{i,l,p}) kwT_1 + \theta_i^r X_{i,t,r}(K-k)WT_1$$

$$D_{regional}(\theta) = \sum_{i=1}^{n} \left((\theta_i^r + \theta_i^p) \sum_{j=1}^{k} s_j + \theta_i^r \sum_{j=k+1}^{K} s_j \right)$$

The maximum rate X_i at which new leading segment set or tailing segment set transmission clusters can be allocated for object i is computed as the inverse of the minimum average time between the initiations of the new clusters for object i. This, in turn, is determined by the arrival rate of requests at the server of interest, and by the size of the transmission cluster "catchup window", defined as the period of time from the beginning of a transmission cluster during which a newly arriving request can be served by that cluster. For example, since the size of the catchup window for a trailing segment set transmission cluster is $\left(W - s_{k+1} + \sum_{j=1}^{k} s_j\right) T_1$, and since the average time from the end of the catchup window until a new client request for object i is $\frac{1}{\lambda_i}$, the maximum allocation rate of new transmission clusters for a trailing segment set that is not cached (and thus must be served remotely), $X_{i,t,R}$ or $X_{i,t,p}$, is computed as:

$$X_{i,t,R} = X_{i,t,p} = \frac{1}{\left(W - s_{k+1} + \sum_{j=1}^{k} s_j\right) T_1 + \frac{1}{\lambda_i}}$$

Note that the above expression gives the allocation rate of transmission clusters if the server had unlimited bandwidth. Since the server has limited bandwidth, the above expression is the *maximum* allocation rate, as there may be queueing of client requests and batching of these requests while queued, or client balking. The rest of the cluster allocation rate equations in the homogeneous model are given in [6].

Both the objective function and the constraints in the homogeneous optimization model are *linear* functions of the binary variables θ. Thus, the minimum cost cache content (i.e., the θ_i values) may be computed through solution of a mixed integer program (MIP) [8]. For the purposes of experimentation, the problem was formulated using the GAMS modeling language [1] and solved using a combination of the XPRESS [4] and CPLEX [2] solvers. Both of these codes use a linear programming [3] based branch and bound solution strategy.

Table 2: **New Parameters for the Request Rate and Server Heterogeneity Model.**

Input	Parameter Definition
f_d	fraction of the total requests that are from clients belonging to the distinct region
Output	**Parameter Definition**
C^d_{remote}, C^{nd}_{remote}	the component of the remote server "cost" (as measured in numbers of channels) apportioned to the distinct proxy server, and to each other proxy server, respectively
$X^{y,z}_{i,l,x}$, $X^{y,z}_{i,t,x}$ $x \in \{R, d, nd\}$ $y, z \in \{R, p, r\}$	maximum rate at which transmission clusters can be allocated for multicasts of the leading/trailing (l/t) segment set for object i, distinguished by the server (R - remote, d - distinct proxy, or nd - other proxy), and by whether requests from a distinct (y) or non-distinct (z) region receive all of the object from the remote site (R), part (only the trailing segment set) from the remote site (p), or receive all of the object from the regional site (r)
$\theta^{y,z}_i$; $y, z \in \{R, p, r\}$	equals 1 if object i is cached at distinct and non-distinct proxy servers according to the superscripts y and z, respectively (R - the object is not cached regionally; p - only part (the leading segment set) is cached regionally; r - the object is fully cached at the respective regional server); equals 0 otherwise.

3 Heterogeneous Client Request Rates and Server Capabilities

In the first heterogeneous system model, proxy servers are assumed to be identical except for a single proxy server that has a "distinct" (higher or lower) client request rate, and possibly also different server bandwidth and storage capacity.

The new parameters that are required in this model are defined in Table 2. In addition, the input and output parameters for the proxy servers ($N_{channels}$, $N_{segments}$, $C_{regional}$, and $D_{regional}$) each have a superscript (d or nd) to denote the type of server (distinct or non-distinct) that the parameter applies to. As before, the key outputs of the model are the θ_i parameters that determine whether each object i is fully or partially cached at each type of regional proxy server.

Object allocations that minimize overall cost are determined by solving the following optimization problem:

$$\min_{\theta} \quad C_{remote}(\theta) + \beta(C^d_{regional}(\theta) + (P-1)C^{nd}_{regional}(\theta))$$

$$\begin{aligned} \text{subject to} \quad & C^d_{regional}(\theta) \le N^d_{channels} \\ & C^{nd}_{regional}(\theta) \le N^{nd}_{channels} \\ & D^d_{regional}(\theta) \le N^d_{segments} \\ & D^{nd}_{regional}(\theta) \le N^{nd}_{segments} \\ & \textstyle\sum_{y,z \in \{R,p,r\}} \theta^{y,z}_i = 1, \quad i = 1, 2, \ldots, n \\ & \theta^{y,z}_i \in \{0, 1\}, \quad y, z \in \{R, p, r\}, \quad i = 1, 2, \ldots, n \end{aligned}$$

Here we use the notation θ to represent the vector whose components are $\theta^{y,z}_i$, $y, z \in \{R, p, r\}$, $i = 1, 2, \ldots, n$.

In heterogeneous systems that have asymmetric client workloads or diverse server capacities or bandwidths, the "socially optimal" proxy server cache contents that minimize total cost of delivery may be quite different from the "individually optimal" cache content that would result if a particular (competitive) regional service provider attempted to minimize its own delivery cost. To assess whether these two types of solutions differ for a given system specified by the model inputs, we also derive individually optimal cache content for the distinct (non-distinct) server(s), under a given fixed set of cache contents, such as the socially optimal allocations, for the non-distinct (distinct) proxy servers. The following optimization problem minimizes the delivery cost of the distinct proxy server, with the superscript "O" in $\theta^{y,O}$ denoting the fixed allocations assumed for the non-distinct proxy servers:

$$\min_{\theta} \quad C^d_{remote}(\theta) + \beta C^d_{regional}(\theta)$$

$$\text{subject to} \quad \begin{aligned} C^d_{regional}(\theta) &\leq N^d_{channels} \\ D^d_{regional}(\theta) &\leq N^d_{segments} \\ \textstyle\sum_{y \in \{R,p,r\}} \theta_i^{y,O} &= 1, \quad i = 1, 2, \ldots, n \\ \theta_i^{y,O} \in \{0,1\}, \quad y &\in \{R,p,r\}, \quad i = 1, 2, \ldots, n \end{aligned}$$

Correspondingly, the optimization problem given below minimizes the delivery cost of the non-distinct proxy server, with the superscript "O" in $\theta^{O,z}$ denoting the fixed allocations assumed for the distinct proxy servers:

$$\min_{\theta} \quad C^{nd}_{remote}(\theta) + \beta C^{nd}_{regional}(\theta)$$

$$\text{subject to} \quad \begin{aligned} C^{nd}_{regional}(\theta) &\leq N^{nd}_{channels} \\ D^{nd}_{regional}(\theta) &\leq N^{nd}_{segments} \\ \textstyle\sum_{z \in \{R,p,r\}} \theta_i^{O,z} &= 1, \quad i = 1, 2, \ldots, n \\ \theta_i^{O,z} \in \{0,1\}, \quad z &\in \{R,p,r\}, \quad i = 1, 2, \ldots, n \end{aligned}$$

Similar to the homogeneous system model summarized in Section 2.2, we derive estimates of the number of channels needed to support a given load at each server, and then determine minimum cost proxy server cache content for each of the above optimization problems through solution of a mixed integer linear program.

Making use of the transmission cluster durations and sizes given in Section 2.2, the number of channels required at the remote server and at the two types of proxy servers, and the storage required at the proxy servers, are given by the following equations:

$$\begin{aligned} C_{remote}(\theta) \quad = \quad & \sum_{i=1}^{n} \Bigg(\Big(\sum_{y,z \in \{R,p\}} \theta_i^{y,z} X_{i,t,R}^{y,z} + \sum_{y \in \{R,p\}} \theta_i^{y,r} X_{i,t,R}^{y,r} + \sum_{z \in \{R,p\}} \theta_i^{r,z} X_{i,t,R}^{r,z} \Big)(K-k)WT_1 \\ & + \Big(\theta_i^{R,R} X_{i,l,R}^{R,R} + \sum_{y \in \{p,r\}} \theta_i^{y,R} X_{i,l,R}^{y,R} + \sum_{z \in \{p,r\}} \theta_i^{R,z} X_{i,l,R}^{R,z} \Big) kwT_1 \Bigg) \end{aligned}$$

$$C^d_{regional}(\theta) = \sum_{i=1}^{n} \left(\sum_{z \in \{R,p,r\}} \theta_i^{r,z} X_{i,t,d}^{r,z} (K-k)WT_1 + \sum_{y \in \{p,r\}} \sum_{z \in \{R,p,r\}} \theta_i^{y,z} X_{i,l,d}^{y,z} kwT_1 \right)$$

$$C^{nd}_{regional}(\theta) = \sum_{i=1}^{n}\left(\sum_{y\in\{R,p,r\}}\theta_i^{y,r}X_{i,t,nd}^{y,r}(K-k)WT_1 + \sum_{y\in\{R,p,r\}}\sum_{z\in\{p,r\}}\theta_i^{y,z}X_{i,l,nd}^{y,z}kwT_1\right)$$

$$D^{d}_{regional}(\theta) = \sum_{i=1}^{n}\sum_{z\in\{R,p,r\}}\left((\theta_i^{r,z}+\theta_i^{p,z})\sum_{j=1}^{k}s_j + \theta_i^{r,z}\sum_{j=k+1}^{K}s_j\right)$$

$$D^{nd}_{regional}(\theta) = \sum_{i=1}^{n}\sum_{y\in\{R,p,r\}}\left((\theta_i^{y,r}+\theta_i^{y,p})\sum_{j=1}^{k}s_j + \theta_i^{y,r}\sum_{j=k+1}^{K}s_j\right)$$

The remote server channel cost that is apportioned to (clients of) the distinct proxy server is the sum of (a) $1/P$ of the channels required to multicast segments that are not cached by any of the servers, and (b) all of the bandwidth required to multicast segments that are not stored at the distinct proxy server but are stored at the other proxy servers:

$$\begin{aligned} C^{d}_{remote}(\theta) &= \sum_{i=1}^{n}\left(\left[\frac{\sum_{y,z\in\{R,p\}}\theta_i^{y,z}X_{i,t,R}^{y,z}}{P} + \sum_{y\in\{R,p\}}\theta_i^{y,r}X_{i,t,R}^{y,r}\right](K-k)WT_1\right. \\ &\quad \left. + \left[\frac{\theta_i^{R,R}X_{i,l,R}^{R,R}}{P} + \sum_{z\in\{p,r\}}\theta_i^{R,z}X_{i,l,R}^{R,z}\right]kwT_1\right) \end{aligned}$$

Similarly, the remote server channel cost apportioned to (clients of) a non-distinct proxy server is given by:

$$\begin{aligned} C^{nd}_{remote}(\theta) &= \sum_{i=1}^{n}\left(\left[\frac{\sum_{y,z\in\{R,p\}}\theta_i^{y,z}X_{i,t,R}^{y,z}}{P} + \frac{\sum_{z\in\{R,p\}}\theta_i^{r,z}X_{i,t,R}^{r,z}}{P-1}\right](K-k)WT_1\right. \\ &\quad \left. + \left[\frac{\theta_i^{R,R}X_{i,l,R}^{R,R}}{P} + \frac{\sum_{y\in\{p,r\}}\theta_i^{y,R}X_{i,l,R}^{y,R}}{P-1}\right]kwT_1\right) \end{aligned}$$

The maximum rates X_i at which transmission clusters can be allocated are determined by the arrival rate of requests at the server of interest, and by the size of the transmission cluster catchup window, as in the homogeneous system model outlined in Section 2. Transmission clusters for leading and trailing segment sets have catchup windows of differing lengths. Further, the length of the catchup window for a leading segment set transmission cluster depends on when the cluster begins in relationship to the end of the catchup window of the corresponding trailing segment set transmission cluster. For each distinct case, the rate X_i may be computed as the inverse of the minimum time between the initiations of new clusters for object i, along similar lines as for the homogeneous system model. The set of equations given below for the X_i employ the superscript notations "*" ("wild-card"), "−" ("don't care"), and "|" ("or") with the following meanings: an equation in which a "*" superscript appears holds for any consistent substitution throughout the equation of the "*" with R, p, or r; an equation in which a "−" superscript appears holds for any (not necessarily consistent) substitution of the "$-$" with R, p, or r; and an equation in which a superscript "$x|y$" appears, where $x, y \in \{R, p, r\}$, holds if either superscript x or y is consistently substituted throughout the equation. The X_i equations for the model with heterogeneous client request rates and server capabilities are:

$$X_{i,t,R}^{R|p,R|p} = \frac{1}{\frac{1}{\lambda_i} + \left(W - s_{k+1} + \sum_{j=1}^{k} s_j\right) T_1}$$

$$X_{i,t,R}^{R|p,r} = \frac{1}{\frac{1}{f_d \lambda_i} + \left(W - s_{k+1} + \sum_{j=1}^{k} s_j\right) T_1}$$

$$X_{i,t,R}^{r,R|p} = \frac{1}{\frac{1}{(1-f_d)\lambda_i} + \left(W - s_{k+1} + \sum_{j=1}^{k} s_j\right) T_1}$$

$$X_{i,t,d}^{r,-} = \frac{1}{\frac{1}{f_d \lambda_i} + \left(W - s_{k+1} + \sum_{j=1}^{k} s_j\right) T_1}$$

$$X_{i,t,nd}^{-,r} = \frac{1}{\frac{P-1}{(1-f_d)\lambda_i} + \left(W - s_{k+1} + \sum_{j=1}^{k} s_j\right) T_1}$$

$$X_{i,l,d}^{r,-} = \frac{1}{\frac{1}{f_d \lambda_i} + \left(X_{i,t,d}^{r,-} w T_1 \frac{w-1}{2} + (1 - X_{i,t,d}^{r,-} w T_1)(w-1)\right) T_1}$$

$$X_{i,l,d}^{p,*} = \frac{1}{\frac{1}{f_d \lambda_i} + \left(X_{i,t,R}^{p,*} w T_1 \frac{w-1}{2} + (1 - X_{i,t,R}^{p,*} w T_1)(w-1)\right) T_1}$$

$$X_{i,l,nd}^{-,r} = \frac{1}{\frac{P-1}{(1-f_d)\lambda_i} + \left(X_{i,t,nd}^{-,r} w T_1 \frac{w-1}{2} + (1 - X_{i,t,nd}^{-,r} w T_1)(w-1)\right) T_1}$$

$$X_{i,l,nd}^{*,p} = \frac{1}{\frac{P-1}{(1-f_d)\lambda_i} + \left(X_{i,t,R}^{*,p} w T_1 \frac{w-1}{2} + (1 - X_{i,t,R}^{*,p} w T_1)(w-1)\right) T_1}$$

$$X_{i,l,R}^{R,R} = \frac{1}{\frac{1}{\lambda_i} + \left(X_{i,t,R}^{R,R} w T_1 \frac{w-1}{2} + (1 - X_{i,t,R}^{R,R} w T_1)(w-1)\right) T_1}$$

$$X_{i,l,R}^{R,p|r} = \frac{1}{\frac{1}{f_d \lambda_i} + \left(X_{i,t,R}^{R,p|r} w T_1 \frac{w-1}{2} + (1 - X_{i,t,R}^{R,p|r} w T_1)(w-1)\right) T_1}$$

$$X_{i,l,R}^{p|r,R} = \frac{1}{\frac{1}{(1-f_d)\lambda_i} + \left(X_{i,t,R}^{p|r,R} wT_1 \frac{w-1}{2} + (1 - X_{i,t,R}^{p|r,R} wT_1)(w-1)\right) T_1}$$

4 Heterogeneous Object Selection Frequencies

In the second heterogeneous system model, the objects and the proxy servers are each partitioned into G equal-sized groups, and each group of proxy servers has a *preference* (i.e., a larger fraction of the regions' client requests) for a distinct group of objects. Each proxy server has the same request rate for each of its (G-1) non-preferred groups of objects. Also, the *relative* selection frequencies of the objects *within a group* are the same for all groups and for all proxy servers.

The new model parameters for the system with the above heterogeneous object selection frequencies are given in Table 3. Note that although each proxy server may optimally cache different object segments, due to the symmetry in the regional client workloads, each proxy server will cache the same segments from its respective preferred and non-preferred groups. This greatly simplifies the model.

Due to the symmetry in the regional client workloads, the socially optimal cache content for each regional/proxy server is also the individually optimal content, which can be determined by solving the following optimization problem:

$$\min_{\theta} \quad C_{remote}(\theta) + P\beta C_{regional}(\theta)$$

$$\begin{array}{ll} \text{subject to} & C_{regional}(\theta) \leq N_{channels} \\ & D_{regional}(\theta) \leq N_{segments} \\ & \sum_{y,z\in\{R,p,r\}} \theta_i^{y,z} = 1, \quad i = 1, 2, \ldots, n \\ & \theta_i^{y,z} \in \{0,1\}, \quad y,z \in \{R,p,r\}, \quad i = 1,2,\ldots,n \end{array}$$

As in the previous models, we formulate estimates of the number of channels needed to support a given load at each server, and then determine optimal object allocations through solution of a mixed integer linear program. The number of channels required at the remote server and at each proxy server, and the storage required at each proxy server, are given by the equations below, where the index i runs only over the objects within a single group since the groups of objects are symmetric:

$$\begin{aligned} C_{remote}(\theta) = \; & G\sum_i \Bigg(\Big(\sum_{y,z\in\{R,p\}} \theta_i^{y,z} X_{i,t,R}^{y,z} + \sum_{y\in\{R,p\}} \theta_i^{y,r} X_{i,t,R}^{y,r} + \sum_{z\in\{R,p\}} \theta_i^{r,z} X_{i,t,R}^{r,z}\Big)(K-k)WT_1 \\ & + \Big(\theta_i^{R,R} X_{i,l,R}^{R,R} + \sum_{y\in\{p,r\}} \theta_i^{y,R} X_{i,l,R}^{y,R} + \sum_{z\in\{p,r\}} \theta_i^{R,z} X_{i,l,R}^{R,z}\Big) kwT_1 \Bigg) \end{aligned}$$

$$\begin{aligned} C_{regional}(\theta) = \; & \sum_i \Bigg(\left[\sum_{z\in\{R,p,r\}} \theta_i^{r,z} X_{i,t,a}^{r,z} + (G-1) \sum_{y\in\{R,p,r\}} \theta_i^{y,r} X_{i,t,na}^{y,r} \right] (K-k)WT_1 \\ & + \left[\sum_{y\in\{p,r\}} \sum_{z\in\{R,p,r\}} \theta_i^{y,z} X_{i,l,a}^{y,z} + (G-1) \sum_{y\in\{R,p,r\}} \sum_{z\in\{p,r\}} \theta_i^{y,z} X_{i,l,na}^{y,z} \right] kwT_1 \Bigg) \end{aligned}$$

Table 3: **New Parameters for the Object Selection Frequency Heterogeneity Model**

Input	Parameter Definition
f_a	fraction of the total requests from a region that are for objects in the group that it has an affinity for
G	number of groups of objects
Output	**Parameter Definition**
$X_{i,l,x}^{y,z}, X_{i,t,x}^{y,z}$ $x \in \{R, a, na\}$ $y, z \in \{R, p, r\}$	maximum rate at which transmission clusters can be allocated for multicasts of the leading/trailing (l/t) segment set for object i distinguished by the server ($x = R$ - remote server, $x = a$ - a proxy server that has an affinity (or preference) for the object's group, or $x = na$ - a proxy server that does not have a preference for the object's group), and by whether requests from a preferring (y) or non-preferring (z) region receive all of the object from the remote site (R) the trailing segment set from the remote site and the leading segment set from the regional site (p) or receive all of the object from the regional site (r)
$\theta_i^{y,z}$; $y, z \in \{R, p, r\}$	equals 1 if object i is cached at each proxy server for which the object's group is the region's preferred group, and at each proxy server for which the object's group is not the preferred group, according to the superscripts y and z, respectively ($y/z = R$ - the object is not cached regionally; $y/z = p$ - only the leading segment set is cached at a preferring/non-preferring proxy server; $y/z = r$ - the object is fully cached at a preferring/non-preferring proxy server); equals 0 otherwise.

$$
\begin{aligned}
D_{regional}(\theta) \quad = \quad & \sum_i \sum_{z \in \{R,p,r\}} \left((\theta_i^{r,z} + \theta_i^{p,z}) \sum_{j=1}^{k} s_j + \theta_i^{r,z} \sum_{j=k+1}^{K} s_j \right) \\
& + (G-1) \sum_i \sum_{y \in \{R,p,r\}} \left((\theta_i^{y,r} + \theta_i^{y,p}) \sum_{j=1}^{k} s_j + \theta_i^{y,r} \sum_{j=k+1}^{K} s_j \right)
\end{aligned}
$$

The maximum rates X_i at which transmission clusters can be allocated is computed in a similar fashion as in the previous models. Using the "wildcard", "don't care", and "or" notation used previously for the request rate and server heterogeneity model, we have:

$$
X_{i,t,R}^{R|p,R|p} = \frac{1}{\frac{1}{\lambda_i} + \left(W - s_{k+1} + \sum_{j=1}^{k} s_j\right) T_1}
$$

$$
X_{i,t,R}^{R|p,r} = \frac{1}{\frac{1}{f_a \lambda_i} + \left(W - s_{k+1} + \sum_{j=1}^{k} s_j\right) T_1}
$$

$$X^{r,R|p}_{i,t,R} = \frac{1}{\frac{1}{(1-f_a)\lambda_i} + \left(W - s_{k+1} + \sum_{j=1}^{k} s_j\right) T_1}$$

$$X^{r,-}_{i,t,a} = \frac{1}{\frac{P}{Gf_a\lambda_i} + \left(W - s_{k+1} + \sum_{j=1}^{k} s_j\right) T_1}$$

$$X^{-,r}_{i,t,na} = \frac{1}{\frac{G-1}{G}\frac{P}{(1-f_a)\lambda_i} + \left(W - s_{k+1} + \sum_{j=1}^{k} s_j\right) T_1}$$

$$X^{r,-}_{i,l,a} = \frac{1}{\frac{P}{Gf_a\lambda_i} + \left(X^{r,-}_{i,t,a} w T_1 \frac{w-1}{2} + (1 - X^{r,-}_{i,t,a} w T_1)(w-1)\right) T_1}$$

$$X^{p,*}_{i,l,a} = \frac{1}{\frac{P}{Gf_a\lambda_i} + \left(X^{p,*}_{i,t,R} w T_1 \frac{w-1}{2} + (1 - X^{p,*}_{i,t,R} w T_1)(w-1)\right) T_1}$$

$$X^{-,r}_{i,l,na} = \frac{1}{\frac{G-1}{G}\frac{P}{(1-f_a)\lambda_i} + \left(X^{-,r}_{i,t,na} w T_1 \frac{w-1}{2} + (1 - X^{-,r}_{i,t,na} w T_1)(w-1)\right) T_1}$$

$$X^{*,p}_{i,l,na} = \frac{1}{\frac{G-1}{G}\frac{P}{(1-f_a)\lambda_i} + \left(X^{*,p}_{i,t,R} w T_1 \frac{w-1}{2} + (1 - X^{*,p}_{i,t,R} w T_1)(w-1)\right) T_1}$$

$$X^{R,R}_{i,l,R} = \frac{1}{\frac{1}{\lambda_i} + \left(X^{R,R}_{i,t,R} w T_1 \frac{w-1}{2} + (1 - X^{R,R}_{i,t,R} w T_1)(w-1)\right) T_1}$$

$$X^{R,p|r}_{i,l,R} = \frac{1}{\frac{1}{f_a\lambda_i} + \left(X^{R,p|r}_{i,t,R} w T_1 \frac{w-1}{2} + (1 - X^{R,p|r}_{i,t,R} w T_1)(w-1)\right) T_1}$$

$$X^{p|r,R}_{i,l,R} = \frac{1}{\frac{1}{(1-f_a)\lambda_i} + \left(X^{p|r,R}_{i,t,R} w T_1 \frac{w-1}{2} + (1 - X^{p|r,R}_{i,t,R} w T_1)(w-1)\right) T_1}$$

5 Conclusions

The optimization models for heterogeneous client populations and proxy server capabilities developed in this technical report have been applied to obtain insights into caching strategies in such systems in [7].

References

[1] A. Brooke, D. Kendrick, and A. Meeraus. *GAMS: A User's Guide.* The Scientific Press, South San Francisco, CA, 1988.

[2] CPLEX Version 6.0. ILOG CPLEX Division, Incline Village, Nevada (http://www.cplex.com/).

[3] G. B. Dantzig. *Linear Programming and Extensions.* Princeton University Press, Princeton, New Jersey, 1963.

[4] *XPRESS-MP User Guide.* Dash Associates, Blisworth House, Blisworth, Northants, UK, 1997. (http://www.dashopt.com/).

[5] D. L. Eager and M. K. Vernon, "Dynamic Skyscraper Broadcasts for Video-on-Demand", *Proceedings of the 4th International Workshop on Multimedia Information Systems (MIS '98)*, Istanbul, Turkey, September 1998, pp. 18-32.

[6] D. L. Eager, M. C. Ferris, and M. K. Vernon, "Optimized Regional Caching for On-Demand Data Delivery", *Proceedings of the IS&T/SPIE Conference on Multimedia Computing and Networking 1999 (MMCN '99)*, San Jose, CA, January 1999, pp. 301-316.

[7] D. L. Eager, M. C. Ferris, and M. K. Vernon, "Optimized Regional Caching in Heterogeneous Video-on-Demand Systems", submitted for publication.

[8] G. L. Nemhauser and L. A. Wolsey. *Integer and Combinatorial Optimization.* Wiley, New York, NY, 1988.