An Evaluation of Alternative Continuous Media Replication Techniques in Wireless Peer-to-Peer Networks

Shahram Ghandeharizadeh and Tooraj Helmi
Department of Computer Science
University of Southern California
Los Angeles, CA 90089, USA

ABSTRACT

This study investigates a novel streaming architecture consisting of home-to-home online (H2O) devices that collaborate to provide on-demand access to a large selection of audio and video clips. An H2O device consists of a high bandwidth wireless communication component, a powerful processor, and gigabytes of storage. This study investigates three families of replication strategies for a H2O cloud. We evaluate these using analytical models. The obtained results demonstrate the superiority of one strategy that determines the number of replicas for a clip i based on (a) the bandwidth required to display clip i proportional to the bandwidth required by the other clips in the database, and (b) the square root of the frequency of access to the clips.

Categories and Subject Descriptors

H.3.2 [**Information Storage And Retrieval**]: Information Storage

General Terms

Algorithm, Design, Performance

Keywords

Replication, Continuous Media, Peer-to-Peer, Wireless

1. INTRODUCTION

Moore's law has made it feasible to consider peer-to-peer network of devices that provide access to a large repository of information. For example, Intel offers a small device that consists of a 500 MHz processor and a wireless component offering transmission rates in the order of tens of Megabits per second, Mbps. This device costs approximately $85 and can be extended with a mass storage device. One application of these devices is to stream continuous media, audio and video clips, for home entertainment systems. With this application, a collection of home-to-home online (H2O) devices collaborate with one another to provide their users to display a clip from a wide variety of titles, see Figure 1. A household may store its personal video library on a H2O cloud,

MobiDE'03, September 19, 2003, San Diego, California, USA.

Figure 1: Home-to-home on-line devices streaming continuous media.

enabling the user to retrieve their content anywhere, e.g., a friend's house. The system might encrypt the content to either protect it from un-authorized access or implement a business model to generate revenues.

We envision the H2O framework to complement the existing wired solutions based on xDSL, cable, cellular base stations, wireless access points, etc. At the same time, we assume each H2O to consist of tens of gigabytes of storage for temporary storage of one or more clips. These devices collaborate by providing their local data to service a display initiated by a different H2O device. A H2O device might participate in three possible roles: (i) as a displaying client that consumes data from one (potentially several other) H2O device, (ii) as a producer of data for consumption by a displaying H2O device, and (iii) as a router of data in order to facilitate delivery of data from a producer to a consumer. Given a fixed number of clips, a key challenge is how to determine the number of replicas for each clip in order to enable all H2O devices to support continuous display of a clip. This is challenging for several reasons. First, when compared with traditional data types such as text and still images, continuous media consist of a sequence of quanta, either audio samples or video frames, that convey meaning when presented at a pre-specified rate [6, 9]. Once the display is initiated, if the data is delivered below this rate then a display might suffer from frequent disruptions and delays, termed hiccups. Second, the available bandwidth between two H2O devices depends on the number of transmitting devices operating in the same radio

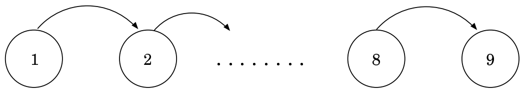

2.a String

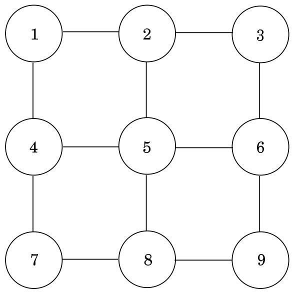

2.b Grid

Figure 2: Two different H2O topologies.

range. Third, the topology of the wireless connectivity between H2O devices impacts how many H2O devices a clip must visit in order to reach its destination. To illustrate, consider the nine H2O devices configured in a string topology, see Figure 2.a. This would resemble homes along a lake front. If $H2O_1$ requests a clip with only one copy residing on $H2O_9$ then the delivery of this clip to $H2O_1$ would consume some bandwidth of each link and visit all seven intermediate nodes. On the other hand, if H2O devices are organized in a grid structure (homes in a city block, see Figure 2.b) then only the bandwidth of four links would be employed to deliver a clip from $H2O_9$ to $H2O_1$.

The primary contributions of this paper is to introduce three different replication strategies based on either size, bandwidth, or display time of the clips that constitute the repository. These strategies are designed for a centralized base station, e.g., a cellular base station (see Figure 1), that publishes audio and video clips across a fixed number of H2O devices. We evaluate these strategies using an analytical model that makes the following simplifying assumptions. First, we assume nodes are organized in a string topology. Second, the available bandwidth between two nodes is fixed. Our primary observation is that replication using bandwidth is superior to the other two alternatives. The performance of this strategy is improved when (a) the frequency of access to the clips is known, and (b) clips are replicated based on the square-root of this frequency.

Many prior studies have analyzed the role of proxy servers and partial caching of continuous media in the context of Internet [15, 13, 1, 12, 5, 11, 14]. Other studies investigated replication techniques in a multi-disk server [10, 16, 7, 2]. These studies assume delivery of data from a producing component (either a proxy server or a disk) does **not** reduce the bandwidth of another data producing component. With a H2O cloud, delivery of a clip (say from $H2O_9$ to $H2O_1$) reduces the bandwidth available to the intermediate nodes, impacting the number of simultaneous displays supported by the system. A recent study [3] investigates replication of data items in an unstructured peer-to-peer system for a non-streaming framework. It employs analytical models to observe that a technique based on the square root of the popularity of a data item yields the best Mean Search Size (MSS), defined as the number of walks necessary to locate the referenced data item. In this study, we do not consider discovery of data item. Instead, we investigate delivery of a discovered data item once it is referenced with a string topology of H2O devices with a finite link bandwidth.

In [8], we describe a replication strategy for a H2O cloud. This strategy is based on the following intuition: the first few blocks of a clip are required more urgently than its last few blocks and should be replicated more frequently in order to minimize startup latency. This study quantified the storage requirement of this approach with different topologies (string, grid, and graph) and showed significant savings in storage when compared with replicating the entire repository on each node. There are several differences between [8] and the work presented here. First, in this study, we investigate clip (instead of block) replication where the clip is replicated across a subset of nodes. Second, we focus on the network bandwidth and how to maximize the number of simultaneous displays supported by a H2O cloud. A comparison of replicating an object at a clip granularity versus block granularity is a future research direction.

The rest of this paper is organized as follows. In Section 2, we describe three alternative replication strategies. Section 3 presents analytical models to estimate the amount of bandwidth consumed to display a clip i with r_i replicas in a string topology consisting of $\mathcal{N}$ nodes. Section 4 presents results from these analytical models to demonstrate the superiority of one of the replication strategies. Brief conclusions and future research directions are contained in Section 6.

2. REPLICATION STRATEGIES

Assume a string topology consisting of $\mathcal{N}$ H2O devices, each with storage capacity c_i bytes. The total storage capacity of the system is C=$\sum_{i=1}^{\mathcal{N}} c_i$. There are m clips in the database, each with a constant bandwidth requirement β_i and size S_i. These two clip parameters specify the display time of a clip, D_i=$\frac{S_i}{\beta_i}$. The frequency of access to clip i is denoted as f_i with $\sum_{i=1}^{m} f_i = 1$. The number of replicas for clip i is denoted as r_i where $r_i \leq \mathcal{N}$. (r_i includes the original copy of clip i.)

A replication strategy determines a r_i value for each clip i with the objective to ensure $\mathcal{N}$ H2O devices can display a clip simultaneously. (A placement strategy assigns a replica of a clip i to a specific node. An investigation of these strategies is beyond the focus of this paper, see Section 6.) We assume the total size of the database exceeds the storage capacity of a H2O device, $\sum_{j=1}^{m} S_j > c_i$ for $1 \leq i \leq \mathcal{N}$. Otherwise, the problem is trivial and the database should be replicated on each device in its entirety. Similarly, we assume there is at least one copy of a clip in the system and that the database size is smaller than the total storage capacity of the system, $\sum_{j=1}^{m} S_j \leq C$. Otherwise, clips cannot be replicated due to insufficient storage capacity. In sum, a replication strategy must construct at least one copy of a clip in the system and no more than $\mathcal{N}$ replicas, $1 \leq r_i \leq \mathcal{N}$.

This section describes replication strategies based on size (S_i), display time (D_i) and display bandwidth requirement (β_i) of a clip. Each might determine the value of r_i based on either a uniform, proportional to the frequency of access (f_i), or proportional to the square root of the frequency of access ($\sqrt{f_i}$). The total expected number of replicas supported by the system is R= $\frac{C}{\sum_{i=1}^{m} f_i S_i}$.

The general framework that embodies each replication strategy consists of two steps. First, it employs a replication strategy to compute r_i for each clip i that constitutes the repository. Second, it compares the storage requirement of this replication strategy with C, total storage capacity of the system. If they are equal then then the algorithm terminates. If a replication strategy's storage requirement is less than C then it computes the remaining idle storage C_{idle}. Next, it removes those clips with $\mathcal{N}$ replicas (changes the value of m and S) and re-applies the replication strategy by: 1) using C_{idle} instead of C, and 2) re-computing a new value for R.

Table 1: Terms and their definitions

Replication Strategies	
SZ	Proportional to clip size.
DT	Proportional to display time.
BW	Proportional to bandwidth required to display a clip.
SZ$Freq$	Proportional to clip size and frequency of access.
DT$Freq$	Proportional to display time and frequency of access.
BW$Freq$	Proportional to bandwidth and frequency of access.
SZ$\sqrt{Freq}$	Proportional to size and square root of frequency.
DT$\sqrt{Freq}$	Proportional to display time and square root of frequency.
BW$\sqrt{Freq}$	Proportional to bandwidth and square root of frequency.
Database Parameters	
m	Number of clips.
S_i	Size of clip i in bytes.
S	Total Size of clips in bytes, $S = \sum_{i=1}^{m} S_i$.
D_i	Display time of clip i.
β_i	Bandwidth requirement of clip i.
f_i	frequency of access to clip i.
r_i	number of replicas for clip i.
System Parameters	
$\mathcal{N}$	Number of nodes in the string topology.
c_i	Storage capacity of node i in bytes.
C	Total storage capacity of the system, $C = \sum_{i=1}^{N} c_i$.
L	Bandwidth of each link in the string topology.
R	Total expected number of replicas in the system, $\frac{\sum_{j=1}^{\mathcal{N}} c_j}{\sum_{i=1}^{m} f_i S_i}$.
Analytical Model	
N_T	Total number of unique paths between two H2O nodes.
δ	The distance between H2O$_j$ and the closest replica of clip i.
χ_j^δ	Number of unique paths between H2O$_j$ and the closest replica of clip i with length $= \delta$.
α_j^δ	Number of unique paths from H2O$_j$ to the closest replica of clip i with length $\geq \delta$.
p_j^δ	Probability of not finding a replica of clip i in fewer than δ hops
$(q_j^\delta \| p_j^\delta)$	Probability of finding a replica of clip i exactly δ hops
P_j^δ	Probability of H2O$_j$ finding a replica of clip i in δ hops and not finding another in fewer hops $P_j^\delta = p_j^\delta \times (q_j^\delta \| p_j^\delta)$.

The algorithm terminates when either it constructs $\mathcal{N}$ replicas for all clips or the storage capacity of the system is exhausted.

When a replication strategy's final storage requirement exceeds C, the algorithm decrements the value of r_i starting with the clip that has the maximum r_i value. It breaks ties (several clips with identical r_i values) by choosing the clip with the largest size. If two equi-sized clips have the same r_i value then it chooses one randomly.

In the following sections, we describe each replication strategy and its variation.

2.1 Replication using clip size (SZ, SZ$_{Freq}$, SZ$_{\sqrt{Freq}}$)

The simplest form of this strategy, termed SZ, replicates clips uniformly based on their size and the total number of replicas supported by the system: r_i=$R \times \frac{S_i}{S}$. If r_i is less than 1 then it is reset to one. If it exceeds $\mathcal{N}$ then it is reset to $\mathcal{N}$.

One may extend this simple strategy to replicate clips proportional to their frequency of access, r_i=$f_i \times R \times \frac{S_i}{S}$. This technique is termed SZ$Freq$. By replicating a frequently accessed clip more often, this strategy strives to maximize the number of simultaneous displays.

Another variation, termed SZ$\sqrt{Freq}$, replicates a clip i proportional to the square root of its frequency of access (f_i), r_i=$\sqrt{f_i} \times R \times \frac{S_i}{S}$. The basic intuition behind SZ$\sqrt{Freq}$ is that our objective is to minimize $\sum_{i=1}^{m} \frac{f_i}{r_i/R}$. This optimization problem is discussed in different context and a square root allocation has been shown to solve this problem [3].

2.2 Replication using clip display time (DT, DT$_{Freq}$, DT$_{\sqrt{Freq}}$)

This family of techniques utilizes display time of clips to control the number of replicas. It is based on the intuition that a clip with a longer display time should be replicated more frequently. DT assumes a uniform replication strategy where r_i=$R \times \frac{D_i}{\sum_{i=1}^{m} D_i}$. DT$Freq$ considers the frequency of access to a clip when computing the number of replicas: r_i=$f_i \times R \times \frac{D_i}{\sum_{i=1}^{m} D_i}$. DT$\sqrt{Freq}$ employs the square root of frequency of access when computing the number of replicas: r_i=$\sqrt{f_i} \times R \times \frac{D_i}{\sum_{i=1}^{m} D_i}$.

2.3 Replication using clip display bandwidth requirement (BW, BW$_{Freq}$, BW$_{\sqrt{Freq}}$)

One may replicate based on the bandwidth consumption of clips in order to replicate those clips with a higher bandwidth requirement more often. This technique and its variations are identical to those of Sections 2.1 and 2.2 where β_i is employed instead of S_i and D_i, respectively.

3. ANALYTICAL MODELS

In this section, we describe analytical models to estimate the bandwidth required to display a clip i given its number of replicas, r_i. These models are average cases based on a randomly chosen H2O device to display clip i and a random placement of its r_i replicas across the nodes. Our analytical models are based on a simple counting methodology, consisting of the following logical steps. First, we enumerate the $\mathcal{N}$ possible nodes that a display might be activated on. We assume that H2O devices are numbered from 1 to $\mathcal{N}$. H2O$_j$ is the device numbered j, $1 \leq j \leq \mathcal{N}$. For each H2O$_j$, we enumerate all possible paths to the **closest** replica of clip i and the probability of this path. Next, we compute the average length for H2O$_j$ to the closest replica of clip i. By multiplying this average with the bandwidth required to display a clip, we obtain the total bandwidth required to display clip i.

Given a displaying H2O, say H2O$_d$, and a data producing H2O device that contains the referenced clip, say H2O$_p$, there exists $\mathcal{N}$ unique possible paths between H2O$_d$ and H2O$_p$. To demonstrate, observe that there exists:

- 1 unique path of length zero when H2O$_d$ finds a clip local to itself, i.e., d=p.
- $2(p-1)$ unique paths where the length of each path varies from one to $p-1$ iff $p \geq 2$. In other words, when $p \geq 2$, there exists two unique paths of length 1, two unique paths of length 2, ..., two unique paths of length $p-1$.
- $\mathcal{N}-2p+1$ unique paths where the length of each path varies from p to $\mathcal{N}-p$ when $\mathcal{N}-p \geq p$. In other words, when $\mathcal{N}-p \geq p$, there exists one unique path of length p, one unique path of length $p+1$, ..., one unique path of length $\mathcal{N}-p$.

The total of these three items is $\mathcal{N}$, i.e., $1+2(p-1)+\mathcal{N}-2p+1 = \mathcal{N}$.

With a r_i value, in order to enumerate all possible H2O$_j$ and δ combinations, we first enumerate all possible number of hops δ between a client and the closest replica of clip i. Observe that the value of δ ranges between zero and $\mathcal{N}-j$ with the string topology. Let χ_j^δ denote the total number of unique paths to the closest replica of clip i that are exactly δ hops in length from H2O$_j$. With H2O$_j$ and a specific value for δ (say $\delta = 5$), there exists either zero, one or two unique paths that are 5 hops long, i.e., $0 \leq \chi_j^5 \leq 2$. With the string topology, the value of χ_j^δ is trivial to derive because there are only two possible scenarios: $\delta \leq \frac{\mathcal{N}}{2}$ and $\delta \geq \frac{\mathcal{N}}{2}$. When $\delta \leq \frac{\mathcal{N}}{2}$, χ_j^δ equals one when either j is less than δ or j is greater than $\mathcal{N}-\delta$. Otherwise, χ_j^δ is two. When $\delta \geq \frac{\mathcal{N}}{2}$, χ_j^δ equals one when j is either less than $\mathcal{N}-\delta$ or greater than δ. Otherwise, χ_j^δ is zero.

We define α_j^δ to denote the number of unique paths between H2O$_j$ and the closest replica of i that are δ or more hops long. To derive the value of α_j^δ, we consider each H2O$_j$ client and its enumerated χ_j^δ value. Thus, $\alpha_j^\delta = \sum_{k=\delta}^{\mathcal{N}} \chi_j^k$. In our analysis, we take advantage of the symmetry in the string topology by breaking it into two parts where the left and the right hand side have an identical bandwidth requirement, e.g., H2O$_1$'s bandwidth requirement is identical to that of H2O$_\mathcal{N}$, H2O$_2$'s bandwidth requirement equals that of H2O$_{\mathcal{N}-1}$, etc. Hence, we consider the scenario where j ranges in value from 1 to $\lceil \frac{\mathcal{N}}{2} \rceil$, resulting in four possible cases:

1. if $\delta \geq \lfloor \frac{\mathcal{N}}{2} \rfloor$ and $\delta <= \mathcal{N}-j$ then $\alpha_j^\delta = \mathcal{N}-j-\delta+1$.
2. if $\delta \geq \lfloor \frac{\mathcal{N}}{2} \rfloor$ and $\delta > \mathcal{N}-j$ then $\alpha_j^\delta = 0$.
3. if $\delta < \lfloor \frac{\mathcal{N}}{2} \rfloor$ and $\delta \leq j-1$ then $\alpha_j^\delta = \mathcal{N}-2\delta+1$.
4. if $\delta < \lfloor \frac{\mathcal{N}}{2} \rfloor$ and $\delta > j-1$ then $\alpha_j^\delta = \mathcal{N}-j-\delta+1$.

Now, we derive the probability of each α_j^δ. This is based on deriving two probabilities. First, p_j^δ denotes the probability of not finding a replica of clip i in fewer than δ hops. The second probability, $(q_j^\delta | p_j^\delta)$, is conditioned using p_j^δ and defines the probability of locating a replica of clip i exactly δ hops away. The value of p_j^δ is dependent on the number of replicas and α_j^δ:

$$p_j^\delta = \begin{cases} 0, \; if \; \alpha_j^\delta < r_i \\ \frac{\alpha_j^\delta}{N_T} \times \frac{\alpha_j^\delta - 1}{N_T - 1} \times .... \times \frac{\alpha_j^\delta - (r_i - 1)}{N_T - (r_i - 1)} = \frac{\alpha_j^\delta!(N_T - r_i)}{(\alpha_j^\delta - r_i)!N_T!} else \end{cases} \tag{1}$$

The value $(q_j^\delta | p_j^\delta)$ is:

$$(q_j^\delta | p_j^\delta) = \begin{cases} 1, \; if \; \alpha_j^\delta \geq r_i > \alpha_j^\delta - \chi_j^\delta \\ 1 - (\frac{\alpha_j^\delta - \chi_j^\delta}{\alpha_j^\delta} \times \frac{\alpha_j^\delta - \chi_j^\delta - 1}{\alpha_j^\delta - 1} \times ... \times \frac{\alpha_j^\delta - \chi_j^\delta - (r_i - 1)}{\alpha_j^\delta - (r_i - 1)}), else \end{cases} \tag{2}$$

A more compact representation is:

$$(q_j^\delta | p_j^\delta) = \begin{cases} 1, \; if \; \alpha_j^\delta \geq r_i > \alpha_j^\delta - \chi_j^\delta \\ 1 - \frac{(\alpha_j^\delta - \chi_j^\delta)!(\alpha_j^\delta - r_i)!}{(\alpha_j^\delta - \chi_j^\delta - r_i)!\alpha_j^\delta!}, else \end{cases} \tag{3}$$

Next, we compute the probability of H2O$_j$ finding clip i in δ hops **and** not finding another replica in fewer than δ hops, denoted P_j^δ. $P_j^\delta = p_j^\delta \times (q_j^\delta | p_j^\delta)$ and is defined as:

$$P_j^\delta = \begin{cases} \frac{(N_T - r_i)!\alpha_j^\delta!}{(\alpha_j^\delta - r_i)!N_T!} - \frac{(N_T - r_i)!(\alpha_j^\delta - \chi_j^\delta)!}{(\alpha_j^\delta - \chi_j^\delta - r_i)!N_T!}, \; if \; r_i \leq \alpha_j^\delta - \chi_j^\delta \\ \frac{\alpha_j^\delta!(N_T - r_i)!}{(\alpha_j^\delta - r_i)!N_T!}, \; if \; \alpha_j^\delta - \chi_j^\delta < r_i \leq \alpha_j^\delta \\ 0, \; if \; r_i > \alpha_j^\delta \end{cases} \tag{4}$$

We compute the expected number of links used to display clip i on H2O$_j$, $E_i^j = \sum_{\delta=1}^{N_T - j} P_j^\delta \times \delta$. The closed form formula for E_i^j depends on whether N_T is even or odd:

$$E_i^j = \begin{cases} \frac{2}{N_T}(\sum_{j=1}^{\lfloor \frac{N_T}{2} \rfloor - 1} \sum_{\delta=1}^{N_T - j} P_j^\delta \times \delta), \; if \; N_T \; is \; even \\ \frac{2}{N_T}(\sum_{j=1}^{\lfloor \frac{N_T}{2} \rfloor - 1} \sum_{\delta=1}^{N_T - j} P_j^\delta \times \delta)) + \frac{\sum_{\delta=1}^{\lfloor \frac{N}{2} \rfloor - 1} P_j^\delta}{N_T}, else \end{cases} \tag{5}$$

Given a clip i with r_i replicas, $\sum_{j=1}^{\mathcal{N}} E_i^j \beta_i$ defines the bandwidth required to display clip i.

4. PERFORMANCE COMPARISON

In our experiments, we assumed a string topology with $\mathcal{N}$ nodes and a fixed bandwidth L between nodes. Each node is configured with 100 Gigabytes of storage. We simulated a skewed distribution of access to the m clips using a Zipfian distribution with a mean of 0.27. This distribution is shown to correspond to sale of movie theater tickets in the United States [4]. An experiments starts with the computation of the number of replicas per clip using one of the strategies with the available storage of $\mathcal{N}$ nodes. Next, each node requests a clip using a random number generator that is conditioned based on the skewed distribution of access. We consider two alternative policies for servicing these $\mathcal{N}$ requests: First Come First Serve (FCFS), and Least Bandwidth First (LBF). With FCFS, requests are scheduled one at a time starting with the first node. Once the available bandwidth is exhausted, the remaining nodes must wait until one or more active displays complete. With LBF, requests are sorted based on the display bandwidth requirement of their referenced clip in ascending order. The simulator schedules requests one at a time until either all requests are scheduled or the available bandwidth is exhausted. The reported average startup latency corresponds to the average delay observed by all requests (including the zero startup latency observed by those requests that were scheduled immediately). Approximating LBF in a distributed manner is a challenging task. We investigate it in order to demonstrate that its impact is marginal with the right replication strategy.

As a yard-stick, we analyzed a Random replication strategy. This strategy selects a clip randomly and constructs a replica for this clip. This replica is assigned to a randomly chosen node that does not contain a replica of this clip. If the total number of replicas for this clip equals $\mathcal{N}$ then this clip is removed from further consideration. Note that the same clip might be chosen repeatedly due to random chance. The termination condition of Random is reached when either the total storage capacity is exhausted or $\mathcal{N}$ replicas are created for all clips. In the reported experiments, we invoked Random 100 times and observed its maximum number of displays. This number is presented in all cases.

We conducted many experiments with different clip sizes, display times, and bandwidth requirements. In the following, we present two sets of experiments that summarize all our observations. These experiments are termed homogeneous and heterogeneous. The homogeneous experiment consisted of 100 video clips with an average bandwidth requirement of 4 Mbps, L=4 Mbps. All clips are equi-sized and have a display time of 120 minutes. Figure 3 shows our obtained results with FCFS and LBF policies with the alternative replications strategies. In these figures, the x-axis is the number of nodes and the y-axis is the number of simultaneous displays that can be activated. The bandwidth of a link between two nodes is 4 Mbps. The obtained results show an identical performance with DT, SZ, and BW replication strategies. This is because: a) clips are identical, and b) there is dependency between the size, bandwidth, and display time of a clip, i.e., size is a function of bandwidth and display time. With a given strategy, say BW, the square root (BW$\sqrt{Freq}$) replication strategy is superior to both proportional (BW$Freq$) and uniform (BW). Using the LBF approach, BW$Freq$ approximates BW$\sqrt{Freq}$. This is because it services those requests that reference frequently accessed clips first in order to initiate the display of a larger number of requests. BW cannot do the same because it replicates the clips uniformly, misleading this strategy to construct a large number of replicas for all clips. It is important to note that a distributed implementation to approximate LBF is non-trivial. Thus, the FCFS results are more realistic than LBF results.

With the heterogeneous experiment, the database consisted of a mix of audio and video clips: 445 video clips and 3555 audio clips. While the video clips continue to require 4 Mbps for their display, the audio clips require 340 Kbps for their display. The 3555 audio clips are evenly divided among three distinct groups: those with a display time of 2, 4, and 8 minutes. The 445 video clips are evenly divided among four unique groups: those with a display time of 30, 60, 90, and 120 minutes. The total size of this repository is 1 Terabyte. We continue to use the Zipfian distribution with a mean of 0.27 among all 4000 clips. Assuming that the 4000 clips are sorted based on their frequency of access, every 10^{th} clip is a video clip assigned in a round-robin manner to the four different video clip types: Clip 10 is 30 minutes long, Clip 20 is 60 minutes long, Clip 30 is 90 minutes long, Clip 40 is 120 minutes long, etc. The audio clips are also assigned in a round-robin manner starting with the shortest clip first: Clip 1 has a display time of 2 minutes, Clip 2 has a display time of 4 minutes, Clip 3 has a display time of 8 minutes, etc. In essence, the short audio (video) clips are accessed more frequently than the longer audio (video) clips.

Figure 4 shows the number of simultaneous displays initiated by each replication strategy with the FCFS policy as a function of the number of nodes using the string topology. The obtained results show replication based on bandwidth (BW) is a superior alternative to replication using either display time (DT) or size (SZ). In order to explain this, recall that BW, DT, and SZ ignore the frequency of access to the clips. The video bandwidth requirement (4 Mbps) is twelve times the audio bandwidth requirement (340 Kbps), motivating BW to replicate video clips 12 times more frequently than the audio clips. With DT, the average video display time (75 minutes) is 14 times the average audio display time (5.33 minutes), resulting in the video clips to be replicated 14 times the audio clips. With SZ, the average video clip size (2.25 Gigabytes)

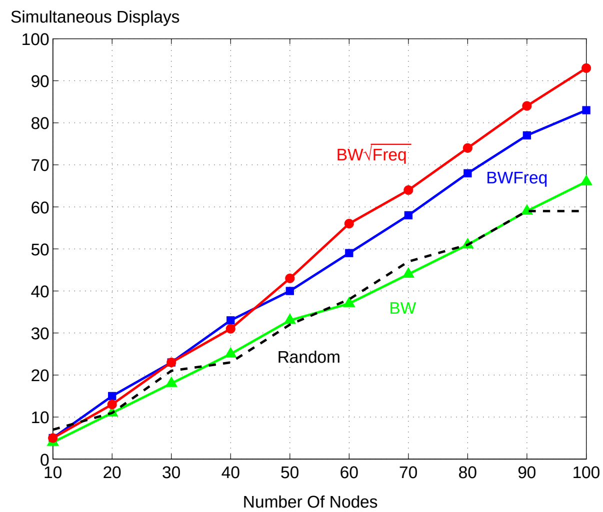

3.a First Come First Serve

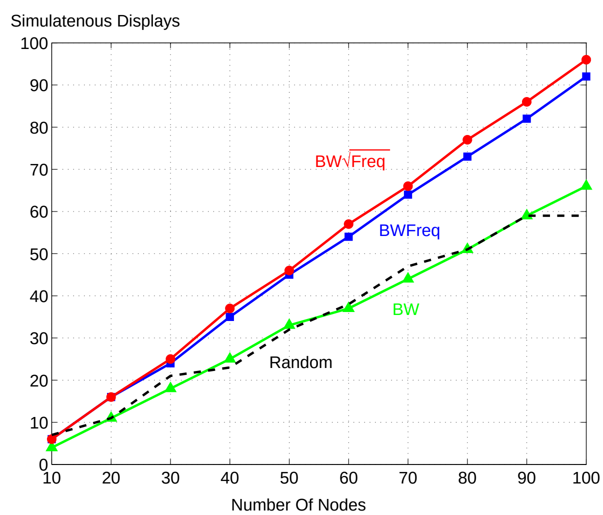

3.b Least Bandwidth First

Figure 3: Number of simultaneous displays with alternative replication strategies and a homogeneous mix of continuous media clips.

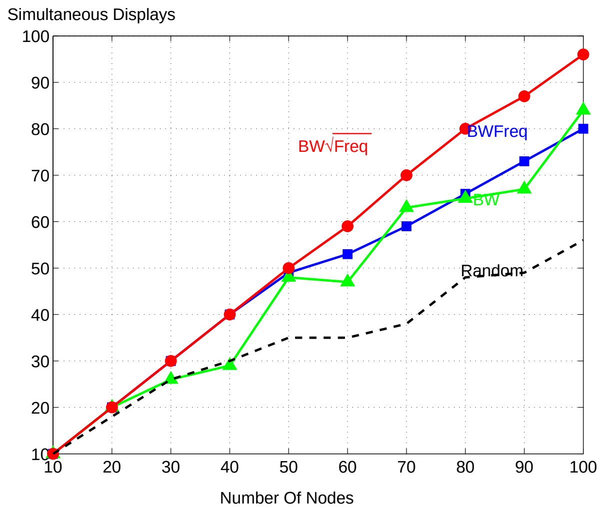

4.a Bandwidth

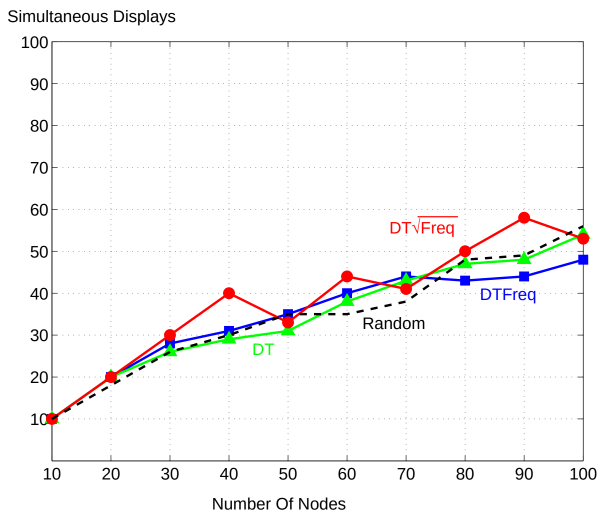

4.b Display time

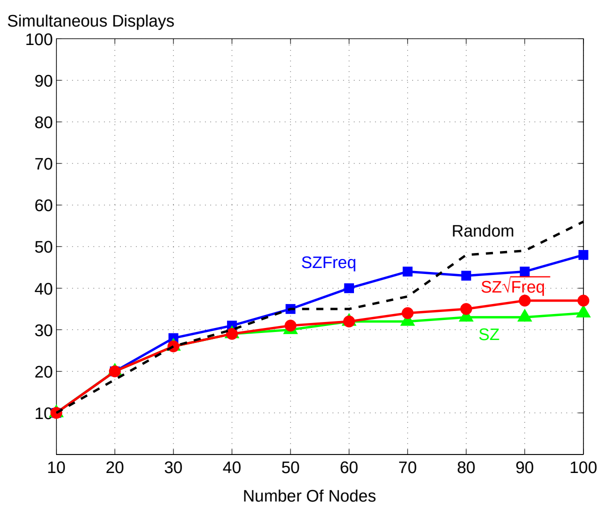

4.c Size

Figure 4: Number of simultaneous displays with alternative replication strategies and a heterogeneous mix of continuous media clips (FCFS policy).

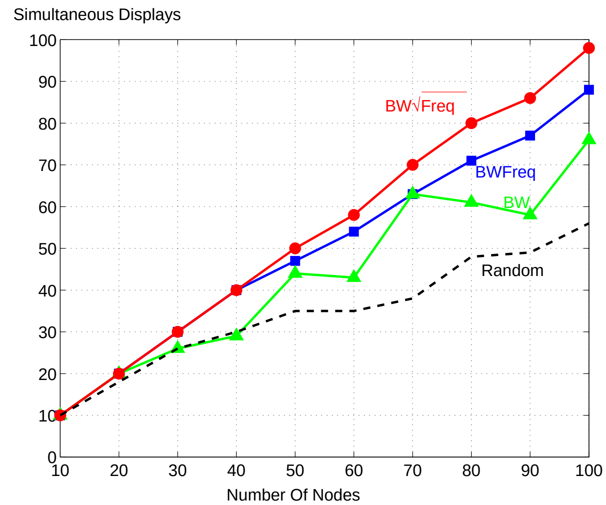

5.a Bandwidth

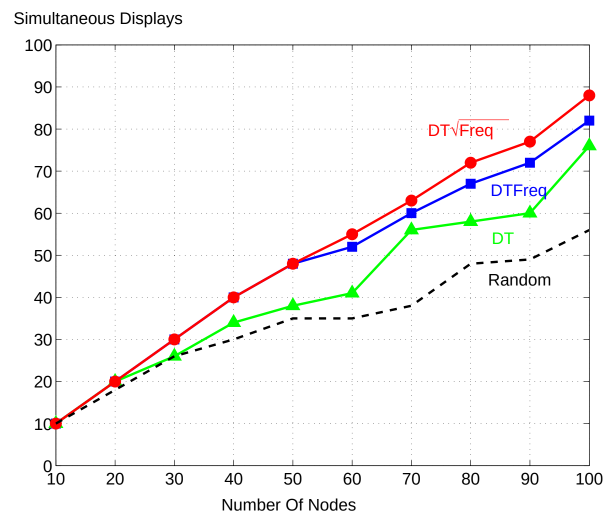

5.b Display time

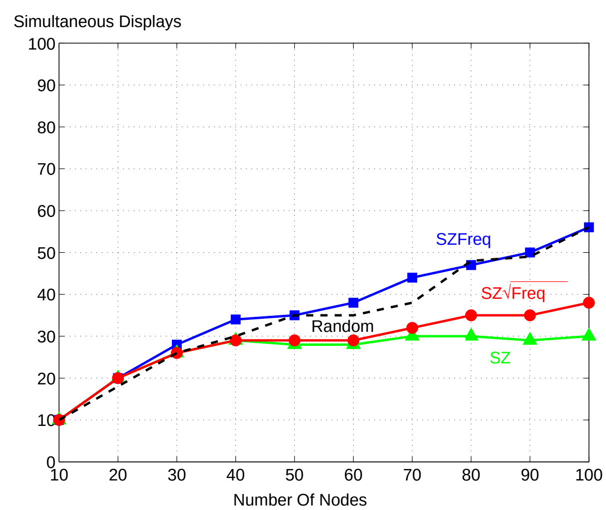

5.c Size

Figure 5: Number of simultaneous displays with alternative replication strategies and a heterogeneous mix of continuous media clips (LBF policy).

is 169 times the average audio clip size (13 Megabytes), resulting in a huge number of replicas for the video clips. With the workload referencing audio clips more frequently, SZ exhausts the available bandwidth with a few displays and cannot support as many simultaneous displays as the other two techniques. SZ$Freq$ utilizes the frequency of access to the clips in order to reduce the impact of the size ratio between video and audio clips. SZ$\sqrt{Freq}$ performs in the middle of SZ and SZ$Freq$ because it gives a lower weight to the frequency of access.

Figure 5 shows the number of simultaneous displays supported with LBF. LBF has a significant impact on replication strategies based on uniform, i.e., BW, DT, and SZ. With BW$\sqrt{Freq}$, the results show that employing LBF provides only marginal benefits when compared with FCFS. This suggests that development of algorithms to approximate LBF in a distributed manner will not provide significant benefit.

Finally, note that a full replication of heterogeneous mix of clips would require one Terabyte of storage from each node. This would enable all nodes to display a clip. BW$\sqrt{Freq}$ approximates this same objective with 100 Gigabytes of storage per node.

5. VALIDATION OF ANALYTICAL MODELS

We used a simulation study to validate the analytical models and its obtained results. The simulation model consists of $\mathcal{N}$ H2O devices organized in a string topology. The radio range of each device is 150 meters and the distance between two devices is 140 meters. The simulation model includes a centralized base station configured with an admission control component. This centralized admission control maintains the string topology of the network and the neighbor relationship between two nodes. Each request by a H2O device to display a clip X is directed to this admission control. Note that a string topology limits the number of possible paths between the displaying node and the closest replica(s) of X to either one (with one replica) or two (with two or more replicas). The admission control maintains the available bandwidth for each path based on the number of displays it has admitted into the system. If this bandwidth exceeds the bandwidth required to display X then the request is admitted and assigned to the node participating in the path with the highest available bandwidth. Otherwise, the request is rejected.

In our experiments, we invoked a replication strategy to compute the number of replicas for each clip. Next, these replicas are randomly assigned to the nodes with at most one replica per node. Subsequently, each node issues a request for a clip conditioned based on the assumed Zipfian distribution of access, forwarding this request to the centralized admission control. The admission control either accepts or rejects this request. The obtained results showed fragmentation of the available network bandwidth due to (a) placement of replicas, and (b) the identity of a clip referenced by a displaying H2O device. This is best illustrated with an example. Consider the topology of Figure 2.a showing a nine node string topology. If a display on node 2 references a clip on node 3 and exhausts the available bandwidth between these two nodes, then node 3 cannot display a clip with its only replica residing on node 1 even though there is bandwidth available on the link connecting node 1 to 2. This limitation is termed network bandwidth fragmentation. The analytical models of Section 3 do not model this limitation. To compensate, at the end of each experiment, we measure both (a) the number of simultaneous displays and (b) the total idle bandwidth in the system. We divide the later with the average bandwidth requirements of the clips to compute the number of displays that would have been supported if bandwidth fragmentation was prevented. This number is added to the total number of displays supported by an experiment.

Preliminary results suggest a maximum of 30% difference between the simulation and analytical models. We hope to conduct a more comprehensive evaluation and include the obtained results in a more extensive version of this study.

6. FUTURE RESEARCH

The primary contributions of this paper is three different replication strategies for continuous media and their evaluation with $\mathcal{N}$ H2O devices organized in a string topology. The primary objective of these techniques is to compute the number of replicas for a heterogeneous mix of continuous media. The obtained results are as follows. First, a replication scheme based on bandwidth (BW) is superior to alternatives that consider either size or display time. Second, this technique maximizes the number of simultaneous displays when it is extended to consider the square root of each clip's access frequency, i.e., BW$\sqrt{Freq}$ outperforms both BW and BW$Freq$.

While these results are specific to the string topology, we hypothesize they hold true for other topologies such as grid, graph, etc. In the immediate future, we intend to extend the analytical models of Section 3 to other topologies. In addition, we intend to use our simulation model for several other purposes. First, to investigate a hybrid replication strategy that considers both the display time of a clip and its bandwidth requirement. The analysis of Section 4 assumed an environment where all requests arrive at the same time. If requests arrive at random times, we speculate both the display time and bandwidth requirements of a clip to become significant. Second, to analyze topics such as placement of data and admission control. Both topics significantly impact the performance observed from a wireless topology. Third, it would be useful to study heterogeneous environments where (1) the storage capacity of H2O nodes might be different, (2) the bandwidth of the link between nodes might be different. The last item is specially important because the bandwidth of a link between two H2O devices is dependent on the number of other H2O devices in their radio range, see Section 1.

7. ACKNOWLEDGMENTS

This research was supported in part by an unrestricted cash gift from Microsoft research. We wish to thank Neha Chaudhry for her assistance with conducting the experiments and producing the graphs presented in this paper.

8. REFERENCES

[1] J. Almeida, D. Eager, and M. Vernon. A Hybrid Caching Strategy for Streaming Media Files. In *Proceedings of SPIE/ACM Conference on Multimedia Computing and Networking*, 2001.

[2] C. Chou, L. Golubchik, and J. C. S. Lui. A Performance Study of Dynamic Replication Techniques in Continuous Media Servers. In *Measurement and Modeling of Computer Systems*, pages 202–203, 1999.

[3] E. Cohen and S. Shenker. Replication Strategies in Unstructured Peer-to-Peer Networks. In *Proceedings of the ACM SIGCOMM*, August 2002.

[4] A. Dan, D. Dias, R. Mukherjee, D. Sitaram, and R. Tewari. Buffering and Caching in Large-Scale Video Servers. In *Proc. of COMPCON*, 1995.

[5] D. Eager, M. Ferris, and M. Vernon. Optimized Regional Caching for On-Demand Data Delivery. In *Proceedings of MMCN*, 1999.

[6] J. Gemmell, H. M. Vin, D. D. Kandlur, P. V. Rangan, and L. A. Rowe. Multimedia Storage Servers: A Tutorial. *IEEE Computer*, 28(5):40–49, 1995.

[7] S. Ghandeharizadeh, S. H. Kim, W. Shi, and R. Zimmermann. On Minimizing Startup Latency in Scalable Continuous Media Servers. In *Proceedings of Multimedia Computing and Networking Confernece*, February 1997.

[8] S. Ghandeharizadeh, B. Krishnamachari, and S. Song. Placement of Continuous Media in Wireless Peer-to-Peer Networks. *IEEE Transactions on Multimedia*, 2004.

[9] S. Ghandeharizadeh and R. Muntz. Design and Implementation of Scalable Continuous Media Servers. *Parallel Computing*, 24(1):91–122, May 1998.

[10] S. Ghandeharizadeh and L. Ramos. An Evaluation of Three Replication Strategies for Continuous Retrieval of Multimedia Data. In *International Symposium on Next Generation Database Systems and Their Applications*, September 1993.

[11] Y. Guo, S. Sen, and D. Towsley. Prefix Caching Assisted Periodic Broadcast: Framework and Techniques to Support Streaming for Popular Videos. In *Proceedings of ICC*, 2002.

[12] S. Sen, J. Rexford, and D.Towsley. Proxy Prefix Caching for Multimedia Streams. In *Proceedings of IEEE INFOCOM*, 1999.

[13] R. Tewari, H. N. Vin, A. Dan, and D. Sitaram. Resource Based Caching for Web Servers. In *Proceedings of SPIE/ACM Conference on Multimedia Computing and Networking*, 1998.

[14] B. Wang, S. Sen, M. Adler, and D. Towsley. Optimal Proxy Cache Allocation for Efficient Sreaming Media Distribution. In *Proceedings of IEEE INFOCOM*, 2002.

[15] Y. Wang, Z. Zhang, D. Du, and D. Su. A Network Conscious Approach to End-to-End Video Delivery over Wide Area Networks Using Proxy Servers. In *Proceedings of IEEE INFOCOM*, 1998.

[16] J. Wolf, P. Yu, and H. Shachnai. DASD Dancing: A Disk Load Blanacing Optimization Scheme for Video-on-Demand Computer. In *Proceedings of ACM SIGMETRICS*, pages 157–166, May 1995.