Comparison of Replication Strategies for Content Availability in C2P2 Networks

Shahram Ghandeharizadeh[1], Shyam Kapadia[1], Bhaskar Krishnamachari[1,2]
[1]Department of Computer Science
[2]Department of Electrical Engineering
University of Southern California
Los Angeles, CA 90089, USA
{shahram,kapadia,bkrishna}@usc.edu

ABSTRACT

This study investigates alternative continuous media replication techniques and their impact on content availability in a mobile car-to-car peer-to-peer (C2P2) network of devices. Using aggregate availability latency as a metric, we compare a simple random replication mechanism with a family of techniques that compute the degree of replication for each title based on its popularity, i.e., frequency of access. We use a simulation study along with some supporting analytical analysis for this comparison. Obtained results demonstrate the following key lesson. When total storage capacity of the network is significantly larger than the clip repository size, a random replication technique is sufficient. Otherwise, there is a large parameter space where the frequency-based replication schemes provide superior performance.

Categories and Subject Descriptors

C.2.1 [Network Architecture and Design]: Wireless Communication

General Terms

Design, Performance, Verification

Keywords

Peer-to-Peer, Content Availability, Replication, Mobility, Simulation, Analysis

1. INTRODUCTION

Advances in computer processing, storage performance and wireless communications have made it feasible to envision on-demand continuous media streaming between mobile vehicles. A vehicle is equipped with a Car-to-Car Peer-to-Peer [8] (C2P2) device consisting of several gigabytes of storage, a fast processor and a wireless interface with bandwidths of several tens of Mbps [1]. The radio range of each C2P2 device might be in the order of a few hundred feet. C2P2 devices collaborate to form an ad-hoc network that facilitates exchange of continuous media from a data producing C2P2 device to a displaying C2P2 device.

A C2P2 device may participate in several different roles simultaneously. First, it may display a title. Second, it may stream a clip from its local storage for display at another C2P2 device. Third, it may route a stream from a data producing C2P2 device to another C2P2 device that is displaying this data. Fourth, a combination of either two or all scenarios may occur simultaneously.

The focus of this paper is on content availability in a C2P2 network. The vehicular entertainment system offers its user a list of available movie titles [1]during the car's journey. A particular title is indicated as available only if sufficient replicas of that title are expected to be encountered in the vicinity of the car to enable successful viewing. However, a title may have a certain time delay after which it is available. To capture this, we define a related QoS metric for content availability. Given that a title is available, the vehicle entertainment system may also provide the earliest time after which the user can view that title. This is termed the availability latency of a title. An availability latency of 0 indicates that the title is available immediately.

Both the list of titles and their availability latencies are dependent on the number of replicas per title that the client C2P2 device might encounter. The number of replicas for a title in turn depend on its frequency of access which reflects its popularity. A title with a high access frequency warrants more replicas than one with a low access frequency. At the same time, the replicas for the titles may be constrained by the system storage. In this study, we introduce storage constraints per C2P2 so that the system storage is limited. Given the frequency of access to the titles, we investigate how many replicas per title should be constructed in a C2P2 network.

We extend the availability latency metric to incorporate the titles' frequency of access. Specifically, the metric used in this paper is the weighted sum of the availability latency of the different movie titles. The weights are proportional to the individual title access frequencies. We address the following questions: (1) Is a random replication scheme which

MDM 2005 Ayia Napa Cyprus

---

[1] We have used the terms titles and clips interchangeably in this paper.

is blind to the popularity of the movie titles good enough to provide a low aggregate availability latency? (2) If sophisticated replication schemes are needed then how large is the design space of interest? (3) Given the frequency of access to the movie titles which specifies their popularity, what would be the optimal number of replicas per title to provide a low aggregate availability latency?

The presence of mobility adds various challenges in the context of this environment. First, the network topology is dynamically changing. In some cases, the network connectivity between multiple devices might be intermittent [12], hence the overall network connectivity can no longer be guaranteed at all times. Second, it is non-trivial to predict the amount of resources available in the future.

Our primary contributions are as follows. In this paper, we investigate frequency-based replication strategies for content availability in mobile ad hoc networks. Using a Markovian mobility model we show that the optimum degree of replication for each clip is a function of its access frequency and the system storage. Specifically, we introduce a family of techniques that compute the degree of replication for a title as a power law function of its frequency of access. The exponent value employed by a technique differentiates it from the other members of this family. By using two replication schemes, namely ones in which the title replication is a linear and a square root function of the title access frequency respectively, along with a random replication scheme which is blind to the title access frequency, we show subsets of the design space in which one outperforms the others. This shows that no one scheme is ideal for all scenarios. Moreover, it is beneficial to use intelligent replication techniques in settings involving reasonable storage constraints.

The rest of this paper is organized as follows. Section 2 gives a brief overview of the related work in the area. Section 3 presents the general framework for the family of frequency-based replication schemes used in our study. Section 4 describes the details of our simulation study and our main results in detail. Section 5 provides an analytical evaluation of the replication schemes. Finally, Section 6 presents brief conclusions, some discussion and our future research directions.

2. RELATED WORK

Techniques to determine number of replicas are similar to assigning seats to different parties as a function of their popularity, i.e., ratio of votes casted for a party to the total vote. Webster's divisor method and its alternatives attributed to Hamilton, Adams and Jefferson allocate storage to each replica (assign seats to a party) as a linear function of its frequency of access (popularity). The divisor technique has been employed to determine the number of replicas of a video clip in a distributed video-on-demand architecture [13]. Our proposed framework captures this technique by setting a key parameter, denoted as n (see Section 3), to one.

Techniques to compute number of replicas for objects have been studied for both peer-to-peer networks [4, 11] and mesh community networks [6, 9]. Mobility of nodes is our primary contribution and separates these prior studies from the work presented here. In the following, we provide an overview of these prior studies. Replication of objects is important to their discovery in an un-structured peer-to-peer network. A smart replication technique minimizes search size, defined as the number of walks required to locate a referenced data item. In [4], the divisor method is compared with one that employs the square root of the frequency of access, demonstrating the superiority of the later.

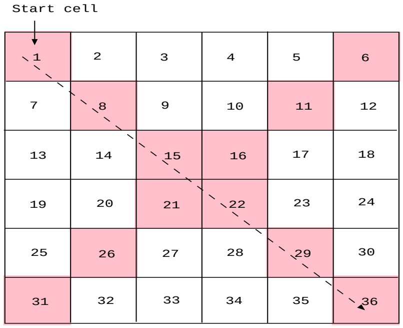

Figure 1: An example 6×6 map.

In mesh community networks, stationary devices configured with wireless devices form an ad-hoc network. In these networks, network bandwidth is a scarce resource that must be managed with care. In addition to the frequency of access to a clip i (f_i), one may determine its number of replicas (r_i) based on its other attributes such as size (S_i), display time (Δ_i), and display bandwidth (β_i) requirement. There exists a family of replication technique that determines the number of replicas based on $(f_i \times \beta_i)^n$ [6, 9]. We show that $n = 0.5$ is typically superior to other alternatives. These prior studies do not consider mobility of devices and the concept of availability latency.

Another related area is that of Intelligent Transportation Systems (ITS). The vehicular network is viewed as a MANET and messages are forwarded from one vehicle to another realizing several applications like vehicle accident notification broadcast, pre-emptive emergency vehicle arrival information etc [2, 3]. The lessons learnt from this work may be directly applicable to some ITS problems.

3. GENERAL FRAMEWORK

This section presents a general framework that supports a family of replication strategies. Table 1 summarizes the notation of the parameters used in the paper. Assume a network of N mobile C2P2 devices, each with storage capacity of α bytes. The total storage capacity of the system is $C = N * \alpha$. There are T clips in the database, each with a display time of Δ_i seconds and display bandwidth requirement of β_i. Hence the size of each clip is given by $S_i = \Delta_i * \beta_i$. The frequency of access to clip i is denoted as f_i with $\sum_{j=1}^{T} f_j = 1$. Let the trip duration of the client C2P2 under consideration be γ.

We now define the normalized frequency of access to the title i, denoted R_i, is:

$$R_i = \frac{(f_i)^n}{\sum_{j=1}^{T} (f_j)^n}; \; 0 \leq n \leq \infty \tag{1}$$

R_i is normalized to a value between 0 and 1. The number of replicas for title i, denoted as r_i, is:

$$r_i = \min\left(N, \max\left(1, \left\lfloor \frac{R_i * N * \alpha}{S_i} \right\rfloor\right)\right) \tag{2}$$

Database Parameters	
T	Number of clips.
S_i	Size of clip i
Δ_i	Display time of clip i.
β_i	Bandwidth requirement of clip i.
f_i	Frequency of access to clip i.
Replication Parameters	
R_i	Normalized frequency of access to clip i, $R_i = \frac{(f_i)^n}{\sum_{j=1}^{T}(f_j)^n}$; $0 \leq n \leq \infty$
r_i	Number of replicas for clip i, $r_i = \min(N, \max(1, \lfloor \frac{R_i * N * \alpha}{S_i} \rfloor))$
n	Characterizes a particular replication scheme.
δ_i	Average availability latency of clip i
δ_{agg}	Aggregate availability latency for replication technique using the n^{th} power, $0 \leq n \leq \infty$, $\delta_{agg} = \sum_{j=1}^{T} \delta_j * f_j$
C2P2 System Parameters	
N	Number of C2P2 devices in the system.
α	Storage capacity per C2P2.
γ	Trip duration of the client C2P2.
C	Total storage capacity of the C2P2 system, $C = N * \alpha$.

Table 1: Terms and their definitions

This defines a family of replication schemes that computes the degree of replication of title i as the n^{th} power of its frequency of access. Hence r_i lies between 1 and N. Note that r_i includes the original copy of a clip. One may simplify Equation 2 by replacing the max function with $\lfloor \frac{R_i * N * \alpha}{S_i} \rfloor$. This would allow the value of r_i to drop to zero for a clip i. This means that there is no copy of the clip in the C2P2 network. In this case, a hybrid framework might provide access to the clip i. For example, a base station employing IEEE 802.16 [10] might facilitate access to a wired infrastructure with remote servers containing the clip i.

The availability latency for a title i, denoted as δ_i, is defined as the time after which a client C2P2 will find at least one replica of the title accessible to it, either directly or via multiple hops, for the title display time (Δ_i). If this condition is not satisfied for a title i, then we set δ_i to γ. This indicates that title i was not available to the client during its journey. Also, if Δ_i exceeds γ for a certain title i then we set δ_i to γ. We are interested in the availability latency observed across all the titles. Hence we augment the δ_i for every title i with its f_i. This is termed the aggregate availability latency (δ_{agg}) metric. It is computed as follows. For each title i, calculate the average availability latency (δ_i) based on the particular replication scheme of interest. Then these availability latencies are combined into a single metric:

$$\delta_{agg} = \sum_{i=1}^{T} \delta_i * f_i \tag{3}$$

Both δ_{agg} and δ_i depend on the value chosen for n. The core problem of interest here is to keep the aggregate availability latency as low as possible by tuning the title replication levels, in the presence of storage constraints. We assume that the database size is smaller than the total storage capacity of the system, $\sum_{i=1}^{T} S_i \leq C$. Otherwise, clips cannot be replicated when at least one replica of a clip must be present in the system. More formally, the optimization problem can be stated as,

$$\textbf{Minimize } \delta_{agg}\textbf{, subject to } \sum_{i=1}^{T} S_i \leq C \tag{4}$$

Implicit in this formulation is the design variable, namely, the desired replication for each title. The value of n in Equation 1 determines a r_i value for each clip i with the objective to minimize[2] δ_{agg}. This minimization is a challenge when the total size of the database exceeds the storage capacity of a C2P2 device, $\sum_{i=1}^{T} S_i > \alpha$. Otherwise, the problem is trivial and can be solved by replicating the clip repository on each device.

The optimization space that defines what value of n provides the best δ_{agg} is quite large and consists of the following parameters: (i) density of C2P2 devices, (ii) title display time, (iii) size of the clip, (iv) display bandwidth per clip, (v) the number of clips, (vi) storage per C2P2 device, (vii) trip duration, (viii) frequency of access to the clips, and (ix) mobility model for the C2P2 devices. This study does not solve the problem directly. Instead we use simulation studies to show how the various parameter settings impact the objective function. The optimization problem can be solved only in idealized settings such as the one we explore analytically in Section 5.

4. SIMULATION STUDY

In this section, we first outline the assumptions that constrain general framework presented in Section 3 for the purpose of the simulation study. Subsequently, we describe our experimental setup and present the simulation results.

4.1 Simulation Model

We assume a repository of homogeneous clips with identical bandwidth requirement, display time, and size ($\beta_i = \beta$, $\Delta_i = \Delta$, $S_i = S$). Figure 1 shows a 6×6 grid used as the map for our experimental study. The map is divided into fixed size cells. Only C2P2s within a cell can communicate with each other either directly if they are in radio-range or via other C2P2s using multi-hop transmissions. In other words, the C2P2s within a cell form a connected sub-network. C2P2s in adjacent cells cannot communicate with each other. Without any loss of generality, to reduce the dimensionality of the problem, we express the title display time Δ as the amount of time required by a C2P2 equipped

---

[2] 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.

vehicle to travel Δ cells. We express α as the number of storage slots per C2P2. Each storage slot stores a clip fragment equivalent to a single cell worth of clip display time. Moreover, we assume the amount of data displayed in each cell is identical. Now, we represent both the size of a clip and the storage slots in terms of the number of cells. This means that a clip has a display time of Δ cells and a C2P2 has α units of cell storage. For example, a clip with display time of 4 cells ($\Delta = 4$) requires 4 storage slots and a C2P2 provides 100 storage slots ($\alpha = 100$).

The trip duration (γ) is also expressed as the number of cells traversed by the client during its journey. We also define availability latency (δ_i) for title i in terms of the number of cells. In other words, δ_i is the number of cells after which a client C2P2 will encounter a replica of the title, either directly or via multiple hops, for the title display time. Hence, the possible values of the availability latency are between 0 and γ. We only consider scenarios in which $\Delta \leq \gamma$. Assume that $\gamma = 6$. For a title i with $\Delta_i = 6$, δ_i is either 0 or 6. $\delta_i = 0$ means that at least one replica of that title was present in each of the 6 cells along the path of the client. $\delta_i = 6$ means that at least one cell along the path of the client was missing a replica of the title. Similarly, for title j with $\Delta_j = 5$, δ_j is either 0, 1 or 6. If $\delta_j = 0$, the client encountered at least one replica of title j along each of the first 5 cells along its path. If $\delta_j = 1$, the client encountered at least one replica of the title along the last 5 cells of its path, but not even a single replica in the first cell. Finally, $\delta_j = 6$ indicates that there were at least 2 cells along the path of the client, in which no replicas of title j were present. Similarly, the availability latency was calculated for each of the other titles for different title display times. In the simulation study, we allow for the possibility that not even a single copy of a title might be present within the C2P2 system. Hence, we rewrite Equation (2) as,

$$r_i = \min\left(N, \left\lfloor \frac{R_i * N * \alpha}{\Delta} \right\rfloor\right) \tag{5}$$

With Equation 5, $r_i \geq 1$ only when $\alpha \geq \Delta$. Otherwise, a C2P2 may not store even one copy of a title. In this case, δ_i is set to γ for the purposes of Equation 3. This indicates that the clip will not be encountered during the entire trip duration.

A Markovian mobility model was used to describe the movement of the cars which was probabilistic in nature. Each cell of the map constitutes a state. These states are self-contained and a transition from one state to another is independent of the previous history of a car in that state. The mobility model is weighted toward the diagonal both from the left to right and vice-versa (indicated by the gray boxes in Figure 1). The aggregate of the transitions from each cell (state) to every other state gives the probability transition matrix $Q = [q_{ij}]$ where q_{ij} is the probability of transition from state i to state j. Using Markov chains, it is possible to estimate the distribution of the steady-state probabilities of being in the various cells, by solving $\Pi = \Pi * Q$, where Π is the vector representing the steady-state probabilities of being in the various cells (states).

4.2 Experimental Study

In this section, we present the results of our simulation study. The experimental setup uses the 6×6 map shown in Figure 1. Assume that the client starts from cell 1 and travels along the path $\{1, 8, 15, 22, 29, 36\}$. Numbers in the bracket indicate the sequence of visited cell ids. Hence $\gamma = 6$.

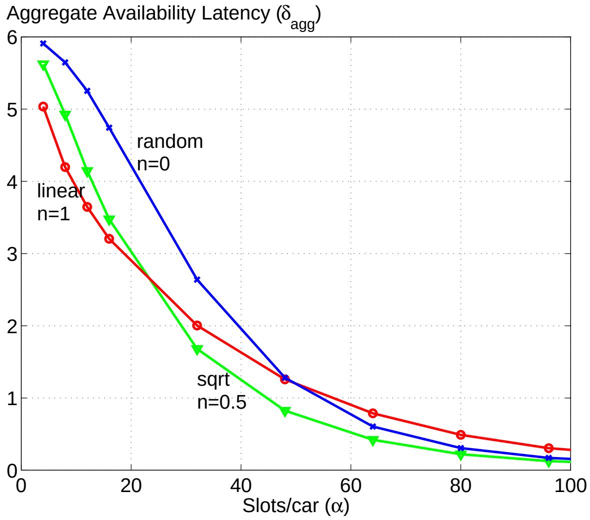

2.a) δ_{agg}

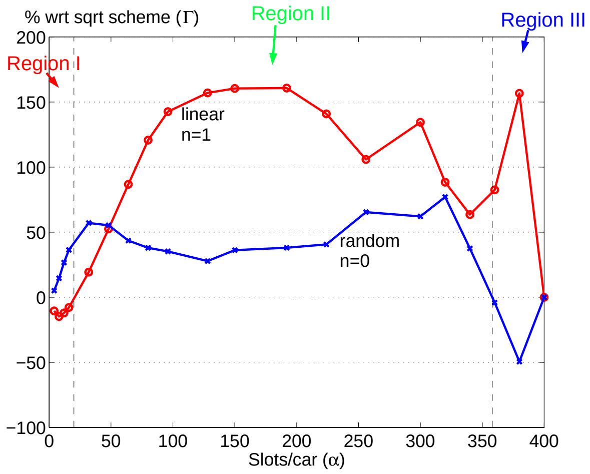

2.b) % Comparison wrt sqrt scheme (Γ)

Figure 2: Figure 2.a shows δ_{agg} of the sqrt, linear and random replication schemes versus α for $\Delta = 4$ and $N = 200$. Figure 2.b shows the % comparison of the linear and random schemes wrt the sqrt scheme for this scenario. Region I indicates the parameter space where $n = 1$ performs the best. Regions II and III indicate the space in which $n = 0.5$ and $n = 0$ perform the best, respectively.

For the initial set of experiments we chose $N = 200$ and $T = 100$. We simulated a skewed distribution of access to the T clips using a Zipfian distribution with a mean of 0.27. The distribution is shown to correspond to sale of movie theater tickets in the United States [5]. We examine a subset of the schemes using values of $n = 0$, $n = 1$ and $n = 0.5$. When $n = 1$, clips are replicated as a linear function of their frequency of access. This technique is termed linear. Similarly, $n = 0.5$ is termed the square root technique. Finally, $n = 0$ indicates that the degree of replication of each title is

$\frac{1}{T}$. This technique can be realized using a random number generator. Thus, it is termed random.

Initially, all cars are distributed across the cells of the map as per the steady state distribution which is determined by a random number generator initialized with a seed. Depending on the particular replication technique the replicas for each title are calculated using Equation 5 and then distributed across the car. A car only contains a maximum of one replica for a particular title. The distribution of title replicas across the cars is uniform. At each step, depending on the current car location, it moves to one of its adjoining cell (including itself) as governed by the mobility model. Another seed determines the choice of which cell a car moves to. Since $\gamma = 6$, each car performs six transitions according to the mobility model. We performed the comparisons for several different title distribution seeds starting from the same initial car positions. Next, we varied the initial car positions by changing the initial seed. Specifically, we chose 50 different initial seeds and for each of these we used 50 seeds that decide the distribution of the title replicas among the cars. Thus, each point in all the presented results is an average of 2500 simulations.

4.3 Simulation Results

We first present an overview of the key lessons learned from the simulation study. Subsequently, we provide a brief explanation of each lesson.

(a) The optimal value of n varies as a function of the scarcity of the network storage (see Figure 2.b). (b) When storage is scarce, the optimal aggregate availability latency is realized by using a higher value of n (see Figure 3 and Figure 4). (c) The random scheme, $n = 0$, is appropriate when storage is abundant relative to the repository size (see Figure 2).

When storage is extremely scarce, with larger title sizes ($\Delta > 1$), linear ($n = 1$) scheme provides the best performance. This is because it allocates more replicas for the popular titles at the cost of assigning very few for the remaining titles. In this case, the contribution to δ_{agg} is a function of the δ for the more popular titles since for the less popular titles there will be insufficient replicas to reduce their δ. On the other hand, since the random scheme is blind to the title access frequencies, on an average, it assigns equal number of replicas for each title thereby providing the worst performance.

The square root ($n = 0.5$) scheme assigns fewer replicas for the popular titles than the linear scheme. As we increase the amount of storage, there is a cut-off point along the storage axis, where allocating more replicas for the popular titles provides negligible improvement in δ_{agg}. It is beyond this point that the square root scheme starts outperforming the linear scheme. This is because the square root scheme can use the extra storage savings for allocating replicas for the less popular titles thereby reducing their δ.

To illustrate, Figure 2 shows the variation of δ_{agg} as a function of α for $\Delta = 4$. Since δ_{agg} is a function of the value of n, hence, here we denote it as $\delta_{agg}(n = i)$. For Figure 2.b, the y-axis represents the percentage comparison of the linear ($n = 1$) and the random ($n = 0$) schemes with respect to the square root ($n = 0.5$) scheme calculated as,

$$\Gamma = \left(\frac{\delta_{agg}(n = i) - \delta_{agg}(n = 0.5)}{\delta_{agg}(n = 0.5)} \right) \times 100 \qquad (6)$$

where $i = \{0, 1\}$.

Figure 2.b shows three distinct regions in which the schemes with $n = 0$, $n = 0.5$ and $n = 1$ perform well under certain parameter settings within the design space. For $\alpha <= 20$, the linear scheme ($n = 1$) performs the best. For $20 <= \alpha <= 360$, the square root scheme ($n = 0.5$) performs the best. Finally, beyond a value of 360 for α, the random scheme ($n = 0$) provides the best performance.

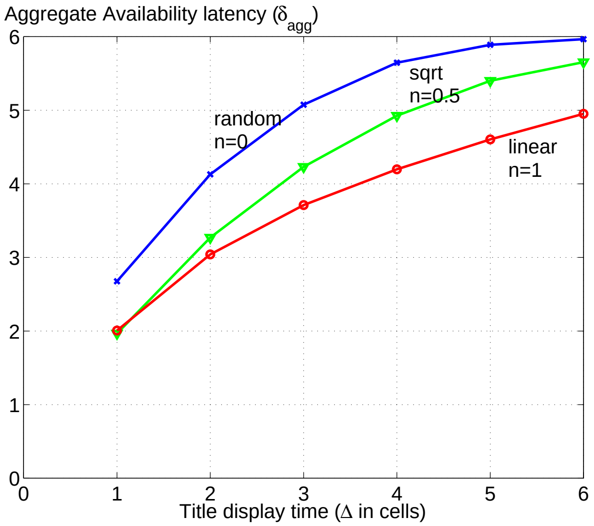

3.a) δ_{agg}

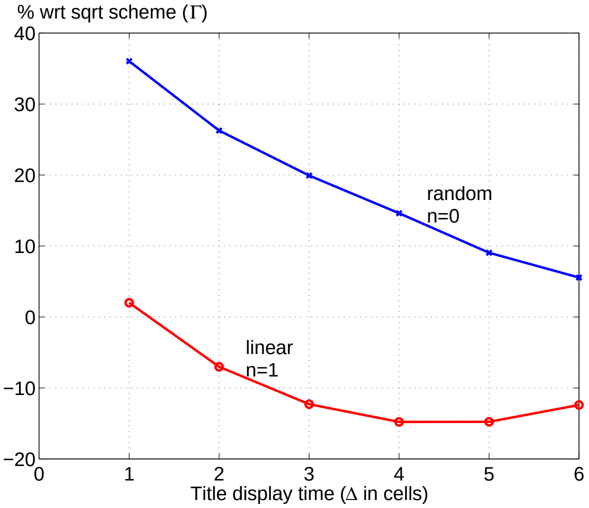

3.b) % Comparison wrt sqrt scheme (Γ)

Figure 3: Figure 3.a shows δ_{agg} of the sqrt, linear and random replication schemes versus Δ for a constant title display time:storage (Δ:α) ratio of 1:2,$N = 200$. Figure 3.b shows the % comparison of the linear and random schemes wrt the sqrt scheme for this scenario.

With $\Delta = x$ and $T = y$, the value of α needed to replicate the entire database on each car is $\alpha_{db} = x * y$. At a certain storage threshold (earlier than α_{db}), the random scheme assigns enough replicas to the popular titles to bring their δ down. In this case, all the titles have the same number of replicas, thereby producing a low δ for every title. Hence, from this point onward, random produces the lowest δ_{agg} and outperforms the other schemes. However, this point requires sufficient storage per car and hence a random scheme may be appropriate only for over-provisioned scenarios. As illustrated in Figure 2.b with $N = 200$, $T = 100$, and $\Delta = 4$, the storage threshold is around 360 slots per car. For $\Delta = 5$, and 6, this threshold is approximately 450 and 540, respectively. These are loose upper bounds.

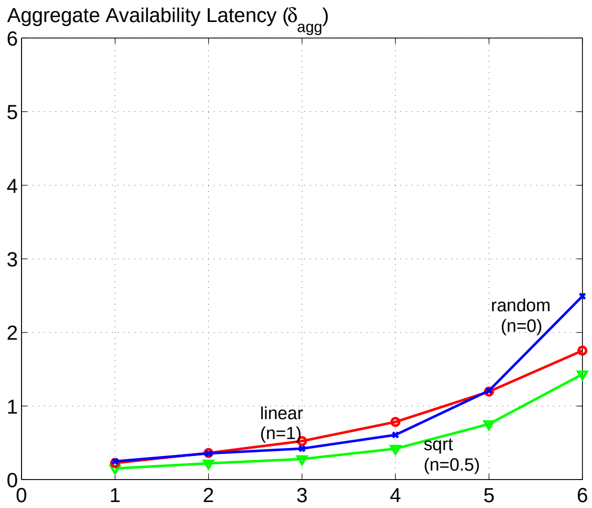

4.a) δ_{agg}

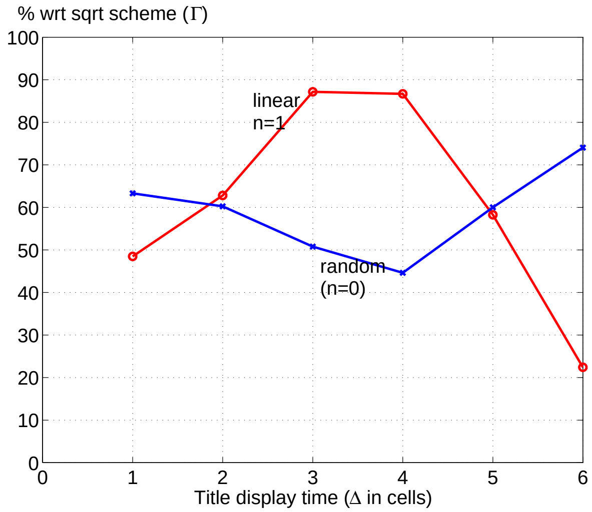

4.b) % Comparion wrt sqrt scheme (Γ)

Figure 4: Figure 4.a shows δ_{agg} of the sqrt, linear and random replication schemes versus Δ for a constant title display time:storage (Δ:α) ratio of 1:16, $N = 200$. Figure 4.b shows the % comparison of the linear and random schemes wrt the sqrt scheme for this scenario.

4.3.1 Aggregate availability latency for a constant Δ and α ratio

A constant title display time and storage ratio (Δ:α) indicates a constant number of titles per car irrespective of the Δ value. Figure 4 shows that for a ratio of $1:16$ (16 titles per car), $n = 0.5$ outperforms the other schemes by an order of magnitude, making the best possible use of the allocated storage. For a ratio of $1:2$, see Figure 3, $n = 1$ scheme performs the best while for a ratio of $1:90$, $n = 0$ provides a performance which is at least as good as the other schemes.

4.3.2 Aggregate availability latency as a function of car density (N)

Car density, which in turn affects the available storage in the system, has a major impact on the performance of δ_{agg} for all the schemes. With the decrease in the car density to $N = 100$, the number of replicas allocated by the schemes is reduced thereby giving comparatively larger values of δ_{agg} across the same storage axis. As α is increased, the drop in the δ_{agg} curves for all the schemes is not as steep as seen in the case with $N = 200$. Again, this is because the number of replicas is not increasing at such a high rate. The storage is reduced by an order of 2, hence a higher value of α is needed to produce the same drop in δ_{agg} as was seen in the case with $N = 200$ cars. This is observed across all values of Δ.

4.3.3 Aggregate availability latency as a function of the repository size (T)

We also performed a series of experiments to study the relative performance of the three schemes with the variation in T. Specifically for a given Δ and N, we varied T from 100 to 1000 in steps of 100. By increasing T, we reduce the number of replicas for a clip. In these experiments, we determined a loose lower bound on the value of α needed for each of the three schemes to perform the best in comparison with the other two. The same set of lessons (see beginning of this section) carries through and hence we have not repeated those results here.

For all experiments, we also calculated the standard deviations (SD) and the standard error of the mean (SEM). The 95% confidence intervals determined as $1.96 * SEM$ are quite small and accordingly the curves are quite smooth. However, the standard deviation is quite large, especially for the cases when δ_{agg} is low for high values of Δ and α. This is because a low latency requires the clip to be present in every cell along the journey depending on the value of Δ. As Δ increases, it becomes increasingly difficult to meet this condition thereby showing a high variance in δ_{agg}. The large SD value is an empirical observation about the nature of the random process.

5. ANALYTICAL EVALUATION

In this section, we present an analytical model to provide insights into the trends observed with the simulation results. We consider a homogeneous repository of clips in cars ($\beta_i = \beta$, $\Delta_i = \Delta$, $S_i = S$). Let the map comprise of c cells. Again, as in the simulation model, here we express γ, Δ, α, δ_i in terms of the number of cells (see Table 1 for notation). Given f_i for title i, the number of replicas for that title, r_i, under a given replication scheme are calculated using Equation 5. For simplicity, we assume that the replicas are uniformly distributed across the c cells. However, in practice, this distribution will converge to the stationary distribution dictated by the mobility model. Here, we present the analysis for clips with a display time of 1 ($\Delta = 1$). The analysis can be easily extended to clips with $\Delta > 1$.

Define, A_k, the event that a title i is encountered for the *first* time in the k^{th} cell (step). Let $P(A_k)$ denote the probability of event A_k occurring. Let p_k denote the probability of encountering title i in the k^{th} cell, given that it was not encountered in the previous $k-1$ cells. Note that p_k is a conditional probability. Then,

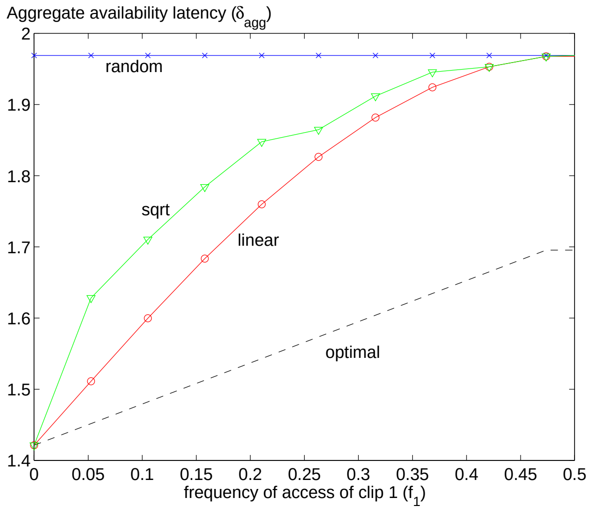

Figure 5: Analytical comparison of the sqrt, linear, random and optimal replication strategies with 10 C2P2 equipped vehicles.

$$p_k = 1 - (1 - \frac{1}{c-k+1})^{r_i};\ 1 \leq k \leq \gamma \tag{7}$$

Also, $P(A_k)$ is a joint probability since encountering a title for first time in the k^{th} cell indicates that it was not encountered in any of the previous $k-1$ cells. Clearly, p_j and p_{j-1} are not independent, hence we use the multiplication rule to obtain the value of $P_(A_k)$ as,

$$P(A_k) = p_k \prod_{j=1}^{k-1} (1 - p_j);\ 1 \leq k \leq \gamma \tag{8}$$

Let $P(A_{\gamma+1})$ denote the probability of not encountering the title i during the entire trip duration (γ). Hence,

$$P(A_{\gamma+1}) = \prod_{j=1}^{\gamma} (1 - p_j) \tag{9}$$

Then, the availability latency (δ_i) for title i is given by,

$$\delta_i = \sum_{k=1}^{\gamma+1} (k-1) P(A_k) \tag{10}$$

Finally using Equation 3, we obtain δ_{agg} for this replication scheme.

Note that the analytical model does not capture the correlation inherent in a realistic mobility model. Moreover, since the parameter space is quite large, analytically computing the storage levels where the replication schemes outperform each other is non-trivial. However, as we show below, the analysis gives some insights about the relative performance of the replication techniques.

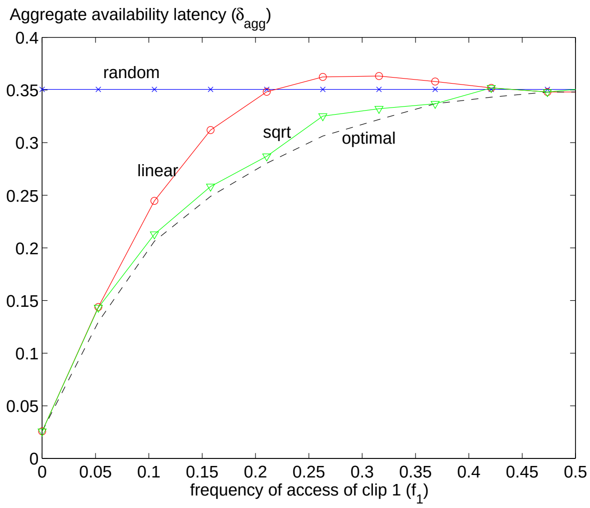

Figure 6: Analytical comparison of the sqrt, linear, random and optimal replication strategies with 50 C2P2 equipped vehicles.

Consider a scenario consisting of only 2 clips ($T = 2$). Let there be N cars distributed uniformly across the c cells. Also let $\gamma = 2$, $\Delta = 1$ and $\alpha = 1$. Hence, each car can have either one of the two clips and the availability latency per title is either 0, 1 or 2. We vary f_1 from 0 to 1. Note that, $f_2 = 1 - f_1$. Also, $0 \leq r_1 \leq N$ and $r_2 = N - r_1$. We obtain δ_i for each title i ={1,2} using Equations 7, 8, 9 and 10. Finally, we use Equation 3 to obtain δ_{agg}.

Assume $c = 10$ cells and consider δ_{agg} with the following values for N:{10,50,100}. As the number of cars increases, the storage and hence the number of replicas per title increases.

Figure 5, 6, 7 shows the comparison of square root, linear, random and optimal strategies for title 1 for three car densities. In this case, the optimal strategy is the one that, for a given f_1, chooses the best r_1 to produce the lowest possible δ_{agg}. Since there are only two clips, the number of replicas for title 2 (r_2) will be the compliment of that seen for title 1, i.e., $N - r_1$. Figure 8 shows optimal number of replicas for title 1 as a function of f_1 (which can be derived from the previous plots, by finding which replication degree minimizes latency for a given title frequency of access).

Linear replication strategy performs better than both the square root and the random schemes with low values of N ($N = 10$). However, it is not the optimal strategy. The optimal strategy in this case is generally to store only the most popular title, as seen from Figure 8. At an intermediate number of cars/storage level, the square root strategy performs the best and matches the optimal strategy quite well. With high values of N ($N = 100$), the random scheme performs the best, approximating the optimal strategy. These observations are consistent with the conclusions from the simulation results.

6. CONCLUSIONS AND FUTURE RESEARCH DIRECTIONS

This study presents a family of replication techniques for

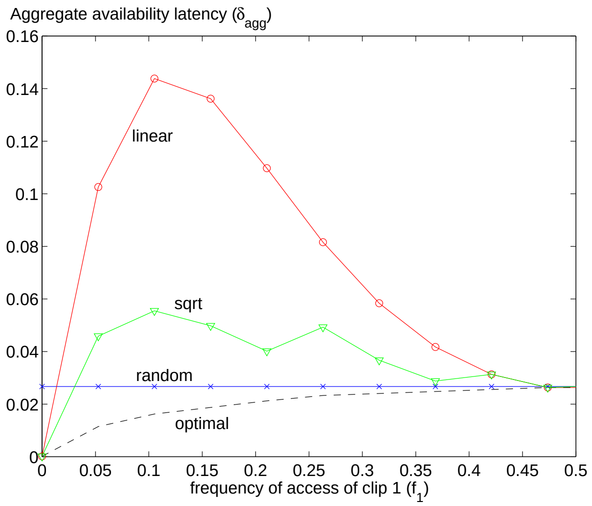

Figure 7: Analytical comparison of the sqrt, linear, random and optimal replication strategies with 100 C2P2 equipped vehicles.

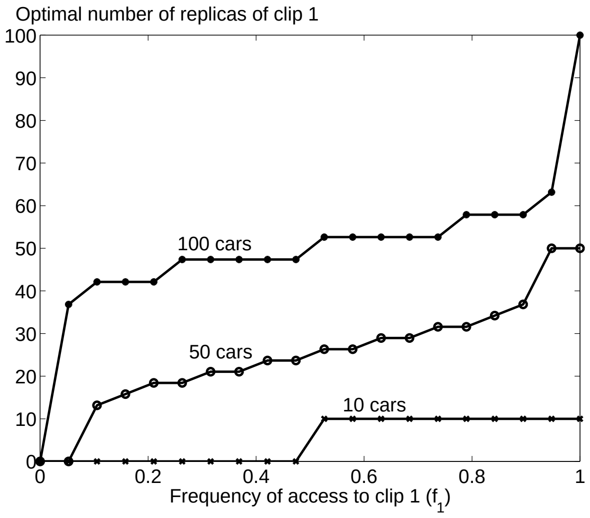

Figure 8: Optimal number of replicas for clip 1 versus its frequency of access for 3 different car densities (analogous to 3 different storage levels in this case)

a mobile ad hoc network of C2P2 devices. These techniques compute the degree of replication for each title as a power law function of its popularity, i.e., frequency of access. The exponent value (n) identifies a specific technique, see Equation 1. We considered three distinct exponent values: random (n=0), square root (n=0.5), and linear (n=1). We use aggregate availability latency as the metric to compare these alternative techniques. The value of this metric with a technique is impacted by a large number of system parameters. These include density of C2P2 devices, title display time, size of clip repository, trip duration, the mobility model, storage per C2P2 device and the popularity of the titles.

Our analysis reveals the following key lesson. While random (n=0) may suffice when storage is abundant relative to the repository size, there is a large parameter space where sophisticated frequency based techniques ($n > 0$) are superior. Moreover, our analytical results reveal that the square-root technique (n=0.5) is not necessarily the optimal one in those regions that it provides the best performance. Some other value of n besides $\{0,\ 0.5,\ 1\}$ may provide a better performance. However, computing this value is non-trivial.

We intend to extend this study in several directions. First, techniques other than random must be provided with the frequency of access to each clip. Our analysis ignores this overhead. Typically, this meta-data will be small and in the order of hundreds of bytes. The architectural framework of [7] assumes a control plane that compliments the ad-hoc data plane. This control plane might be realized using the low-bandwidth base stations of cellular telephones. It would facilitate exchange of meta-data such as frequency of access for the clips, the car density information etc.

Second, this study considered a homogeneous repository of clips with equal title display time (Δ). We intend to extend our evaluation to a heterogeneous repository of clips with different Δ.

Third, our study ignores placement of data. With mobility, the initial placement of the replicas gets quickly diffused. The mobility model dictates how replicas will be distributed over a geographical area over time. In this study, we have considered a Markov mobility model on a grid that serves as a map for a region. In order to make the model more realistic we can superimpose a square grid on the topographical map of a region e.g. an urban area within the city of Los Angeles. In this case, most of the transitions within the map will be weighted along the major freeways and the side streets would only provide a minor transition bias. Moreover, some of the vehicles may follow a deterministic route, e.g. buses, airport shuttle vans etc and may have different trip durations (γ). Our future work includes modeling a heterogeneous mix of known and unknown vehicular routes with different trip durations. We intend to develop dynamic data reorganization schemes that will reactively or preferably proactively reorganize the placement of data to ensure that the system promised service criteria are met.

Finally, vehicular ad-hoc networks are an emerging area and even though this study considers audio and video content, the nature of the content might be vary. For example, broadcast of traffic and emergency information might be facilitated using C2P2 networks. However, the lessons about the impact of different replication strategies on data availability will be applicable over a wide range of information content.

7. ACKNOWLEDGMENTS

This research was supported in part by NSF grant IIS-0307908, National Library of Medicine LM07061-01, and an unrestricted cash gift from Microsoft Research.

8. REFERENCES

[1] S. Bararia, S. Ghandeharizadeh, and S. Kapadia. Evaluation of 802.11a for streaming data in ad-hoc networks. In *4th Workshop of Applications and Services in Wireless Networks (ASWN) Boston, MA*, 2004.

[2] L. Briesemeister and G. Hommel. Role-based multicast in highly mobile but sparsely connected ad hoc networks. In *Proceedings of MobiHOC*, pages 45–50, August 2000.

[3] L. Briesemeister, L. Schafers, and G. Hommel. Disseminating Messages among Highly Mobile Hosts based on Inter-Vehicle Communication. In *Proceedings of IEEE Intelligent Vehicle Symposium*, pages 522–527, October 2000.

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

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

[6] S. Ghandeharizadeh and T. Helmi. An Evaluation of Alternative Continuous Media Replication Techniques in Wireless Peer-to-Peer Networks. In *Third International ACM Workshop on Data Engineering for Wireless and Mobile Access (MobiDE, in conjunction with MobiCom'03)*, September 2003.

[7] S. Ghandeharizadeh, S. Kapadia, and B. Krishnamachari. Pavan: A policy for content availability in vehicular ad-hoc networks. In *First ACM Workshop on Vehicular Ad Hoc Networks (VANET),Philadelphia, PA, October 2004*. ACM Press, 2004.

[8] S. Ghandeharizadeh and B. Krishnamachari. C2P2: A Peer-to-Peer Network for On-Demand Automobile Information Services. In *First International Workshop on Grid and Peer-to-Peer Computing Impacts on Large Scale Heterogeneous Distributed Database Systems (Globe'04)*, 2004.

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

[10] IEEE 802.16 Working Group. Emerging IEEE 802.16 WirelessMAN Standards for Broadband Wireless Access. *Intel Technology journal, Vol. 08, No.03*, 2004.

[11] Q. Lv, P. Cao, E. Cohen, K. Li, and S. Shenker. Search and Replication in Unstructured Peer-to-Peer Networks. In *Proceedings of the 16th annual ACM International Conference on Supercomputing*, 2002.

[12] A. Spyropoulous, K. Psounis, and C. Raghavendra. Single-copy routing in intermittently connected mobile networks. In *IEEE SECON 2004*, 2004.

[13] 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.