

An Evaluation of Availability Latency in Carrier-based Vehicular Ad-Hoc Networks

Shahram
Ghandeharizadeh
Dept of Computer Science
Univ of Southern California
Los Angeles, CA 90089, USA
shahram@usc.edu

Shyam Kapadia
Dept of Computer Science
Univ of Southern California
Los Angeles, CA 90089, USA
kapadia@usc.edu

Bhaskar Krishnamachari
Dept of Computer Science
Dept of Electrical Engineering
Univ of Southern California
Los Angeles, CA 90089, USA
bkrishna@usc.edu

ABSTRACT

On-demand delivery of audio and video clips in peer-to-peer vehicular ad-hoc networks is an emerging area of research. Our target environment uses data carriers, termed zebroids, where a mobile device carries a data item on behalf of a server to a client thereby minimizing its availability latency. In this study, we quantify the variation in availability latency with zebroids as a function of a rich set of parameters such as car density, storage per device, repository size, and replacement policies employed by zebroids. Using analysis and extensive simulations, we gain novel insights into the design of carrier-based systems. Significant improvements in latency can be obtained with zebroids at the cost of a minimal overhead. These improvements occur even in scenarios with lower accuracy in the predictions of the car routes. Two particularly surprising findings are: (1) a naive random replacement policy employed by the zebroids shows competitive performance, and (2) latency improvements obtained with a simplified instantiation of zebroids are found to be robust to changes in the popularity distribution of the data items.

1. INTRODUCTION

Technological advances in areas of storage and wireless communications have now made it feasible to envision on-demand delivery of data items, for e.g., video and audio clips, in vehicular peer-to-peer networks. In prior work, Ghandeharizadeh *et al.* [8] introduce the concept of a Car-to-Car-Peer-to-Peer (C2P2) device equipped in each vehicle. The notable features of a C2P2 include a mass storage device offering hundreds of gigabytes of storage, a fast processor, and several types of networking cards. In this study, we assume each C2P2 is configured with two types of networking cards: 1) a low-bandwidth networking card with a long radio-range in the order of miles that enables a C2P2 device to communicate with a nearby cellular or WiMax base station, and 2) a high-bandwidth networking card with a limited radio-range in the order of hundreds of feet.

The low-bandwidth connection serves as the control plane, enabling C2P2 devices to exchange meta-data with one or more centralized servers. This connection offers bandwidths in the order of tens and hundreds of Kilobits per second. The high bandwidth

connection supports data rates in the order of tens to hundreds of Megabits per second and represents the ad hoc peer to peer network between the vehicles. We term this the data plane and, as suggested by its name, use it to exchange data items between devices. The technical feasibility of such a two-tier architecture is presented in [6], where some back-of-the-envelope calculations are performed to indicate that the bandwidth of the control plane is sufficient for exchange of control information needed for realizing such an application.

In a typical scenario, a C2P2 device presents a passenger with a list of data items,¹ showing both the name of each title and its availability latency. The latter, denoted as δ , is defined as the earliest time at which the client encounters a copy of its requested data item. A data item is available immediately when it resides in the local storage of the C2P2 device serving the request. Due to C2P2 storage constraints, the entire data item repository of interest to a user may not be stored locally. In that case, availability latency is the time from request initiation until when the C2P2 device will encounter another vehicle containing the requested data item.

The availability latency for an item is a function of the current location of the client, its destination and travel path, the mobility model of the C2P2 equipped vehicles, the number of replicas constructed for the different data items, and the placement of data item replicas across the vehicles. A method to improve the availability latency is to employ data carriers which transport a replica of the requested data item from a server car containing it to a client that requested it. We term these data carriers as ‘zebroids’.

Selection of zebroids is facilitated by the two-tiered architecture. The control plane enables centralized information gathering at a dispatcher present at a base-station.² Some examples of control information are currently active requests, travel path of the clients and their destinations, and paths of the other cars. For each client request, the dispatcher may choose a set of z carriers that collaborate to transfer a data item from a server to a client (**z-relay zebroids**). Here, z is the number of zebroids such that $0 \leq z < N$, where N is the total number of cars. Note that $z = 0$ means no carriers, requiring a server to deliver the data item directly to the client. The chosen relay team of z zebroids hand over the data item transitively to one another to arrive at the client, thereby reducing availability latency (see Section 4.1 for details). To increase robustness, the dispatcher may employ multiple relay teams of z -carriers for every request. This may be useful in scenarios where the dispatcher has lower prediction accuracy in the information about the routes of the

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
Copyright 200X ACM X-XXXXX-XX-X/XX/XX ...\$5.00.

¹Without loss of generality, we will use the term data item in this paper with the understanding that a data item can be an audio clip, video clip, generic data document, etc.

²There may be dispatchers deployed at a subset of the base-stations for fault-tolerance and robustness. Dispatchers between base-stations may communicate via the wired infrastructure.

cars. Finally, storage constraints may require a zebroid to evict existing data items from its local storage to accommodate the client requested item.

In this study, we quantify the following main factors that affect availability latency in the presence of zebroids: (i) data item repository size, (ii) car density, (iii) storage capacity per car, (iv) client trip duration, (v) replacement scheme deployed by the zebroids, and (vi) accuracy of the car route predictions. For a significant subset of these factors, we address some key questions pertaining to use of zebroids both via analysis and extensive simulations.

Our main findings are as follows. A naive random replacement policy employed by the zebroids shows competitive performance in terms of availability latency. With such a policy, substantial improvements in latency can be obtained with zebroids at a minimal replacement overhead. In more practical scenarios, where the dispatcher has inaccurate information about the car routes, zebroids continue to provide latency improvements. A surprising result is that changes in popularity of the data items do not impact the latency gains obtained with a simple instantiation of z-relay zebroids called one-instantaneous zebroids (see Section 4.1). This study suggests a number of interesting directions to be pursued to gain better understanding of design of carrier-based systems that improve availability latency.

The rest of this paper is organized as follows. Section 2 provides a brief overview of the related work in the area. Section 3 provides an overview of the terminology along with the factors that impact availability latency in the presence of zebroids. Section 4 describes how the zebroids may be employed. Section 5 provides details of the analysis methodology employed to capture the performance with zebroids. Section 6 describes the details of the simulation environment used for evaluation. Section 7 enlists the key questions examined in this study and answers them via analysis and simulations. Finally, Section 8 presents brief conclusions and future research directions.

2. RELATED WORK

Recently, several novel and important studies such as ZebraNet [9], DakNet [12], Data Mules [13], Message Ferries [18], SWIM [14], and Seek and Focus [15] have analyzed factors impacting intermittently connected networks consisting of data carriers similar in spirit to zebroids. Table 1 provides an overview of these studies. Factors considered by each study are dictated by their assumed environment and target application. A novel characteristic of our study is the impact on availability latency for a given database repository of items. In future work, it may be useful to integrate these diverse studies along with our work under a comprehensive general model/framework that incorporates all possible factors, environmental characteristics, and application requirements.

Below, we provide an overview of the different projects and their target applications. In ZebraNet, data sensed by sensors attached to zebras is collected by humans as they drive by in a vehicle. In DakNet, vehicles are used to transport data between villages and cities using a store and forward mechanism. Message Ferries capture a more generalized scenario where the movement of the ferries can be controlled to carry data from a source node to a destination node. With Data Mules, intermediate carriers that follow a random walk mobility model are used to carry data from static sensors to base-stations. When all nodes move as per this mobility model, end-to-end routing can be performed using a Seek and Focus strategy presented in [15]. With the Shared Wireless Infostation Model (SWIM), data collected by sensors on whales is replicated when two sensors are in the vicinity of each other and ultimately relayed to a small number of static on-shore base-stations when the whales come to the surface.

We now describe how the previous studies are different from

ours, detailing the novel features of our study. None of these studies predict the future movement trajectory of the nodes to accomplish data delivery as we do in our study with zebroids. Moreover, while DakNet, Message Ferries, and Seek and Focus employ end-to-end data delivery, studies such as Data Mules, ZebraNet and SWIM require data from many nodes to be sent to a single (or a few) base-station(s) yielding a many-to-one or many-to-some delivery mode. In our study, the mode of data delivery is any-to-one since a request for a certain data item can be satisfied by any node that stores that item.

Studies such as ZebraNet, Message Ferries, and Seek and Focus seek to optimize the energy usage, especially with sensors. This is not a constraint with a C2P2 environment. Instead, we seek to optimize meeting time or latency to request satisfaction in the presence of storage constraints. This metric is similar to that considered in the seek and focus study which ignores storage constraints per node. Also, even though the Data Mules study considers a storage constraint per node, it does not optimize for latency subject to this constraint. Instead, that study provides the scaling behavior needed in the storage per sensor (and mule) to maintain the data delivery ratio over a certain threshold.

While reliable data delivery is required with DakNet, Seek and Focus, and Message Ferries, storage constraints may cause data loss, preventing Data Mules and ZebraNet from achieving the same. In our study, the notion of reliable data delivery has a different connotation. This is because for each data item request, the client stipulates a deadline, which is the maximum amount of time it is willing to wait to encounter a copy of the requested item. Zebroids are employed to improve the availability latency for each request, increasing the likelihood of the client's deadline being met.

Finally, studies like ZebraNet, DakNet and Data Mules can tolerate high delay as long as the appropriate data is delivered. However, the Seek and Focus and Ferries approach use delay as a metric of interest. Also, with zebroids, SWIM, and ZebraNet, multiple replicas of the data are employed to facilitate reduction in the delivery latency. Replicas are created when two nodes encounter each other. While there are no storage constraints with SWIM, with ZebraNet, excess data is dropped from the sensors using a drop-tail like scheme. Moreover, none of the studies mentioned above use a two-tier architectural framework like ours that facilitates intelligent selection of data carriers and creation of appropriate data replicas to reduce latency for currently active requests. With storage constraints, we also explore a wide variety of replacement policies deployed at the zebroids.

In prior work, Ghandeharizadeh *et al.* [5] examine the performance of replacement policies that employ location demographic information about the areas native to a zebroid in order to perform a replacement. The primary difference from our work is that in their environment data items may be lost from the system.

3. OVERVIEW AND TERMINOLOGY

Table 2 summarizes the notation of the parameters used in the paper. Below we introduce some terminology used in the paper.

Assume a network of N C2P2-equipped cars, each with storage capacity of α bytes. The total storage capacity of the system is $S_T = N \cdot \alpha$. There are T data items in the database, each with size S_i . The frequency of access to data item 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 data item i , denoted by R_i , as:

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

Study	Potentially Mobile Nodes	Mobility Model	Delivery	Energy Efficiency	Optimize Delay	How many copies created?	Storage Constraint
ZebraNet [9]	All	Controlled + species dependent	Many to One	X		Many	X
DakNet [12]	One	Controlled	One to One			One	
Message Ferries [18]	All	Random + Controlled	One to One	X	X	One	
SWIM [14]	Most	Random without predictions	Many to Some		X	Many	
Data Mules [13]	Some	Random without predictions	Many to One			One	X
Seek and Focus [15]	All	Random without predictions	One to One	X	X	One	
Our Work	All	Random with predictions	Any to One		X	One or More	X

Table 1: Related studies on intermittently connected networks.

Database Parameters	
T	Number of data items.
S_i	Size of data item i
f_i	Frequency of access to data item i .
Replication Parameters	
R_i	Normalized frequency of access to data item i
r_i	Number of replicas for data item i
n	Characterizes a particular replication scheme.
$\bar{\delta}_i$	Average availability latency of data item i
δ_{agg}	Aggregate availability latency, $\delta_{agg} = \sum_{j=1}^T \bar{\delta}_j \cdot f_j$
C2P2 System Parameters	
G	Number of cells in the map (2D-torus).
N	Number of C2P2 devices in the system.
α	Storage capacity per C2P2.
γ	Trip duration of the client C2P2.
S_T	Total storage capacity of the C2P2 system, $S_T = N \cdot \alpha$.

Table 2: Terms and their definitions

The exponent n characterizes a particular replication technique. A square-root replication scheme is realized when $n = 0.5$ [3]. This serves as the base-line for comparison with the case when zebroids are deployed. R_i is normalized to a value between 0 and 1. The number of replicas for data item i , denoted as r_i , is: $r_i = \min(N, \max(1, \lfloor \frac{R_i \cdot N \cdot \alpha}{S_i} \rfloor))$. This captures the case when at least one copy of every data item must be present in the ad-hoc network at all times. In cases where a data item may be lost from the ad-hoc network, this equation becomes $r_i = \min(N, \max(0, \lfloor \frac{R_i \cdot N \cdot \alpha}{S_i} \rfloor))$. In this case, a request for the lost data item may need to be satisfied by fetching the item from a remote server.

The availability latency for a data item i , denoted as δ_i , is defined as the earliest time at which a client C2P2 will find the first replica of the item accessible to it. If this condition is not satisfied, then we set δ_i to γ . This indicates that data item i was not available to the client during its journey. Note that since there is at least one replica in the system for every data item i , by setting γ to a large value we ensure that the client's request for any data item i will be satisfied. However, in most practical circumstances γ may not be so large as to find every data item.

We are interested in the availability latency observed across all data items. Hence, we augment the average availability latency for every data item i with its f_i to obtain the following weighted availability latency (δ_{agg}) metric: $\delta_{agg} = \sum_{i=1}^T \bar{\delta}_i \cdot f_i$

We have identified the following critical parameters that affect availability latency in the presence of zebroids: (i) title database size (T), (ii) car density (N), (iii) storage per car (α), (iv) trip du-

ration (γ), (v) replacement scheme employed by the zebroids (see Section 4.2), and (vi) prediction accuracy of the car routes. We study the variation in availability latency as a function of each of these parameters in the presence of zebroids. Next, we present our solution approach describing how zebroids are selected.

4. SOLUTION APPROACH

4.1 Zebroids

When a client references a data item missing from its local storage, the dispatcher identifies all cars with a copy of the data item as servers. Next, the dispatcher obtains the future routes of all cars for a finite time duration equivalent to the maximum time the client is willing to wait for its request to be serviced. Using this information, the dispatcher schedules the quickest delivery path from any of the servers to the client using any other cars as intermediate carriers. Hence, it determines the optimal set of forwarding decisions that will enable the data item to be delivered to the client in the minimum amount of time. Note that the latency along the quickest delivery path that employs a relay team of z zebroids is similar to that obtained with epidemic routing [17] under the assumptions of infinite storage and no interference.

A simple instantiation of z -relay zebroids occurs when $z = 1$ and the transfer of a copy of the requested data item, from a server to a zebroid in its vicinity, occurs at the instant when the client issues a request. We refer to such a zebroid as a **one-instantaneous zebroid**. It should also be noted that in some cases, the dispatcher might have inaccurate information about the routes of the cars. Hence, a zebroid scheduled on the basis of this inaccurate information may not rendezvous with the client for whom it was employed.

The time required to transfer a data item from a server to a zebroid depends on the data item size and the available link bandwidth. With data items of smaller size, it is reasonable to assume that this transfer time is small, especially in the presence of the high bandwidth provided by the data plane. Use of zebroids for data items of larger size remains a future research direction.

Initially, the data item replicas are distributed across the C2P2s using a replication scheme defined by Equation 1. This scheme computes the number of data item replicas as a function of their popularity. It is static since once allocated the number of replicas in the system do not change and no replacements are performed. Hence, this is referred to as the 'no-zebroids' environment. We quantify the performance of the various replacement policies with reference to this base-line that does not employ zebroids.

4.2 Carrier-based Replacement policies

The replacement policies considered in this paper are reactive since a replacement occurs only in response to a request issued for a certain data item. When the local storage of a zebroid is completely occupied, it needs to replace one of its existing items to carry the client requested data item. For this purpose, the zebroid must select an appropriate candidate for eviction. This decision process is analogous to that encountered in operating system paging where the goal is to maximize the cache hit ratio to prevent disk access delay [16]. We present below a list of carrier-based replacement policies employed in our study which are adapted from different page replacement policies.

1. **Least recently used (LRU)** LRU-K [11] maintains a sliding window containing the time stamps of the K^{th} most recent references to data items. During eviction, the data item whose K^{th} most recent reference is furthest in the past is evicted. Here, we consider the case with $K = 1$. Depending on whether the evictions are based on the least recently used data item across all client requests (**lru-global**) or only the individual client's requests (**lru-local**), we consider global or local variants of the LRU policy.
2. **Least frequently used (LFU)** (a) **Local (lfu-local)**: Each C2P2 keeps track of the least frequently used data item within its local repository. During eviction³, this is the candidate replica that is replaced. (b) **Global (lfu-global)**: The dispatcher maintains the frequency of access to the data items based on requests from all clients. When a zebroid contacts the dispatcher for a victim data item, the dispatcher chooses the data item with the lowest frequency of access.
3. **Random policy (random)** In this case, the chosen zebroid evicts a data item replica from its local storage chosen uniformly at random.

The replacement policies incur the following overheads. First, the complexity associated with the implementation of a policy. Second, the bandwidth used to transfer a copy of a data item from a server to the zebroid. Third, the average number of replacements incurred by the zebroids. Note that in the no-zebroids case neither overhead is incurred.

The metrics considered in this study are aggregate availability latency, δ_{agg} , percentage improvement in δ_{agg} with zebroids as compared to the no-zebroids case and average number of replacements incurred per client request which is an indicator of the overhead incurred by zebroids.

Note that the dispatchers with the help of the control plane may ensure that no data item is lost from the system. In other words, at least one replica of every data item is maintained in the ad-hoc network at all times. In such cases, even though a car may meet a requesting client earlier than other servers, if its local storage contains data items with only a single copy in the system, then such a car is not chosen as a zebroid.

5. ANALYSIS METHODOLOGY

Here, we present the analytical evaluation methodology and some approximations as closed-form equations that capture the improvements in availability latency that can be obtained with both one-instantaneous and z-relay zebroids. First, we present some preliminaries of our analysis methodology.

- Let N be the number of cars in the network performing a 2D random walk on a $\sqrt{G} \times \sqrt{G}$ torus. An additional car serves

³The terms eviction and replacement are used interchangeably.

as a client yielding a total of $N + 1$ cars. Such a mobility model has been used widely in the literature [15, 13] chiefly because it is amenable to analysis and provides a baseline against which performance of other mobility models can be compared.

- We assume that all cars start from the stationary distribution and perform independent random walks. Although for sparse density scenarios, the independence assumption does hold, it is no longer valid when N approaches G .
- Let the size of data item repository of interest be T . Also, data item i has r_i replicas. This implies r_i cars, identified as servers, have a copy of this data item when the client requests item i .

All analysis results presented in this section are obtained assuming that the client is willing to wait as long as it takes for its request to be satisfied (unbounded trip duration $\gamma = \infty$). With the random walk mobility model on a 2D-torus, there is a guarantee that as long as there is at least one replica of the requested data item in the network, the client will eventually encounter this replica [2]. Later, we extend our analysis to consider finite trip duration γ .

Consider a scenario where no zebroids are employed. In this case, the expected availability latency for the data item is the expected meeting time of the random walk undertaken by the client with any of the random walks performed by the servers. Aldous *et al.* [2] show that the the meeting time of two random walks in such a setting can be modelled as an exponential distribution with the mean $C = c \cdot G \cdot \log G$, where the constant $c \simeq 0.17$ for $G \geq 25$. The meeting time, or equivalently the availability latency δ_i , for the client requesting data item i is the time till it encounters any of these r_i replicas for the first time. This is also an exponential distribution with the following expected value (note that this formulation is valid only for sparse cases when $G \gg r_i$): $\bar{\delta}_i = \frac{cG \log G}{r_i}$

The aggregate availability latency without employing zebroids is then this expression averaged over all data items, weighted by their frequency of access:

$$\delta_{agg}(no - zeb) = \sum_{i=1}^T \frac{f_i \cdot c \cdot G \cdot \log G}{r_i} = \sum_{i=1}^T \frac{f_i \cdot C}{r_i} \quad (2)$$

5.1 One-instantaneous zebroids

Recall that with one-instantaneous zebroids, for a given request, a new replica is created on a car in the vicinity of the server, provided this car meets the client earlier than any of the r_i servers. Moreover, this replica is spawned at the time step when the client issues the request. Let \bar{N}_i^c be the expected total number of nodes that are in the same cell as any of the r_i servers. Then, we have

$$\bar{N}_i^c = (N - r_i) \cdot \left(1 - \left(1 - \frac{1}{G}\right)^{r_i}\right) \quad (3)$$

In the analytical model, we assume that \bar{N}_i^c new replicas are created, so that the total number of replicas is increased to $r_i + \bar{N}_i^c$. The availability latency is reduced since the client is more likely to meet a replica earlier. The aggregated expected availability latency in the case of one-instantaneous zebroids is then given by,

$$\delta_{agg}(zeb) = \sum_{i=1}^T \frac{f_i \cdot c \cdot G \cdot \log G}{r_i + \bar{N}_i^c} = \sum_{i=1}^T \frac{f_i \cdot C}{r_i + \bar{N}_i^c} \quad (4)$$

Note that in obtaining this expression, for ease of analysis, we have assumed that the new replicas start from random locations in the torus (not necessarily from the same cell as the original r_i

servers). It thus treats all the N_i^c carriers independently, just like the r_i original servers. As we shall show below by comparison with simulations, this approximation provides an upper-bound on the improvements that can be obtained because it results in a lower expected latency at the client.

It should be noted that the procedure listed above will yield a similar latency to that employed by a dispatcher employing one-instantaneous zebroids (see Section 4.1). Since the dispatcher is aware of all future car movements it would only transfer the requested data item on a single zebroid, if it determines that the zebroid will meet the client earlier than any other server. This selected zebroid is included in the N_i^c new replicas.

5.2 z-relay zebroids

To calculate the expected availability latency with z-relay zebroids, we use a coloring problem analog similar to an approach used by Spyropoulos *et al.* [15]. Consider a data item i requested by the client. Recall that, there are N total cars and r_i replicas for data item i . Assume that each of these r_i replicas is colored red, while the other cars including the client are colored blue. Whenever a red car encounters a blue car, the latter is colored red.

The expected number of steps until the client is colored red then gives the average availability latency with z-relay zebroids. If at a given step, there are k red cars ($k \geq r_i$), then there will be $N - k$ blue cars. Recall that movements of the cars can be modelled as exponential distributions. Hence, by the property of exponential distribution, the average time until any of the k red cars meets any of the $N + 1 - k$ blue cars is $\frac{C}{k \cdot (N + 1 - k)}$. Now, the expected time until all the cars are colored red is $\sum_{k=r_i}^N \frac{C}{k \cdot (N + 1 - k)}$

Note that the client may be colored red in any one of these steps with equal probability. Consequently, the expected time till the client is colored red is given by,

$$\bar{\delta}_i = \frac{C}{N + 1 - r_i} \sum_{m=r_i}^N \sum_{k=r_i}^m \frac{1}{k \cdot (N + 1 - k)} \quad (5)$$

Evaluating the above expression, we get,

$$\bar{\delta}_i = \frac{C}{N + 1} \cdot \frac{1}{N + 1 - r_i} \cdot [N \cdot \log \frac{N}{r_i} - \log(N + 1 - r_i)] \quad (6)$$

Now, the aggregate availability latency (δ_{agg}) with z-relay zebroids is obtained by definition,

$$\delta_{agg}(zeb) = \sum_{i=1}^T [f_i \cdot \frac{C}{N + 1} \cdot \frac{1}{N + 1 - r_i} \cdot (N \cdot \log \frac{N}{r_i} - \log(N + 1 - r_i))] \quad (7)$$

6. SIMULATION METHODOLOGY

The simulation environment considered in this study comprises of vehicles such as cars that carry a fraction of the data item repository. A prediction accuracy parameter inherently provides a certain probabilistic guarantee on the confidence of the car route predictions known at the dispatcher. A value of 100% implies that the exact routes of all cars are known at all times. A 70% value for this parameter indicates that the routes predicted for the cars will match the actual ones with probability 0.7. Note that this probability is spread across the car routes for the entire trip duration. We now provide the preliminaries of the simulation study and then describe the parameter settings used in our experiments.

- Similar to the analysis methodology, the map used is a 2D torus. A Markov mobility model representing a unbiased 2D

random walk on the surface of the torus describes the movement of the cars across this torus.

- Each grid/cell is a unique state of this Markov chain. In each time slot, every car makes a transition from a cell to any of its neighboring 8 cells. The transition is a function of the current location of the car and a probability transition matrix $Q = [q_{ij}]$ where q_{ij} is the probability of transition from state i to state j . Only C2P2 equipped cars within the same cell may communicate with each other.
- The parameters γ, δ have been discretized and expressed in terms of the number of time slots.
- A C2P2 device does not maintain more than one replica of a data item. This is because additional replicas occupy storage without providing benefits.
- Either one-instantaneous or z-relay zebroids may be employed per client request for latency improvement.
- Unless otherwise mentioned, the prediction accuracy parameter is assumed to be 100%. This is because this study aims to quantify the effect of a large number of parameters individually on availability latency.

Here, we set the size of every data item, S_i , to be 1. α represents the number of storage slots per C2P2. Each storage slot stores one data item. γ represents the duration of the client's journey in terms of the number of time slots. Hence the possible values of availability latency are between 0 and γ . δ is defined as the number of time slots after which a client C2P2 device will encounter a replica of the data item for the first time. If a replica for the data item requested was encountered by the client in the first cell then we set $\delta = 0$. If $\delta > \gamma$ then we set $\delta = \gamma$ indicating that no copy of the requested data item was encountered by the client during its entire journey. In all our simulations, for illustration we consider a 5×5 2D-torus with γ set to 10. Our experiments indicate that the trends in the results scale to maps of larger size.

We simulated a skewed distribution of access to the T data items that obeys Zipf's law with a mean of 0.27. This distribution is shown to correspond to sale of movie theater tickets in the United States [4]. We employ a replication scheme that allocates replicas for a data item as a function of the square-root of the frequency of access of that item. The square-root replication scheme is shown to have competitive latency performance over a large parameter space [7]. The data item replicas are distributed uniformly across the C2P2 devices. This serves as the base-line no-zebroids case. The square-root scheme also provides the initial replica distribution when zebroids are employed. Note that the replacements performed by the zebroids will cause changes to the data item replica distribution. Requests generated as per the Zipf distribution are issued one at a time. The client car that issues the request is chosen in a round-robin manner. After a maximum period of γ , the latency encountered by this request is recorded.

Initially, all cars are distributed across the map as per the steady-state distribution governed by Q . This initial placement of cars across the map is determined by a random number generator initialized with a seed. All results presented in this section are averages over 10 such seeds each invoking 20,000 requests. Hence, each point in all the presented results is an average of 200,000 requests.

The 95% confidence intervals are determined for all sets of results. These intervals are quite tight for the metrics of latency and replacement overhead, hence, we only present them for the metric that captures the percentage improvement in latency with respect to the no-zebroids case.

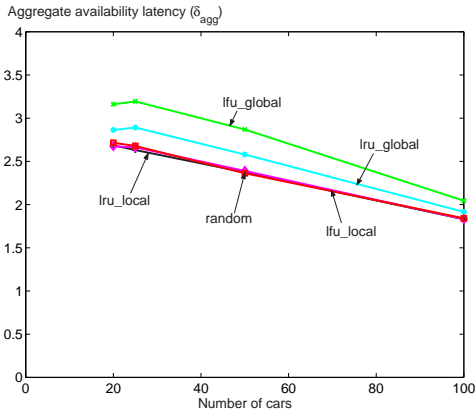


Figure 1: Figure 1 shows the availability latency when employing one-instantaneous zebroids as a function of (N, α) values, when the total storage in the system is kept fixed, $S_T = 200$.

7. RESULTS

In this section, we describe our evaluation results where the following key questions are addressed. With a wide choice of replacement schemes available for a zebroid, what is their effect on availability latency? In fact, a more central question may be: Do zebroids provide significant improvements in availability latency? What is the associated overhead incurred in employing these zebroids? What happens to these improvements in scenarios where a dispatcher may have imperfect information about the car routes? What inherent trade-offs exist between car density and storage per car with regards to their combined as well as individual effect on availability latency in the presence of zebroids? We present both simple analysis and detailed simulations to provide answers to these questions as well as gain insights into design of carrier-based systems.

7.1 How does a replacement scheme employed by a zebroid impact availability latency?

For illustration, we present ‘scale-up’ experiments where one-instantaneous zebroids are employed (see Figure 1). By scale-up, we mean that α and N are changed proportionally to keep the total system storage, S_T , constant. Here, we set $T = 50$ and $S_T = 200$. We choose the following values of $(N, \alpha) = \{(20, 10), (25, 8), (50, 4), (100, 2)\}$. The figure indicates that a random replacement scheme shows a competitive performance. This is because of several reasons.

Recall that the initial replica distribution is set as per the square-root replication scheme. The random replacement scheme does not alter this distribution since it makes replacements blind to the popularity of a data item. However, the replacements cause dynamic data re-organization so as to better serve the currently active request. Moreover, the mobility pattern of the cars is random, hence, the locations from which the requests are issued by clients are also random and not known a priori at the dispatcher. These findings are significant because a random replacement policy can be implemented in a simple decentralized manner.

The lru-global and lfu-global schemes provide a latency performance that is worse than random. This is because these policies rapidly develop a preference for the more popular data items thereby creating a larger number of replicas for them. During eviction, the more popular data items are almost never selected as a replacement candidate. Consequently, there are fewer replicas for the less popular items. Hence, the initial distribution of the data

item replicas changes from square-root to that resembling linear replication. The higher number of replicas for the popular data items provide marginal additional benefits, while the lower number of replicas for the other data items hurts the latency performance of these global policies. The lfu-local and lru-local schemes have similar performance to random since they do not have enough history of local data item requests. We speculate that the performance of these local policies will approach that of their global variants for a large enough history of data item requests. On account of the competitive performance shown by a random policy, for the remainder of the paper, we present the performance of zebroids that employ a random replacement policy.

As part of our future work, it remains to be seen if there are other more sophisticated replacement schemes that may have a performance better than random. Moreover, the distribution of replicas seems to have a profound impact on the availability latency in the presence of zebroids. Exploring and analyzing the cumulative effect of different replica distributions with zebroids presents a promising future research direction. A concrete goal is the investigation of a replacement scheme that over time converges to a replica distribution resembling that provided by a square-root replication scheme.

7.2 Do zebroids provide significant improvements in availability latency?

We find that in many scenarios employing zebroids provides substantial improvements in availability latency.

7.2.1 Analysis

We first consider the case of one-instantaneous zebroids. Figure 2.a shows the variation in δ_{agg} as a function of N for $T = 10$ and $\alpha = 1$ with a 10×10 torus using Equation 4. Both the x and y axes are drawn to a log-scale. Figure 2.b show the % improvement in δ_{agg} obtained with one-instantaneous zebroids. In this case, only the x-axis is drawn to a log-scale. For illustration, we assume that the T data items are requested uniformly.

Initially, when the network is sparse the analytical approximation for improvements in latency with zebroids, obtained from Equations 2 and 4, closely matches the simulation results. However, as N increases, the sparseness assumption for which the analysis is valid, namely $N \ll G$, is no longer true. Hence, the two curves rapidly diverge. The point at which the two curves move away from each other corresponds to a value of $\delta_{agg} \leq 1$. Moreover, as mentioned earlier, the analysis provides an upper bound on the latency improvements, as it treats the newly created replicas given by \bar{N}_i^c independently. However, these \bar{N}_i^c replicas start from the same cell as one of the server replicas r_i . Finally, the analysis captures a one-shot scenario where given an initial data item replica distribution, the availability latency is computed. The new replicas created do not affect future requests from the client.

Next, we consider the case where z-relay zebroids are employed (see Figure 3). Similar observations, like the one-instantaneous zebroid case, apply since the simulation and analysis curves again start diverging when the analysis assumptions are no longer valid. However, the key observation here is that the latency improvement with z-relay zebroids is significantly better than the one-instantaneous zebroids case, especially for lower storage scenarios. This is because in sparse scenarios, the transitive hand-offs between the zebroids creates higher number of replicas for the requested data item, yielding lower availability latency. Moreover, it is also seen that the simulation validation curve for the improvements in δ_{agg} with z-relay zebroids approaches that of the one-instantaneous zebroid case for higher storage (higher N values). This is because one-instantaneous zebroids are a special case of z-relay zebroids.

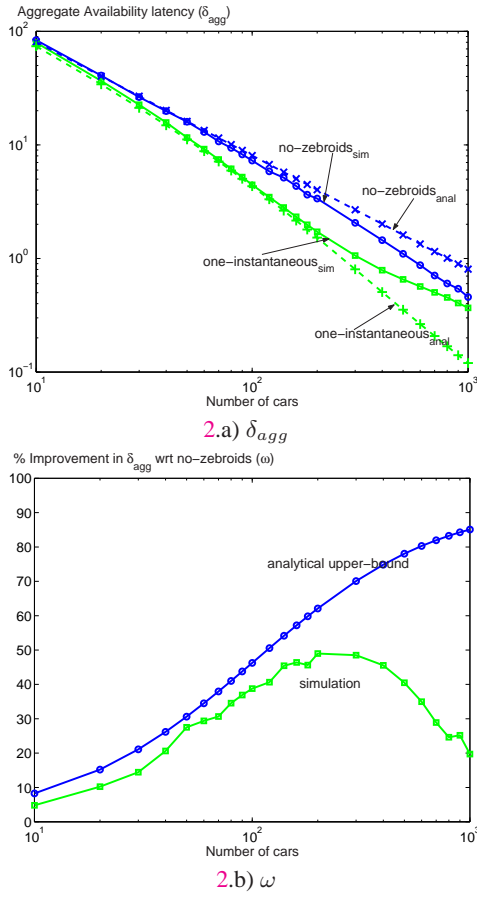


Figure 2: Figure 2 shows the latency performance with one-instantaneous zebroids via simulations along with the analytical approximation for a 10×10 torus with $T = 10$.

7.2.2 Simulation

We conduct simulations to examine the entire storage spectrum obtained by changing car density N or storage per car α in order to also capture scenarios where the sparseness assumptions for which the analysis is valid do not hold. We separate the effect of N and α by capturing the variation of N while keeping α constant (case 1) and vice-versa (case 2) both with z-relay and one-instantaneous zebroids. Here, we set the repository size as $T = 25$. Figure 4 and 5 respectively capture the two cases mentioned above. With Figure 4.b, keeping α constant, initially increasing car density has higher latency benefits because increasing N introduces more zebroids in the system. As N is further increased, ω reduces because the total storage in the system goes up. Consequently, the number of replicas per data item goes up thereby increasing the number of servers. Hence, the replacement policy cannot find a zebroid as often to transport the requested data item to the client earlier than any of the servers. On the other hand, the increased number of servers benefits the no-zebroids case in bringing δ_{agg} down. The net effect results in reduction in ω for larger values of N . Similar trends are seen by keeping N constant and increasing α (see Figure 5.b).

The trends mentioned above are similar to that obtained from the analysis. However, somewhat counter-intuitively with relatively higher system storage, z-relay zebroids provide slightly lower improvements in latency as compared to one-instantaneous zebroids. We speculate that this is due to the different data item replica distributions enforced by them. Note that replacements performed by

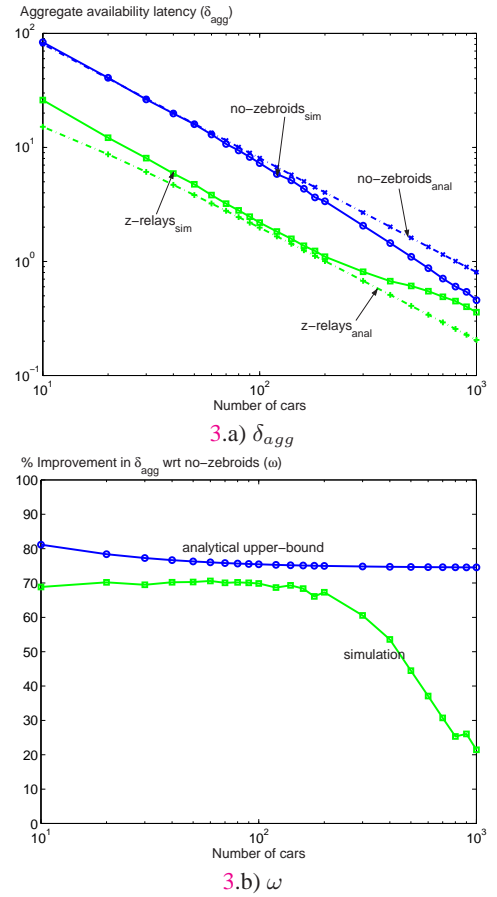


Figure 3: Figure 3 shows the latency performance with z-relay zebroids via analysis and simulations for a 10×10 torus with $T = 10$.

the zebroids cause fluctuations in these replica distributions which may effect future client requests. We are currently exploring suitable choices of parameters that can capture these changing replica distributions.

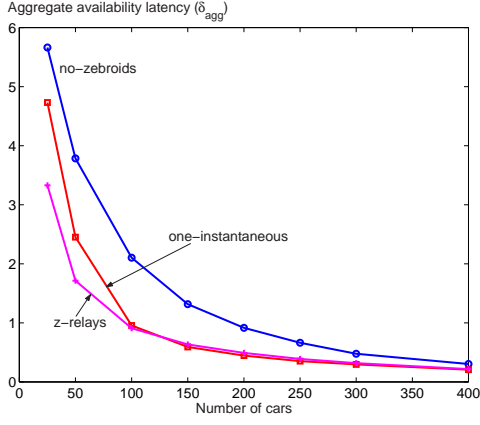
7.3 What is the overhead/cost associated with achieving improvements in latency with zebroids?

We find that the improvements in latency with zebroids are obtained at a minimal replacement overhead (< 1 per client request).

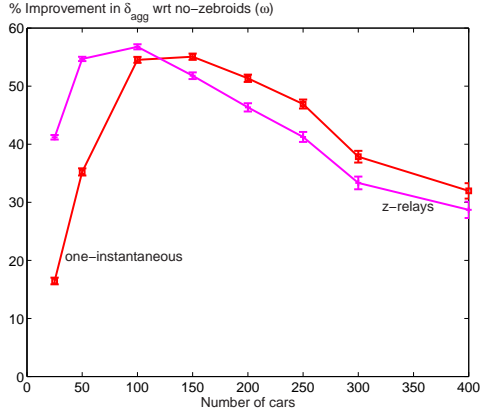
7.3.1 Analysis

With one-instantaneous zebroids, for each client request a maximum of one zebroid is employed for latency improvement. Hence, the replacement overhead per client request can amount to a maximum of one. Recall that to calculate the latency with one-instantaneous zebroids, \bar{N}_i^c new replicas are created in the same cell as the servers. Now a replacement is only incurred if one of these \bar{N}_i^c newly created replicas meets the client earlier than any of the r_i servers.

Let X_{r_i} and $X_{\bar{N}_i^c}$ respectively be random variables that capture the minimum time till any of the r_i and \bar{N}_i^c replicas meet the client. Since X_{r_i} and $X_{\bar{N}_i^c}$ are assumed to be independent, by the property of exponentially distributed random variables we have,



4.a



4.b

Figure 4: Figure 4 depicts the latency performance with both one-instantaneous and z-relay zebroids as a function of the car density when $\alpha = 2$ and $T = 25$.

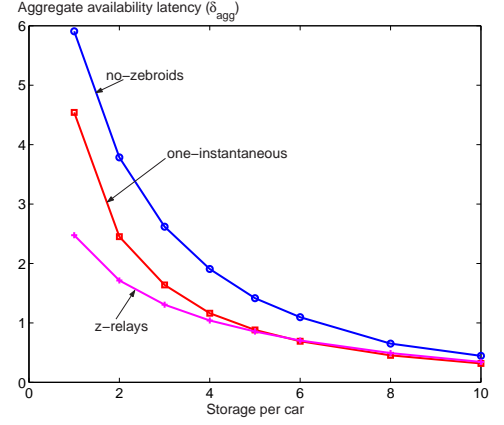
$$\text{Overhead/request} = 1 \cdot P(X_{\overline{N_i^c}} < X_{r_i}) + 0 \cdot P(X_{r_i} \leq X_{\overline{N_i^c}}) \quad (8)$$

$$\text{Overhead/request} = \frac{\frac{r_i}{C}}{\frac{r_i}{C} + \frac{N_i^c}{C}} = \frac{r_i}{r_i + N_i^c} \quad (9)$$

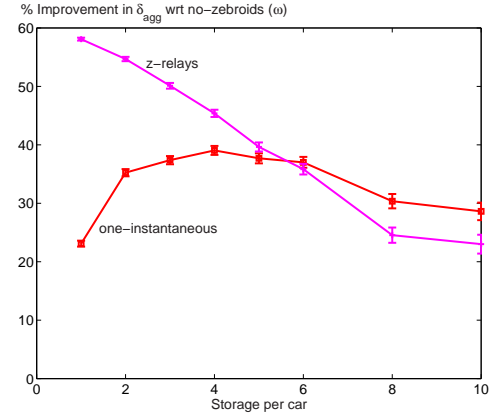
Recall that the number of replicas for data item i , r_i , is a function of the total storage in the system i.e., $r_i = k \cdot N \cdot \alpha$ where k satisfies the constraint $1 \leq r_i \leq N$. Using this along with Equation 2, we get

$$\text{Overhead/request} = 1 - \frac{G}{G + N \cdot (1 - k \cdot \alpha)} \quad (10)$$

Now if we keep the total system storage $N \cdot \alpha$ constant since G and T are also constant, increasing N increases the replacement overhead. However, if $N \cdot \alpha$ is constant then increasing N causes α to go down. This implies that a higher replacement overhead is incurred for higher N and lower α values. Moreover, when $r_i = N$, this means that every car has a replica of data item i . Hence, no zebroids are employed when this item is requested, yielding an overhead/request for this item as zero. Next, we present simulation results that validate our analysis hypothesis for the overhead associated with deployment of one-instantaneous zebroids.



5.a



5.b

Figure 5: Figure 5 shows latency performance with both one-instantaneous and z-relay zebroids as a function of α when $N = 50$ and $T = 25$.

7.3.2 Simulation

Figure 6 shows the replacement overhead with one-instantaneous zebroids when (N, α) are varied while keeping the total system storage constant. The trends shown by the simulation are in agreement with those predicted by the analysis above. However, the total system storage can be changed either by varying car density (N) or storage per car (α). Figures 7.a and Figure 7.b respectively indicate the replacement overhead incurred with both one-instantaneous and z-relay zebroids when α is kept constant and N is varied and vice-versa.

We present an intuitive argument for the behavior of the per-request replacement overhead curves. When the storage is extremely scarce so that only one replica per data item exists in the C2P2 network, the number of replacements performed by the zebroids is zero since any replacement will cause a data item to be lost from the system. The dispatcher ensures that no data item is lost from the system. At the other end of the spectrum, if storage becomes so abundant that $\alpha = T$ then the entire data item repository can be replicated on every car. The number of replacements is again zero since each request can be satisfied locally. A similar scenario occurs if N is increased to such a large value that another car with the requested data item is always available in the vicinity of the client. However, there is a storage spectrum in the middle where replacements by the scheduled zebroids result in improvements in δ_{agg} (see Figures 4.b and 5.b).

Moreover, we observe that for sparse storage scenarios, the higher

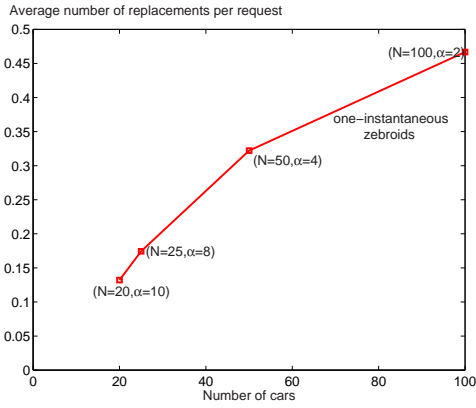


Figure 6: Figure 6 captures replacement overhead when employing one-instantaneous zebroids as a function of (N, α) values, when the total storage in the system is kept fixed, $S_T = 200$.

improvements with z-relay zebroids are obtained at the cost of a higher replacement overhead when compared to the one-instantaneous zebroids case. This is because in the former case, each of these z zebroids selected along the lowest latency path to the client needs to perform a replacement. However, the replacement overhead is still less than 1 indicating that on an average less than one replacement per client request is needed even when z-relay zebroids are employed.

Note that the average replacement per request metric does not explicitly capture the bandwidth overhead associated with the transfer of items to the zebroids. This bandwidth overhead may be significant in the case of multiple simultaneous active requests. We intend to explicitly incorporate these bandwidth considerations in our model as part of our future research.

7.4 What happens to the availability latency with zebroids in more practical scenarios with inaccuracies in the car route predictions?

We find that zebroids continue to provide improvements in availability latency even with lower accuracy in the car route predictions. We use a single parameter p to quantify the accuracy of the car route predictions. This parameter inherently provides a certain probabilistic guarantee on the confidence of the car route predictions for the entire trip duration.

7.4.1 Analysis

Since p represents the probability that a car route predicted by the dispatcher matches the actual one, hence, the latency with zebroids can be approximated by,

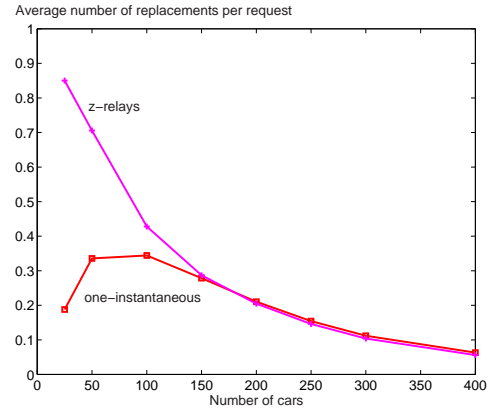
$$\delta_{agg}^{err} = p \cdot \delta_{agg}(zeb) + (1 - p) \cdot \delta_{agg}(no - zeb) \quad (11)$$

$$\delta_{agg}^{err} = p \cdot \delta_{agg}(zeb) + (1 - p) \cdot \frac{C}{r_i} \quad (12)$$

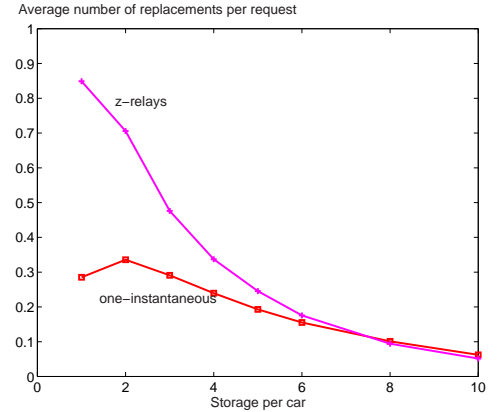
Expressions for $\delta_{agg}(zeb)$ can be obtained from Equations 4 (one-instantaneous) or 7 (z-relay zebroids).

7.4.2 Simulation

Figure 8 shows the variation in δ_{agg} as a function of this route prediction accuracy metric with one-instantaneous zebroids. We



7.a



7.b

Figure 7: Figure 7.a and Figure 7.b respectively show the replacement overhead with zebroids for the cases when N is varied keeping $\alpha = 2$ and α is varied keeping $N = 50$.

observe a smooth reduction in the improvement in δ_{agg} as the prediction accuracy metric reduces. For zebroids that are scheduled but fail to rendezvous with the client due to the prediction error, we tag any such replacements made by the zebroids as failed. It is seen that failed replacements gradually increase as the prediction accuracy reduces.

In this study, we have considered a metric that probabilistically governs errors in the car route predictions. Another possible choice for the metric is similar to that used by Jun *et al.* [10] where the car routes are assumed to follow a Gaussian distribution defined by a mean and a variance. The estimates about the mean and the variance can be built at the dispatcher based on the history of the individual car movements. Exploring such alternate choices of prediction control metrics presents a promising future research direction.

7.5 Under what conditions are the improvements in availability latency with zebroids maximized?

Surprisingly, we find that the improvements in latency obtained with one-instantaneous zebroids are independent of the input distribution of the popularity of the data items.

7.5.1 Analysis

The fractional difference (labelled ω) in the latency between the no-zebroids and one-instantaneous zebroids is obtained from equa-

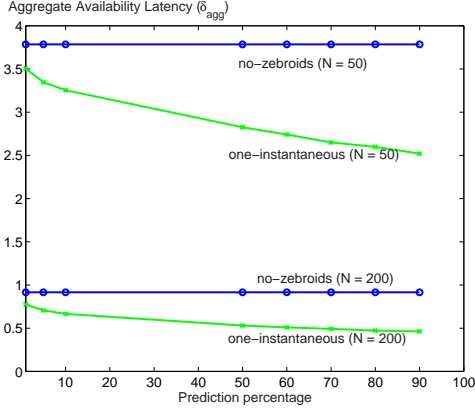


Figure 8: Figure 8 shows δ_{agg} for different car densities as a function of the prediction accuracy metric with $\alpha = 2$ and $T = 25$.

tions 2, 3, and 4 as

$$\omega = \frac{\sum_{i=1}^T \frac{f_i \cdot C}{r_i} - \sum_{i=1}^T \frac{f_i \cdot C}{r_i + (N - r_i) \cdot \left(1 - \left(1 - \frac{1}{G}\right)^{r_i}\right)}}{\sum_{i=1}^T \frac{f_i \cdot C}{r_i}} \quad (13)$$

Here $C = c \cdot G \cdot \log G$. This captures the fractional improvement in the availability latency obtained by employing one-instantaneous zebroids. Let $\alpha = 1$, making the total storage in the system $S_T = N$. Assuming the initial replica distribution is as per the square-root replication scheme, we have, $r_i = \frac{\sqrt{f_i \cdot N}}{\sum_{j=1}^T \sqrt{f_j}}$. Hence, we get $f_i = \frac{K^2 \cdot r_i^2}{N^2}$, where $K = \sum_{j=1}^T \sqrt{f_j}$. Using this, along with the approximation $(1 - x)^n \simeq 1 - n \cdot x$ for small x , we simplify the above equation to get,

$$\omega = 1 - \frac{\sum_{i=1}^T \frac{r_i}{1 + \frac{N - r_i}{G}}}{\sum_{i=1}^T r_i} \quad (14)$$

In order to determine when the gains with one-instantaneous zebroids are maximized, we can frame an optimization problem as follows:

$$\text{Maximize } \omega, \text{ subject to } \sum_{i=1}^T r_i = S_T \quad (15)$$

THEOREM 1. *With a square-root replication scheme, improvements obtained with one-instantaneous zebroids are independent of the input popularity distribution of the data items.*

Proof: Using the Lagrangian multipliers method, the optimization can be expressed as:

$$\text{Max} \left\{ 1 - \sum_{i=1}^T \frac{G \cdot r_i}{N \cdot (G + N - r_i)} + \lambda \cdot \left[\sum_{i=1}^T r_i - N \right] \right\} \quad (16)$$

We solve for r_i to obtain:

$$r_i = G + N - G \cdot \sqrt{\frac{G + N}{G \cdot N \cdot \lambda}} \quad (17)$$

Note that in Equation 17, while r_i is independent of i , it is the same for all titles i . It can be verified that the maximum ω occurs at this value of r_i since $\frac{\partial^2 \omega}{\partial r_i^2} < 0$. This implies that as long as the constraint $\sum_{i=1}^T r_i = N$ is satisfied, improvements in latency can be maximized. \square

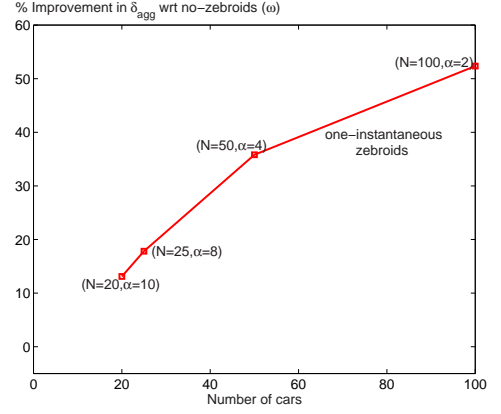


Figure 9: Figure 9 captures improvement in availability latency with one-instantaneous zebroids as a function of (N, α) values, when the total storage in the system is kept fixed, $S_T = 200$.

7.5.2 Simulation

We perform simulations with two different frequency distribution of data items: Uniform and Zipfian (with mean=0.27). Similar latency improvements with one-instantaneous zebroids are obtained in both cases. This result has important implications. In cases with biased popularity toward certain data items, the aggregate improvements in latency across all data item requests still remain the same. Even in scenarios where the frequency of access to the data items changes dynamically, zebroids will continue to provide similar latency improvements.

7.6 What combinations of car-density and storage per car offer the highest improvements in latency with zebroids for the same total system storage?

Our findings indicate that higher latency improvements can be obtained with zebroids when there are more cars with lower storage than fewer ones with higher storage.

7.6.1 Analysis

Consider the case of one-instantaneous zebroids. The fractional difference (labelled ω) in δ_{agg} between the no-zebroids and one-instantaneous zebroids cases is obtained in Equation 13. Using the approximation $(1 - x)^n \simeq 1 - n \cdot x$ for small x , we simplify the above equation to get,

$$\omega = 1 - \frac{\sum_{i=1}^T \frac{G \cdot f_i}{r_i \cdot (G + N - r_i)}}{\sum_{i=1}^T \frac{f_i}{r_i}} \quad (18)$$

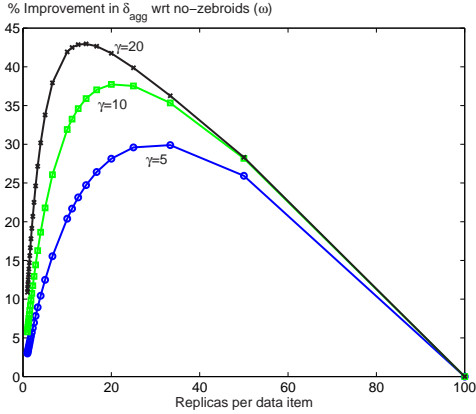
Recall that the number of replicas for data item i , r_i , is a function of the total storage in the system i.e., $r_i = k \cdot N \cdot \alpha$ where k has to satisfy the constraint that $1 \leq r_i \leq N$. Given a total system storage $N \cdot \alpha$, we find that except for N all other terms in Equation 18 are constant. Also, with increasing N , ω increases. However, for a constant $N \cdot \alpha$, if N increases, α has to reduce. This indicates for a given system storage, higher improvements in latency with zebroids are obtained with higher car density and lower storage per car.

7.6.2 Simulation

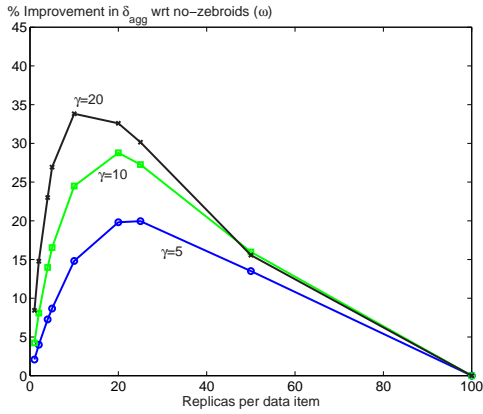
Figure 9 validates the insight obtained from the analysis in that the improvements in latency go up with higher N and lower α values when $N \cdot \alpha = 200$. The increase in N increases the zebroid

density enabling the dispatcher to almost always find a zebroid that can deliver the requested data item to the client earlier than any of the potential servers. This trade-off between the two system parameters of number of cars and storage per car may have important implications in the design of carrier-based networks that improve availability latency.

Although we have assumed a constant storage per car for all cars, in practical scenarios, there may be variable storage per car. Part of the C2P2 device storage space may be reserved by a user for his preferred titles which the user may not seek to erase/evict. These considerations may create a more heterogeneous environment with variable storage per car. In addition, zebroids may need to take into account user preferences prior to making these replacements. Incorporating all these considerations into the model is likely to pay richer dividends in estimating latency in real deployments.



10.a) Analysis



10.b) Simulation

Figure 10: Figures 10.a and 10.b % show improvement in δ_{agg} with one-instantaneous zebroids for different client trip durations in case of 10×10 torus with a fixed car density, $N = 100$.

7.7 What is the impact of different trip durations and repository sizes on availability latency in the presence of zebroids?

7.7.1 Analysis

In most practical scenarios, the client will wait for a maximum duration within which it expects its issued request to be satisfied. Here, we consider the case where the client has a finite trip duration γ , similar to that considered in the simulation environment ($\gamma =$

10). The availability latency, δ_i , can be any value between 0 and $\gamma - 1$. If the client's request is not satisfied, we set $\delta_i = \gamma$ indicating that the client's request for item i was not satisfied.

Recall that latency in the case of a 2D-random walk on a torus can be modelled as an exponential distribution as:

$$P(\delta_i > t) = \lambda \cdot \exp(-\lambda \cdot t) \quad (19)$$

where $\lambda = \frac{r_i}{c \cdot G \cdot \log G}$. The average availability latency with finite trip duration γ is then given by,

$$\bar{\delta}_i = \int_0^\gamma x \cdot \lambda \cdot \exp(-\lambda \cdot t) dx + \int_\gamma^\infty \gamma \cdot \lambda \cdot \exp(-\lambda \cdot t) dx \quad (20)$$

Hence, we get

$$\bar{\delta}_i = \frac{c \cdot G \cdot \log G}{r_i} \cdot [1 - \exp(\frac{-\gamma \cdot r_i}{c \cdot G \cdot \log G})] \quad (21)$$

The aggregate availability latency with finite trip duration is then given by,

$$\delta_{agg}(no-zeb) = \sum_{i=1}^T f_i \cdot \frac{c \cdot G \cdot \log G}{r_i} \cdot [1 - \exp(\frac{-\gamma \cdot r_i}{c \cdot G \cdot \log G})] \quad (22)$$

In the presence of one-instantaneous zebroids, the aggregate availability latency can be obtained using a procedure similar to that used in Section 5.1, giving

$$\delta_{agg}(zeb) = \sum_{i=1}^T f_i \cdot \frac{c \cdot G \cdot \log G}{(r_i + \bar{N}_i^e)} \cdot [1 - \exp(\frac{-\gamma \cdot (r_i + \bar{N}_i^e)}{c \cdot G \cdot \log G})] \quad (23)$$

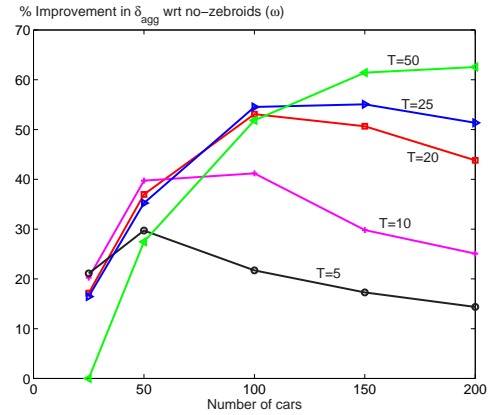


Figure 11: Figure 11 shows improvement in availability latency as a function of the car density for different repository sizes with $\alpha = 2$ and $\gamma = 10$.

The above equations yield the improvements in latency with finite trip duration. We consider a 10×10 torus with $N = 100$ cars each with one storage slot ($\alpha = 1$). The distribution of data item replicas is assumed to be uniform. Hence, as we increase the size of the data item repository (T) the number of replicas per data item decreases. Specifically, the value of T is varied as $\{1, 2, 4, 10, 20, 25, 50, 100\}$. Then, the number of replicas per data item changes as $\{100, 50, 25, 10, 5, 4, 2, 1\}$. Figure 10.a and 10.b capture the latency performance obtained via analysis and simulations respectively for different trip durations.

We now describe the behavior of the curves for a given finite trip duration γ . When $T = 1$, every car has a copy of the item, hence, no car can serve as a zebroid. As T increases, replicas per item go down. Hence, some cars can potentially serve as zebroids, providing some improvements in availability latency. As T is further increased, a peak is reached beyond which the improvements start diminishing. This is because if T is large then the number of server replicas per item becomes small. Hence, the likelihood of finding another car in the vicinity of such a server that will meet the client earlier also reduces. Moreover, as γ increases, peak improvements in latency are obtained with higher T .

7.7.2 Simulation

Here, we present simulation results that capture the effect of different repository sizes on the availability latency when the trip duration $\gamma = 10$ (see Figure 11). For a fixed storage per car, sufficient car density is needed to provide higher improvements in latency for a given repository size. This implies that from a system designer's point of view, if an estimate of the total car density is known, then sufficient gains in latency with zebroids can be realized by adjusting the repository size of titles presented to the users.

While a homogeneous repository of data items has been assumed throughout this study, sizes of data items such as audio clips are typically smaller than video clips. One way in which our model can be extended to consider such a heterogeneous repository is to assume that every data item can be divided into a set of constant-sized blocks. Different blocks of an item may be stored across different cars. During data delivery, zebroids will be scheduled to ensure timely delivery of the various blocks of a requested data item to a client. We intend to consider repositories with different-sized data items as part of our future work.

8. CONCLUSIONS AND FUTURE RESEARCH DIRECTIONS

In this study, we examined the improvements in latency that can be obtained in the presence of carriers that deliver a data item from a server to a client. We quantified the variation in availability latency as a function of a rich set of parameters such as car density, storage per car, title database size, and replacement policies employed by zebroids. Our key findings are summarized below. A naive random replacement policy employed by the zebroids exhibits competitive latency benefits at a minimal replacement overhead. Zebroids continue to provide improvements even in the presence of lower accuracy in the predictions of the car routes. Improvements in latency obtained with one-instantaneous zebroids are independent of the input distribution of the popularity of the data items. Also, for a given total system storage, presence of more cars with lower storage as compared to fewer cars with more storage yields higher improvements in availability latency.

A contribution of this work is to identify a significant research direction for each question presented and discussed in Section 7. Comprehensive analytical models that enable a study of the different factors for each question remain an important future research topic. Below we summarize some key future research directions we intend to pursue. First, we are currently investigating the effect of different replica distributions on availability latency in the presence of zebroids. Second, we intend to explore alternate definitions for the parameter that controls the prediction accuracy of the car routes known at the dispatcher. Third, to better reflect reality we would like to validate the observations obtained from this study with some real world simulation traces of vehicular movements (for example using CORSIM [1]). This will also serve as a validation for the utility of the Markov mobility model used in this study. Fourth, we intend to include bandwidth constraints in our

model so that it can be better equipped for scenarios where multiple simultaneous requests are active in the system.

9. REFERENCES

- [1] Federal Highway Administration. Corridor simulation (corsim/tsis). Version 5.1, <http://www.ops.fhwa.dot.gov/trafficanalysisstools/corsim.htm>.
- [2] D. Aldous and J. Fill. Reversible markov chains and random walks on graphs. Under preparation.
- [3] E. Cohen and S. Shenker. Replication strategies in unstructured peer-to-peer networks. In *SIGCOMM*, pages 177–190. ACM Press, 2002.
- [4] A. Dan, D. Dias, R. Mukherjee, D. Sitaram, and R. Tewari. Buffering and Caching in Large-Scale Video Servers. In *COMPCON*, 1995.
- [5] S. Ghandeharizadeh and S. Kapadia. An Evaluation of Location-Demographic Replacement Policies for Zebroids. In *IEEE CCNC*, 2006.
- [6] S. Ghandeharizadeh, S. Kapadia, and B. Krishnamachari. Pavan: a policy framework for content availability in vehicular ad-hoc networks. In *VANET*, New York, NY, USA, 2004. ACM Press.
- [7] S. Ghandeharizadeh, S. Kapadia, and B. Krishnamachari. Comparison of Replication Strategies for Content Availability in C2P2 networks. In *MDM*, May 2005.
- [8] S. Ghandeharizadeh and B. Krishnamachari. C2p2: A peer-to-peer network for on-demand automobile information services. In *Globe. IEEE*, 2004.
- [9] P. Juang, H. Oki, Y. Wang, M. Martonosi, L. Peh, and D. Rubenstein. Energy-efficient computing for wildlife tracking: design tradeoffs and early experiences with zebnet. *SIGARCH Comput. Archit. News*, 2002.
- [10] H. Jun, M. Ammar, and E. Zegura. Power Management in Delay Tolerant Networks: A Framework and Knowledge-Based Mechanisms. In *SECON*, September 2005.
- [11] E. O'Neil, P. O'Neil, and G. Weikum. The lru-k page replacement algorithm for database disk buffering. In *ACM SIGMOD*, pages 297–306, 1993.
- [12] A. Pentland, R. Fletcher, and A. Hasson. DakNet: Rethinking Connectivity in Developing Nations. *Computer*, 37(1):78–83, 2004.
- [13] R. Shah, S. Roy, S. Jain, and W. Brunette. Data mules: Modeling and analysis of a three-tier architecture for sparse sensor networks. *Elsevier Ad Hoc Networks Journal*, 1, September 2003.
- [14] T. Small and Z. J. Haas. The shared wireless infostation model: a new ad hoc networking paradigm (or where there is a whale, there is a way). In *MobiHoc*, pages 233–244, New York, NY, USA, 2003. ACM Press.
- [15] T. Spyropoulos, K. Psounis, and C. Raghavendra. Single-Copy Routing in Intermittently Connected Mobile Networks. In *SECON*, April 2004.
- [16] A. Tanenbaum. *Modern Operating Systems, 2nd Edition, Chapter 4, Section 4.4*. Prentice Hall, 2001.
- [17] A. Vahdat and D. Becker. Epidemic routing for partially-connected ad hoc networks. Technical report, Department of Computer Science, Duke University, 2000.
- [18] W. Zhao, M. Ammar, and E. Zegura. A message ferrying approach for data delivery in sparse mobile ad hoc networks. In *MobiHoc*, pages 187–198, New York, NY, USA, 2004. ACM Press.