



# Adaptive multiversion data broadcast organizations

Oleg Shigiltchoff<sup>a</sup>, Panos K. Chrysanthis<sup>a,\*</sup>, Evaggelia Pitoura<sup>b</sup>

<sup>a</sup>Department of Computer Science, University of Pittsburgh, Pittsburgh, PA 15260, USA

<sup>b</sup>Department of Computer Science, University of Ioannina, GR 45110 Ioannina, Greece

---

## Abstract

Broadcasting provides an efficient means for disseminating information in both wired and wireless setting. In this paper, we propose a suite of broadcast organization schemes for multiversion data broadcast, i.e., data broadcast in which more than one value is broadcast per data item. Besides increasing the concurrency of client transactions, multiversion broadcast provides clients with the possibility of accessing multiple server states. For example, such a functionality is essential to support applications that require access to data sequences and have limited local memory to store the previous versions, such as in the case of sensor networks. We identify two basic multiversion organizations, namely vertical and horizontal broadcasts and propose an efficient compression scheme applicable to both. We also consider a multiversion client data cache and introduce appropriate cache replacement techniques. Finally, we propose an adaptive scheme that dynamically selects the appropriate broadcast organization based on the client access pattern. We provide performance evaluation results for both flat and broadcast disk organizations and for a variety of client query patterns.

© 2003 Elsevier Ltd. All rights reserved.

*Keywords:* Mobile computing; Data dissemination; Multiversion data; Wireless communications; Broadcast organization; Adaptive broadcast

---

## 1. Introduction

A major problem on the Internet is the scalable dissemination of information. This problem is particularly acute with the presences of mobile users. The traditional unicast pull framework simply does not scale up and it is not suitable for mobile devices due to their inherent resource limitations including power of mobile devices and the capacity of the wireless links. A solution

to this scalability problem is to use multicast communication for both wireline and wireless devices. In particular, in the context of wireless and mobile environments, *broadcast push* [1] takes both the communication and energy limitations into account, exploiting the asymmetry in wireless communication and the reduced energy consumption in the receiving mode. Servers have both much larger bandwidth available than client devices and more power to transmit large amounts of data.

In broadcast push, the server repeatedly sends information to a client population without explicit client requests. Clients monitor the broadcast channel and retrieve the data items they need as

---

\*Corresponding author.

*E-mail addresses:* oleg@cs.pitt.edu (O. Shigiltchoff),  
panos@cs.pitt.edu (P.K. Chrysanthis),  
pitoura@cs.uoi.gr (E. Pitoura).

they arrive on the broadcast channel. Applications of broadcast push typically involve a small number of servers and a much larger number of clients with similar interests. Examples include stock trading, electronic commerce applications, such as auctions and electronic trading, and traffic control information systems. Any number of clients can monitor the broadcast channel. If broadcast data are organized according to the clients' interests, such a scheme makes an effective use of the low-wireless bandwidth. Besides wireless settings, where there is physical support for broadcast, broadcast push provides a scalable means to disseminate data in web environments as well.

To reduce access time, clients may cache data of interest locally. However, in the case of data updates, the problem arises of keeping the data cached in the client consistent with the updated data on the server [2–4]. The same problem also exists in the context of broadcast push, even without client caching. Broadcasting is a form of a cache “on the air”.

In our previous work, we proposed maintaining multiple versions of data items on the broadcast channel as well as in the client cache for concurrency control purposes [5] (similarly to multiversion schemes in traditional unicast pull models, e.g., [6]). With multiple versions, more client transactions read consistent data (i.e., data values that belong to the same server database state) and complete their operation successfully. The time overhead induced by the multiple versions is smaller than the overall time lost for aborts and subsequent recoveries in the absence of multiple versions [5]. Furthermore, multiple versions increase client's tolerance to network disconnections that are common in wireless communications [7].

Besides increasing the concurrency of client transactions and their tolerance to disconnections, multiversion broadcast provides clients with the possibility of accessing multiple server states. For example, such a functionality is essential to support applications that require access to data sequences and have limited local memory to store the previous versions as is the case with data streams.

The performance (characterized by the client access time and power consumption) of multiversion broadcast is directly related to the way data are organized in the broadcast. The main contributions of this paper are:

1. The introduction of two different basic broadcast organizations for multiversion broadcast, namely *Vertical* and *Horizontal*.
2. The development of a compression scheme along the lines of *run length encoding* (RLE) [8], applicable to both of the proposed multiversion broadcast organizations. Our compression scheme exploits the fact that adjacent versions of an item may have the same value and reduces the broadcast size by not explicitly sending unchanged parts of older versions.
3. The proposal of different cache replacement policies for multiversion broadcast that consider versions and apply our compression scheme to maximize the number of hits.
4. The evaluation of the conditions under which each of the proposed broadcast organizations performs better, and the introduction of an adaptive broadcast organization scheme that improves performance in dynamically changing environments.

The remainder of this paper is structured as follows. In Section 2, we present the system model. Section 3 describes different multiversion broadcast organizations for both single and multiple disk organizations as well as our compression scheme. Section 4 evaluates client side cache. Adaptive broadcast is introduced in Section 5. Section 6 presents our experimental platform, while our experimental results are discussed in Section 7. Section 8 discusses related work, and Section 9 concludes the paper.

## 2. System model

In a broadcast dissemination environment, a data server periodically broadcasts data items to a large client population. Each period of the broadcast is called a *broadcast cycle* or *bcycle*, while the content of the broadcast is called a *bcast*. Each

client listens to the broadcast and fetches data as they arrive. In this way, data can be accessed concurrently by any number of clients without any performance degradation (compared to a “pull” or on-demand approach). However, access to data is strictly sequential, since clients need to wait for the data of interest to appear on the channel. We assume that all updates are performed at the server and disseminated from there.

Without loss of generality, in this paper, we consider a model in which the bcast disseminates a fixed number of data items. However, the data values (values of the data items) may or may not change between two consecutive bicycles. In multiversion broadcast, the server maintains multiple versions of each data item and broadcasts a fixed number of versions for each data item. At each cycle, the oldest version of the data is discarded and a new, the most recent, version is included in the broadcast. The number  $m$  of older versions that are retained can be seen as a property of the server: an  $m$ -multiversion server is a server that broadcasts  $m$  values of each item. In terms of client transaction consistency, an  $m$ -multiversion server guarantees the consistency of all transactions with span  $m$  or smaller [5]. The *span* of a client transaction  $T$ , is defined to be the maximum number of different bicycles from which  $T$  reads data.

The client listens to the broadcast and searches for data elements based on the pair of values (data id and version number). Clients do not need to listen to the broadcast continuously. Instead, they tune-in to read specific items. Such selective tuning is important especially in the case of portable mobile computers, since they most often rely for their operation on the finite energy provided by batteries and listening to the broadcast consumes energy. Indexing has been used to support selective tuning and reduce power consumption, often at the cost of access time [37]. In this paper, we focus only on broadcast organization and how to reduce its size without adopting any indexing scheme.

The logical unit of a broadcast is called a *bucket*. Buckets are the analog to blocks for disks. Each bucket has a header that includes useful information. The exact content of the bucket header depends on the specific broadcast organization.

Information in the header usually includes the position of the bucket in the bcast as an offset time-step from the beginning of the broadcast as well as the offset to the beginning of the next broadcast.

The main question in multiversion broadcast is how to organize the broadcast, that is, where to place the new and the old versions. In the next section, we elaborate on this issue, considering in addition broadcast compression as a method to reduce the size of the broadcast.

### 3. Multiversion broadcast organization

#### 3.1. Basic organization

A set of multiversion data to be broadcast can be represented as a two-dimensional array, where dimensions correspond to broadcast version numbers ( $Vno$ ) and data ids ( $Did$ ), and the array elements are the data values ( $Dval$ ) of the items. That is,  $Dval[i, k] = v$  means that the  $k$ -version of the  $i$ -data item is equal to  $v$ . Data items can appear in any order. Versions appear in descending order with the most recent version appearing in the left most column and the oldest version in the right most column. This data representation can be extended to any number of data items and versions.

A simple sequential broadcast can be generated by linearizing the two-dimensional array in two different ways: horizontally or vertically. In *Horizontal broadcast*, a server broadcasts all versions (with different  $Vno$ ) of a data item with a particular  $Did$ , then all versions (with different  $Vno$ ) of the next data item with the next  $Did$  and so on. In *Vertical broadcast*, a server broadcasts all data items (with different  $Did$ ) having a particular  $Vno$ , then all data items (with different  $Did$ ) having the next  $Vno$  and so on. Formally, the Horizontal broadcast transmits  $[Did[Vno, Dval]^*]^*$  sequences, whereas the Vertical broadcast transmits  $[Vno[Did, Dval]^*]^*$  sequences.

Consider for example, the two-dimensional array of Fig. 1 representing four data items selected to be broadcast with  $Did = \{3, 5, 8, 9\}$ , each having four versions with  $Vno = 3$  being the

	Vno = 3	Vno = 2	Vno = 1	Vno = 0
Did = 3	Dval = 1	Dval = 1	Dval = 1	Dval = 1
Did = 5	Dval = 8	Dval = 8	Dval = 8	Dval = 5
Did = 8	Dval = 6	Dval = 1	Dval = 1	Dval = 2
Did = 9	Dval = 5	Dval = 4	Dval = 4	Dval = 4

Fig. 1. Example of an array representing a sequence of four versions of four data items.

most recent one. For the Horizontal broadcast, the data values appear on the bcst in the following order (the complete bcst will include also the data ids and version numbers as indicated above):

**1 1 1 1 8 8 8 5 6 1 1 2 5 4 4 4**

while for the Vertical broadcast, the data values on the bcst are placed in the following order:

**1 8 6 5 1 8 1 4 1 8 1 4 1 5 2 4**

The resulting bcsts have the same size for both organizations, but differ in the order in which they broadcast the data values.

The broadcast order affects the client access time. For instance, say a client who tunes in at the beginning of our example bcst wants to access  $Vno = 3$  of data element with  $Did = 9$ . In the Vertical broadcast, the client waits time equal to the time for broadcasting 3 data items before the item of interest appears in the bcst, while in the Horizontal broadcast, it waits time equal to the time for broadcasting 12 data items. We expect that different strategies are more appropriate for different applications. If users require different versions of a particular data (for example, the history of a stock index change), Horizontal broadcast is preferable. If users need the most recent data (for example, the current stock indexes), Vertical broadcast is expected to be more efficient. Our experimental results in Section 7 verify this intuition.

### 3.2. Compressed organization

In both cases, Horizontal and Vertical, the broadcast size and consequently the access time can be reduced by using some compression scheme. A good compression scheme should reduce the broadcast as much as possible with minimal, if any, impact on the client. That is, it

should not require additional processing at the client, so it should not trade access time for processing time. The following is a simple compression scheme that exhibits the above properties.

Our compression scheme is based on the observation that data values do not always change from one version to another. In other words,  $Dval[i, k] = Dval[i, k - 1] = \dots = Dval[i, k - N] = v$ , where the value of the item having  $Did = i$  remains equal to  $v$  for  $N$  consequent versions. Then, when broadcasting data, there is no reason to broadcast all versions of a data item if its  $Dval$  does not change. Instead, the compressed scheme broadcasts a  $Dval$  only if it is different from the  $Dval$  of the previous version. In order not to lose information (as well as to support selective tuning), it also broadcasts the number of versions having the same  $Dval$ .

The generation of a compressed broadcast proceeds in two steps. In the first step, the compressed values of the data elements are produced and in the second phase, these values are broadcast based on the selected organization. In the first step, any sequence of elements with repetitive data values is replaced by the first element of the sequence and the length of the sequence.

The array of Fig. 1 is redrawn in Fig. 2 to capture the first step of the compression algorithm. For example, 1 1 1 1 becomes  $1 \times 3$  and 8 8 8 becomes  $8 \times 2$ .

In the second step of the compression, the Horizontal and Vertical broadcasts are produced by using horizontal and vertical linearization of the arrays, respectively. Formally, the Horizontal broadcast produces  $[Did[Vno(Dval \times Repetition)]^*]^*$  sequences, while the Vertical broadcast creates  $[Vno[Did(Dval \times Repetition)]^*]^*$  sequences. Obviously, we do not include into the broadcast those versions, which already have been included “implicitly” with other versions.

Did = 3	1 x 3		
Did = 5	8 x 2		5
Did = 8	6	1 x 1	2
Did = 9	5	4 x 2	

Fig. 2. The example array after the first step of the compression.

In the example of Fig. 2, the compressed Horizontal broadcast produced in the second step is

$$1 \times 3 \ 8 \times 2 \ 5 \ 6 \ 1 \times 1 \ 2 \ 5 \ 4 \times 2$$

and the compressed Vertical

$$1 \times 3 \ 8 \times 2 \ 6 \ 5 \ 1 \times 1 \ 4 \times 2 \ 5 \ 2.$$

In the case of the Vertical broadcast, it also makes sense to rearrange the sequence of broadcast data elements within a single-version sweep and make them dependent not on *Did* but on the number of implicitly broadcast elements. Applying this reordering to our example, the resulting Vertical broadcast is

$$1 \times 3 \ 8 \times 2 \ 6 \ 5 \ 4 \times 2 \ 1 \times 1 \ 5 \ 2.$$

In this example,  $4 \times 2$  and  $1 \times 1$  switch their positions, because we broadcast implicitly two 4's and only one 1. The idea is that we broadcast first as “dense” data as possible, because when a client begins to read the broadcast, the client has higher chances to find the necessary data elements in “more dense” data. Of course, this optimization works under the assumption that client access data uniformly.

For providing a rough estimation of the reduction of the broadcast size (and, consequently, of the duration of the bicycle) due to our compression scheme, let us introduce *Randomness* and *Randomness degree* to represent the repetitiveness of data from one version to the other. Let  $Randomness[i, k] = 0$  if  $Dval[i, k] = Dval[i, k - 1]$  and  $Randomness[i, k] = 1$  otherwise. Then, the average randomness  $Randomness\ degree[i]$  of the data element *i* over all its versions represents how frequently the value of this element changes. We can use this average randomness as a parameter describing the probability that  $Dval[i, m]$  is not equal to  $Dval[i, m - 1]$ . For instance,  $Randomness\ degree[i] = 0$  means that  $Dval[i, m]$

$= Dval[i, m - 1]$ , for any *m*. The smaller the degree of randomness, the higher the gain of our compression scheme. Hence, we can expect that the broadcast of data having many “static” elements (for example, a cartoon clip with one-color background or a stock index of infrequently traded companies, etc.) improves the “density” of broadcast data.

As an example, consider the broadcast of data item *j* assuming  $Randomness\ degree[j] = 0.1$ . Then, in average out of 100 versions we have 10 versions with values different from the values of the previous versions and 90 versions repeating their values. This means that instead of broadcasting 100 data values, we broadcast only 10. We can roughly estimate that the overhead created by the auxiliary symbols will not exceed 1 symbol per “saved” data item from the broadcast. Assuming, one data item consumes 16 bytes and one auxiliary symbol consumes 1 byte, the gain is  $100 \times 16 / (10 \times 16 + 90 \times 1) = 6.4$ , which corresponds to a 84% reduction of the broadcast length. Similarly, the broadcast shrinks about 45%, for  $Randomness\ degree = 0.5$  and about 9%, for  $Randomness\ degree = 0.9$ . This reduction holds for both Vertical and Horizontal broadcast.

### 3.3. Broadcast encoding

In order to make our broadcast fully self-descriptive, we add all necessary information about version numbers and data items. One of our design principles has been to make the system flexible, allowing a client to understand the content of a broadcast without requiring the client explicitly to be told of the organization of the broadcast. For this purpose, we use four auxiliary symbols:

$$\# (Did), \ \mathbf{V} (Vno), \\ = (\text{Assignment to } Dval), \ \wedge (\text{Number of repetitions})$$



Using these symbols, the *uncompressed bcst* for Horizontal broadcast in our example is fully encoded as

$$\begin{aligned} V3\#3 &= 1V2\#3 = 1V1\#3 = 1V0\#3 = 1V3\#5 \\ &= 8V2\#5 = 8V1\#5 = 8V0\#5 = 5 \\ V3\#8 &= 6V2\#8 = 1V1\#8 = 1V0\#8 = 2V3\#9 \\ &= 5V2\#9 = 4V1\#9 = 4V0\#9 = 4 \end{aligned}$$

and for Vertical broadcast

$$\begin{aligned} V3\#3 &= 1\#5 = 8\#8 = 6\#9 = 5V2\#3 = 1\#5 \\ &= 8\#8 = 1\#9 = 4V1\#3 = 1\#5 = 8\#8 \\ &= 1\#9 = 4V0\#3 = 1\#5 = 5\#8 = 2\#9 = 4 \end{aligned}$$

$V3$ ,  $V2$ ,  $V1$  and  $V0$  are the version numbers. They determine  $Vno$  of the data elements which follow it in the broadcast.  $\#3=1$  means the element having  $Did=3$  of the corresponding version (broadcast before) is equal to 1. So,  $V3\#3=1\#5=8$  means  $Dval[Did=3, Vno=3]=1$  and  $Dval[Did=5, Vno=3]=8$ . Note that for Vertical broadcast we do not need to include the version number in the broadcast before each data element, but for Horizontal broadcast we have to do this. Because of this need of some extra auxiliary symbols, a Horizontal broadcast is usually longer than its corresponding Vertical broadcast. However, given that the size of an auxiliary symbol is much smaller (which is typically the case) than the size of a data element, this difference in length becomes very small.

In the case of *compressed bcst*, the symbol  $\wedge$  is used to specify that the following versions of a data item have the same value. The other auxiliary symbols are also used to give a client the complete information about  $Did$ ,  $Vno$ , and  $Dval$  in a uniform format for both the compressed and uncompressed multiversion broadcast organizations. Returning to our example broadcasts, the compressed Horizontal broadcast is encoded as

$$\begin{aligned} V3\wedge 3\#3 &= 1V2V1V0V3\wedge 2\#5 = 8V2V1V0\wedge 0\#8 \\ &= 5V3\wedge 0\#8 = 6V2\wedge 1\#8 = 1V1 \\ V0\wedge 0\#8 &= 2V3\wedge 0\#9 = 5V2\wedge 2\#9 = 4V1V0 \end{aligned}$$

whereas the compressed Vertical broadcast as:

$$\begin{aligned} V3\wedge 3\#3 &= 1\wedge 2\#5 = 8\wedge 0\#8 = 6\#9 = 5V2\wedge 2\#9 \\ &= 4\wedge 1\#8 = 1V1V0\wedge 0\#5 = 5\#8 = 2 \end{aligned}$$

Considering the Vertical bcst as an example, let us clarify some details of the broadcast. It starts from the version 0. First, it broadcasts the data elements with the most repetitive versions.  $V3\wedge 3\#3=1\wedge 2\#5=8\wedge 0\#8=6\#9=5$  means that versions 3, 2, 1, 0 of data element 3 are 1, versions 3, 2, 1 of data element 5 are 8, version 3 of data element 8 is 6, version 3 of data element 9 is 5.  $V2\wedge 2\#9=4\wedge 1\#8=1$  means that versions 2, 1, 0 of data element 9 are 4 and versions 2, 1 of data element 8 are 1. We do not broadcast versions 2 of data elements 3 and 5 because we broadcast them together with versions 3.

### 3.4. Broadcast disk organization

In this section, before we describe how horizontal and vertical organizations can be used with broadcast disks, we briefly review the basic idea of broadcast disks using an example; for a complete definition of broadcast disks refer to [9,35]. In a broadcast disk organization, the items of the broadcast are divided in ranges of similar access probabilities. Each of these ranges is placed on a separate disk. In the example of Fig. 3, the buckets of the first disk  $Disk_1$  are broadcast three times as often as those in the second disk  $Disk_2$ . To achieve these relative frequencies, the disks are split into smaller equal sized units called chunks; the number of chunks per disk is inversely proportional to the relative frequencies of the disks. In the example, the number of chunks is 1 (chunk 1) and 3 (chunks 2a, 2b and 2c) for  $Disk_1$  and  $Disk_2$ , respectively. Each bcst is generated by broadcasting one chunk from each disk and cycling through all the chunks sequentially over all disks. A minor cycle is a sub-cycle that consists of one chunk from each disk. In the example of Fig. 3, there are three minor cycles.

A direct application of the broadcast disks on multiversion broadcast is to base the distribution of each  $Dval[i, k]$  on its access probability. Then, the  $Dval[i, k]$  sequences inside each bucket or inside each minor cycle can follow either a Horizontal or a Vertical organization. However, this approach requires knowing the access probabilities of each version of each data item. This assumption is rather unrealistic in many settings.

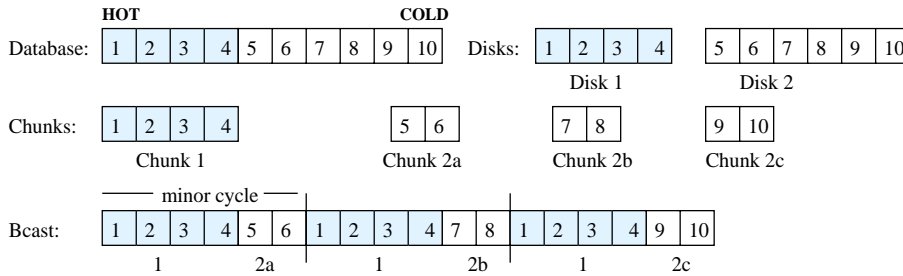


Fig. 3. Broadcast disk organization.

We describe next three “heuristic” approaches that differ on the amount of information about the client access pattern that is required at the server side.

Assume for simplicity that we have only two disks: a “fast” and a “slow” disk. One approach is to determine the *Dids* of the frequently accessed data elements and place all versions of these data items on the “fast” disk. All remaining data items are placed on the “slow” disk. Within each disk, data are organized either horizontally or vertically. This first approach is easily extended to multiple disks.

The second approach builds on the assumption that the most recent versions of items are usually the most frequently requested ones and places the most recent version of all data items on the “fast” disk. Clearly, this means that the data elements on the “fast” disk are vertically organized since they all belong to the same version. All previous versions of the data items are placed in the “slow” disk and can be organized either horizontally or vertically. This second approach can also be generalized for multiple disks, by associating different “temperatures” (degrees of popularity) to different versions. For example, in the case of a 7-multiversion server and three disks, “fast”, “medium-speed” and “slow”, the most recent versions of all data items are placed on the “fast” disk, the previous two most recent versions of all the data items are placed on the “medium-speed” disk and the remaining four oldest versions of all the data items are placed on the “slow” disk.

A third approach, that in some sense combines the above two ones, again uses a third “medium-speed” disk having speed between the “fast” and

the “slow” disks. The approach distinguishes the *Did* of the data items based on their popularity. It places the most recent version of the hot (i.e., most frequently accessed) items on the “fast” disk, using a vertical organization. The other versions of the most frequently accessed items are placed on the “medium-speed” disk along with the most recent version of the remaining cold data items. Finally, the remaining versions of the cold data items are placed on the “slow” disk. Again this approach can be easily generalized for multiple disks and both horizontal and vertical organization can be adopted for each broadcast disk.

#### 4. Client side cache

To improve performance, clients may cache items of interest locally. Caching reduces the latency in answering queries and the need to access the broadcast channel for every data item. In this section, we describe how caching can be used in conjunction with multiversion broadcast.

The simplest way to integrate caching with multiversion broadcast is to assume that data elements are the units of caching and adopt the traditional *least recently used (LRU)* replacement method to discard elements from the cache when the cache runs out of space. Each cache entry has three basic fields for each multiversion data element: *Did*, *Vno*, and *Dval* and a fourth field: *t*, which shows at which broadcast cycle the element was used the last time. Following the LRU policy, when the cache becomes full, the cache element with the oldest (i.e., smallest) *t* is selected for replacement.

Compression can be applied to the caching content as well. Our compression scheme requires that each cache entry has one additional field  $f$ , which contains information on how many adjacent versions of this data element have the same  $Dval$ . Thus, each entry in the compressed LRU scheme is of the form  $[Did, Vno, Dval, t, f]$ . For example, a cache entry  $[i, k, u, t, 2]$  indicates that the  $k$ ,  $k - 1$ , and  $k - 2$  versions of item  $i$  have the same value  $u$ . In the case of compressed broadcast, data elements can be directly cached. In the case of uncompressed broadcast, data elements need to be combined with the corresponding already cached elements.

In [10], we considered a number of variations of compressed-LRU. Our results showed that the compressed-LRU variant with autoprefetch exhibits the best performance with respect to client's perceived latency. With autoprefetching, a client acts pro-actively by prefetching newest version of data that have already been cached. This is a reasonable choice, when energy is not a major issue, for instance, in the case of stationary wireless or docked mobile devices. With autoprefetching, when the client reads the broadcast and it appears that the newest version of a data element has the same  $Dval$  as the newest cached value for this element, besides the  $Vno$  field, the  $f$  field of the corresponding cache entry is incremented as well indicating that there is now a longer sequence of versions with the same  $Dval$ . Autoprefetching increases the effective number of elements stored in the cache.

In this paper, we assume compressed-LRU with autoprefetch but consider different rules for cache replacement. There is a number of different criteria on which to decide which cache entry should be replaced. We suggest using the values of the following cache fields as the basis for replacement:

1.  $t$ : This corresponds to the standard LRU replacement strategy.
2.  $Vno$ : This approach flushes the oldest cache elements on the premise that they will not be of interest to the client.
3.  $f$ : This approach tries to keep in the cache the entries with the most compressed data, effec-

tively, storing in the cache as much data as possible.

We call the replacement strategy which uses the  $t$  (100)-replacement (this is equivalent to LRU replacement), the strategy which uses the  $Vno$  field (010)-replacement, and the strategy which uses the  $f$  field (001)-replacement. It is also possible to combine the three strategies. To this end, we define  $Score$  as

$$Score = A * t + B * Vno + C * f$$

The lower the score, the higher the probability that cache entry will be replaced. Coefficients A, B, and C show which of the three parameter is more significant and define its weight in selecting the item to be replaced. The (ABC) abbreviation shows the weight of each replacement mechanism. In this paper, we consider (100)-replacement, (010)-replacement, (001)-replacement as well as the "aggregated" (111)-replacement.

## 5. Dynamic model of the adaptive broadcast

### 5.1. Client access behavior

Clients have different interests. We consider three types of access behavior. In *historical* access, clients are interested in the "history" of a particular data item, that is, they are interested in reading various versions of this item. For instance, a client may wish to monitor the progress of the values of some data item over a period of time (e.g., a specific stock price, the weather condition at a given city, the output of a particular sensor). In *snapshot* access, clients are interested in a specific version of a set of data items, for instance, they are interested in the current versions of the items or in their values at a specific time instance. For example, clients may be interested in weather conditions in a set of cities in some earlier time, the oldest available value for a number of stock prices, or the current values of a set of sensors. Finally, in the *browse* type of access, a client is interested in a randomly chosen version of a randomly chosen data item.



We also distinguish between *dynamic* and *static* access. In the case of static access, the client knows all data it wants to access before the broadcast. In contrast, in dynamic access, the client determines the next data to access after it reads the previous item.

Our experimental results show that the Horizontal organization is more appropriate for the historical type of access, whereas the Vertical broadcast organization is more appropriate for the snapshot one. Thus, the choice between Horizontal and Vertical broadcast should be based on the access type of the majority of the clients. However, this may change over time. To this end, we introduce adaptive broadcast.

### 5.2. Adaptive broadcast

In adaptive broadcast, the type of broadcast organization (i.e., Vertical or Horizontal) is dynamically determined by the access behavior of the majority of the clients. For adaptive broadcast to be supported, information about the type of client behavior must be available at the server. This information is minimal, just the client's access type: historical, or snapshot. The browse type of access is not affected by the type of the broadcast organization, thus clients with such behavior are not considered in determining the broadcast organization type. To convey their access type, clients send messages via the uplink channel to the server specifying their access type when this changes. The server keeps information reflecting the interest of the majority of its clients. At the beginning of each broadcast cycle, if the server

notices a switch in its clients' interests, it moves from one type of broadcast organization to the other.

There are various reasons for changes of the client access type. New mobile clients may enter the area of coverage of the server, or previously supported mobile clients may leave the server's area. Client's interests may also change over time. A set of client access types constitutes the clients' accumulative profile, or simply profile. We consider three different profile types depending on how the profile changes over time:

1. *linear profile*,
2. *binary profile*,
3. *random profile*.

The *linear profile* simulates the case when the overall client interest migrates gradually from one access type to the other. It can be considered as a synthetic profile. In the *binary profile*, all clients synchronously change their interest. This is a synthetic profile. The *random profile* models the case in which clients independently change their access interests with some probability. The main parameter of a random profile is the probability that a client changes its interests between two consecutive broadcast cycles, broadcast cycles  $i$  and  $i + 1$ . We call this parameter *Instability*. Note that Instability is different from *Randomness degree*.

As an example, consider Table 1 that depicts the profiles of a set of five clients. The table also shows how this profile changes over four broadcast cycles. Assume that 1 denotes snapshot access and 0 denotes historical access. For instance, in the

Table 1  
Clients profile

	Broadcast cycle											
	Linear profile				Binary profile				Random profile			
	1	2	3	4	1	2	3	4	1	2	3	4
Client 1	1	1	1	1	1	0	1	0	1	0	0	1
Client 2	1	1	1	0	1	0	1	0	0	1	0	0
Client 3	1	1	0	0	1	0	1	0	1	1	0	1
Client 4	1	0	0	0	1	0	1	0	0	0	1	1
Client 5	0	0	0	0	1	0	1	0	1	0	1	0

linear profile, client 1 uses snapshot access for all broadcast cycles. Client 2 uses snapshot access for broadcast cycles 1, 2, and 3, and historical access for broadcast cycle 4.

Table 1 depicts a linear profile where the access type changes from snapshot to historical. In particular, for the first two broadcast cycles, the majority of the clients follow a snapshot access pattern, so the server uses the Vertical broadcast organization, whereas, for the last two broadcast cycles the majority changes, so the server switches to Horizontal broadcast. In the binary profile in Table 1, the access behavior of all five clients changes at each broadcast cycle. Finally, in the random profile, changes of access type occur randomly.

## 6. Experimental testbed

Our simulator consists of a broadcast server which broadcasts a specified number of versions of a set of data items, and a client who receives the data. The number of data items in the set is determined by the *Size* parameter and the number of versions by the *Versions* parameter. The communication is based on the client–server mechanism via sockets. For simplicity the data values are integer numbers from 0 to 9. The server repeatedly broadcasts data, each broadcast cycle has a number, determined by the *Bcast cycle* parameter.

The simulator runs the server in two modes, corresponding to the two broadcast organizations, namely Vertical broadcast and Horizontal broadcast (determined by the *Bcast type* parameter). The broadcast is either Compressed or basic Sequential (determined by the *Compression* parameter). The server generates broadcast data with different degrees of randomness (from 0 to 1), which is determined by the parameter *Randomness degree*.

The client searches the data by using three different access types: Browse, Snapshot and Historical (determined by the *Access type* parameter). The client generates the data elements it needs to access (various versions of data items) before tuning into the broadcast. The parameter *Elements* determines the number of data elements

to be requested by the client. For the Browse type of access, the data items and their versions are determined randomly to simulate the case when all versions of all data items are equally important for a client. For the Snapshot and the Historical accesses, the requested data elements are grouped into a number of strides (determined by *StrideN*), each containing *l* elements (determined by *StrideL*). (Clearly,  $StrideL * StrideN = Elements$ .) For example, if  $StrideN = 2$  and  $StrideL = 5$ , for Snapshot access, the client searches for 2 versions (determined randomly with Zipf distribution) of 5 consecutive data elements. For Historical access, the client tries to find 5 versions of 2 data items (determined randomly with Zipf distribution). In both cases, this pattern can be written as  $Elements = 5 \times 2$ .

The client may tune in at any point in the broadcast, but it starts its search for data elements at the beginning of the next broadcast. Thus, if a client does not tune in at the beginning of a broadcast, it sleeps until it wakes up at the beginning of the next broadcast which is determined by the next broadcast pointer in the header of each bucket. A client reads a broadcast until all the desired data elements are found. In this way, it is guaranteed that the desired data elements are found within a single broadcast. This scheme is applicable for both static and dynamic access. Note that for the broadcast with dynamic access, the performance is determined by how fast the data is found in the last broadcast cycle. Therefore, by simulating the last broadcast, it is possible to estimate the performance for both cases. While the client is reading, it counts the number and type of characters it reads. This can be converted into *Access time* – the time elapsed between the time the client starts its search and until it reads its last requested data element, given a specific data transmission rate. In our study, access time is the measure of performance for both response time and power consumption (recall that we do not consider selective tuning in this paper, hence a client stays in active mode throughout its search). The smaller the access time, the better the performance and the smaller the consumption of energy. We assume that the auxiliary characters (e.g., #, =, ^, V) consume one time unit and the

Table 2  
Simulation parameters

Parameter	Description	Default value
<i>Compression</i>	Basic Sequential broadcast Compressed broadcast	Compressed
<i>Disk</i>	0–Single disk 1–Multiple disks	0
<i>HDSize</i>	Fast disk size	
<i>Frequency</i>	Relative rotation frequency of fast disk relative to slow disk	
<i>Bcast cycle</i>	Broadcast cycle	
<i>Bcast type</i>	Vertical broadcast Horizontal broadcast	Vertical
<i>Size</i>	Number of data items	10
<i>Versions</i>	Number of versions	
<i>Randomness degree</i>	0–1, (0: all versions have the same value, 1: versions are completely independent)	0.5
<i>Length</i>	Size of a data element (size of an auxiliary symbol is 1)	16
<i>Elements</i>	Number of the requested data items	
<i>Access type</i>	Browse access Snapshot access Historical access	Browse
<i>StrideN</i>	Number of strides for Snapshot/ Historical access	
<i>StrideL</i>	Length of the strides for Snapshot/ Historical access	
<i>Tries</i>	Repetition of each experiment to reduce deviations	80
<i>Instability</i>	Probability that a client changes its access pattern	

data elements may consume 4, 16, 64, etc. time units, depending on the complexity of the data. The *Length* parameter is used to specify the size of the data element. In the experiments reported in this paper, we have chosen *Length* to be 16, which may correspond to 16 bytes.

For the broadcast disk case, we implemented two different speed disks (“fast” for all versions of frequently accessed data and “slow” for all versions of infrequently accessed data). *HDSize*, is the size of the fast disk, which corresponds to the number of values, which are “hot”. *Frequency* is the relative frequency of how faster (more frequently) the “hot” data are broadcast, compared to the “cold” data. In this case, the client generates requests to data not uniformly, but following the Zipf-distribution. We varied the *HDSize* parameter from 1 to 20 and the *Frequency* parameter from 1 to 6. The disks were implemented for both Vertical and Horizontal broadcasts. In real world, the clients are expected to let the server know about their distributions, and on the base of that knowledge the server is expected to

place the most frequently accessed data to the faster rotating disk. In our experiments, we show that, besides single disks, compression improves the performance in the case of broadcast disks as well.

In order to estimate confidence intervals, we performed the measurements 80 times (parameter *Tries*). Then, we calculated the average access time and the corresponding standard deviation which are shown in our graphs. The discussed parameters and their default values are summarized in Table 2.

## 7. Performance results

### 7.1. Single disk broadcast

In this section, we report on the results of our experiments that demonstrate the applicability of our proposed multiversion broadcast organizations and the advantages of our compression scheme.

### 7.1.1. Performance of the compression scheme

For this set of experiments, we consider the Vertical broadcast organization and the Browse access for the clients. Our results are also applicable for the other broadcast organizations and access types. *Size* is set to 90 and for the first two sets of experiments *Elements* is set to 5.

The effectiveness of the Compressed Broadcast depends on the size of the data elements on the broadcast (represented by *Length* parameter) as opposed to the size of the auxiliary symbols (note that we assume that such symbols consume one time unit). Fig. 4 shows the dependence of the access time on the size of the data items for the Compressed and the Sequential (i.e., uncompressed) server broadcasts. Compression reduces the client's access time for any size of data (about 2 times for *Randomness degree* = 0.5). The greatest gain in terms of absolute access time occurs for the largest data sizes as expected.

The effectiveness of our compression depends on the frequency of updates. The smaller the update frequency, the larger the benefits of compression. To model this, we use the average randomness *Randomness degree*[*i*] of the data element *i* to represent how frequently the value of this element changes, in particular to express the probability that any two consequent versions of *k*,




Fig. 4. Dependence of the access time on the size of the data items.




Fig. 5. Compression performance for different Randomness degree.

$Dval[k, j]$  and  $Dval[k, j - 1]$  are not equal. Fig. 5 depicts the dependence of the performance of compression on the *Randomness degree*. For *Randomness degree* = 0.0, we observe about 10 times improvement. When the versions become more different, the performance of the compressed broadcast worsens, getting close to that of the sequential broadcast as *Randomness degree* approaches 1.

The experiment shows that the proposed compressed scheme works best if data do not change from one version to another in all time interval. However, even if they do change, the Compressed broadcast just converts to a simple Sequential broadcast. The auxiliary symbols overhead is so small that the fact that even at *Randomness degree* = 1 there exist some data items for which the data values are the same for adjacent version numbers (and so, we still have some minimal compression) is enough to have some minor performance improvement. This is a situation to be expected in practice.

The dependence of the *Access time* on the number of elements requested (given by *Elements* parameter) is shown in Fig. 6. We see that at the beginning, the increase of the number of elements requested leads to a significant increase of the access time, but later (for *Elements* higher than 4)




Fig. 6. Performance for different number of searched elements.

the access time increases more slowly. This behavior does not look very surprising if we consider the access time as the time needed to search from the beginning of a bcast to “the furthest data element”. The other requested data elements are “in between” and are “picked up” on the way. As the number of *Elements* increases, the place where the last searched element was “picked up” shifts toward the end of the broadcast, making the *Access time* “saturated”. The important feature is that the absolute difference between the access time for the Compressed broadcast and the Sequential broadcast is the largest for *Elements* higher than 4. However, the relative difference stays approximately the same (slightly greater than 2 times).

The results presented in Figs. 4–6 are obtained for the Vertical broadcast and the Browse access type, but qualitatively the tendencies observed are valid for all other broadcast organizations and access schemes.

7.1.2. Broadcast organization and client access type

Fig. 7 shows the dependence of the *Access time* on *Randomness degree* when the server uses the Vertical broadcast, whereas Fig. 8 when the server uses the Horizontal broadcast. *Elements* is set to 20. The main conclusion is that for the Snapshot access type, the most efficient broadcast organiza-




Fig. 7. Vertical broadcast at different Randomness degree.

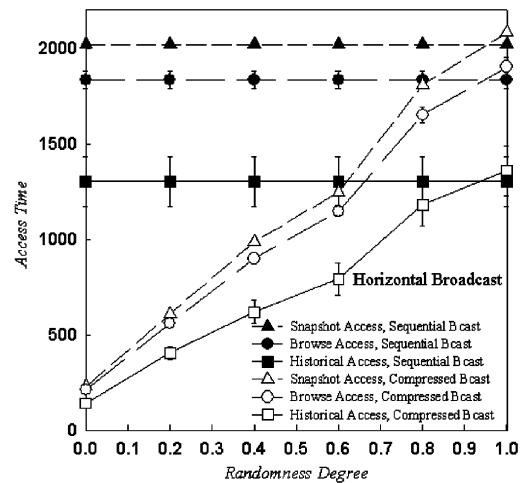


Fig. 8. Horizontal broadcast at different Randomness degree.

tion is the Vertical organization, whereas for Historical access type, the most efficient broadcast organization is the Horizontal one.

In particular, Fig. 7 shows that for the Vertical broadcast the most efficient access scheme is the Snapshot access and the worst is the Historical access (about 1.5 times worse than the Snapshot access). The Browse access is somewhere in between (about 1.4 times worse than the Snapshot access) closer to the Historical access. Fig. 8 shows



the opposite results. The best scheme for the Horizontal broadcast is the Historical access, and the worst is the Snapshot Access (about 1.4 times worse than the Historical access).

These results are valid for both Compressed and Sequential broadcasts. The interesting feature is that for small values of *Randomness degree*, it is more important for the performance whether the broadcast is Compressed or Sequential than whether the access scheme “corresponds” to the broadcast. We can observe from the figures that for *Randomness degree* less than 0.7 the *Access time* for any access type is smaller in the case of the Compressed broadcast. But for *Randomness degree* higher than 0.7, there are cases when the Sequential broadcast with the “right” access scheme can beat the Compressed broadcast with the “wrong” access scheme. Hence, in order to have the best performance, the broadcast organization and access scheme should have “similar patterns”, either Vertical broadcast organization and Snapshot access, or Horizontal broadcast organization and Historical access.

## 7.2. Broadcast disks

In this section, we restrict our discussion to the first approach for distributing data between the “fast” and the “slow” disk presented in Section 3.4. Similar results apply to the other approaches as well. The goal of our experiments are two-fold. First, we test whether the popular technique of broadcast disks is “compatible” with our compression scheme. Second, we want to get a better insight of the use of the broadcast disks for disseminating multiple versions of data. Specifically, two important parameters in this setting are the size of the fast disk and the relative speed of the two disks. We performed two experiments to see how these parameters affect performance of multiversion broadcast.

We set the relative speed of the disks (*Frequency*) as a constant equal to 2 and varied the size of the fast disk (*HDSize*). The results are presented in Fig. 9. We see that for both *Elements* values 2 and 5, the Compressed broadcast outperforms the Sequential broadcast. However for *Elements* = 2 the curve has the minimum at *HDSize* = 10 and




Fig. 9. Performance of different absolute sizes of fast and slow disk for broadcast disks.

for *Elements* = 5 the curve has the minimum at *HDSize* = 1. This means that for smaller *Elements* values, broadcast disks may improve the performance, but for larger values of *Elements*, the broadcast disks scheme creates more overhead (due to the increase of the overall broadcast size created by the repetitive broadcast of the values on the fast disk) than gains. Higher *Elements* values lead to higher probability that the necessary data are located on “slow” disks, so the overhead of repetitive fast disks overwhelms the gains produced by the different speed of the disk rotation. This observation is valid for both the Compressed and the Sequential broadcast.

In the second set of experiments, we set *HDSize* as a constant (equal to 5, 10 and 15) and varied *Frequency*. The results are shown in Fig. 10. We intentionally set *Elements* parameter to 2, because the previous experiment showed that higher values are not appropriate for the broadcast disks scheme. The results from Fig. 10 show that an optimal relative frequency of the disks does exist for certain values of *HDSize*. In this experiment, it is 3 for *HDSize* = 5, 2 for *HDSize* = 10, and the disks scheme does not improve performance (*Frequency* = 1 in fact means that the broadcast is flat) for *HDSize* = 15. Hence, we found another parameter responsible for the decision if disks

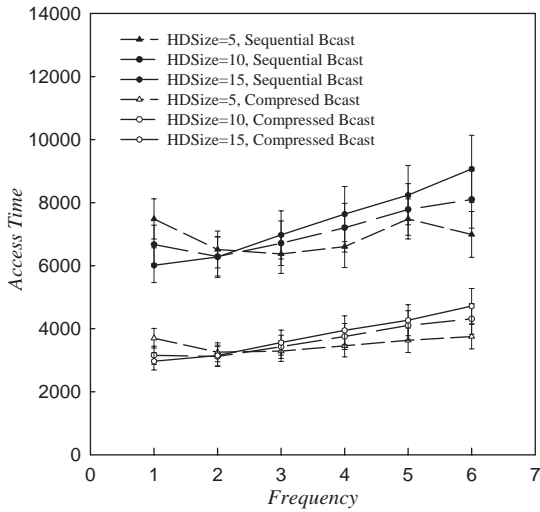


Fig. 10. Performance of different relative frequencies (fast, slow disks) for broadcast disks.

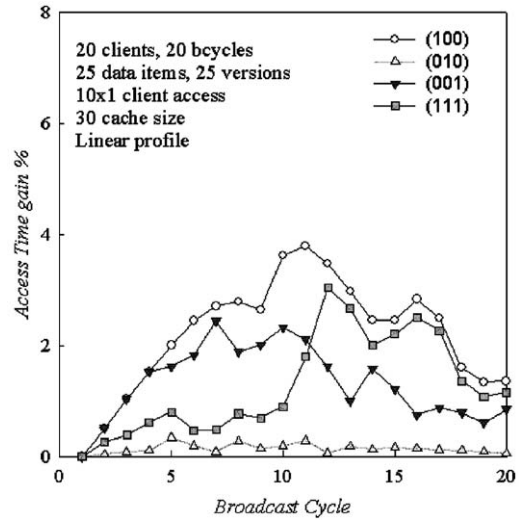


Fig. 11. Access time gain for different cache organizations in the case of the linear profile.

scheme should be used or should not. This result is valid for both the Sequential and the Compressed broadcasts.

### 7.3. Client side cache and adaptive broadcast

In this section, we consider two techniques which can reduce the effect of the “incompatibility” of the broadcast organization and the client access pattern. One technique is the client side cache and the other is the server adaptive broadcast. In all the following experiments, *Size* and *Versions* are each set to 25.

We consider an adaptive compressed broadcast organization. Figs. 11–12 (*Elements* = 10 × 1, *Cachesize* = 30) show the Access time gains and the Hit rates for the Linear profile of 20 clients for 20 broadcast cycles. Four different cache organizations are considered. The results for the other two profiles, Binary and Random, of 20 clients for 20 broadcast cycles are shown in Figs. 13–14 (*Elements* = 15 × 2, *Cachesize* = 90) and 15–16 (*Elements* = 4 × 2, *Cachesize* = 60), respectively. For all profiles, the highest hit rate was obtained for (100) and (001) cache organizations (15–20%) and the lowest for (010) cache organization (around zero). The “aggregate” (111) cache organization behaved a little worse than (100)

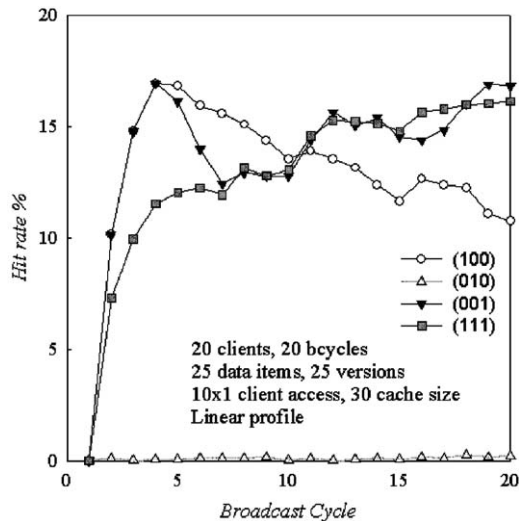


Fig. 12. Hit rate for different cache organizations in the case of the linear profile.

and (001). Similar results were obtained for the access time gain, which was 2–4% for (100) and (001) cache organizations and practically zero for (010). The (111) cache organization again was a little worse than (001) and (100) types.

Similar experiments were performed with non-adaptive server broadcast. The behavior of the cache was the same: (100) and (001) caches had the

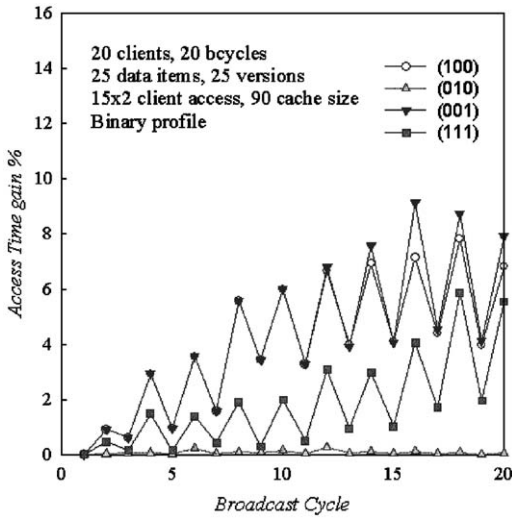


Fig. 13. Access time gain for different cache organizations in the case of the binary profile.

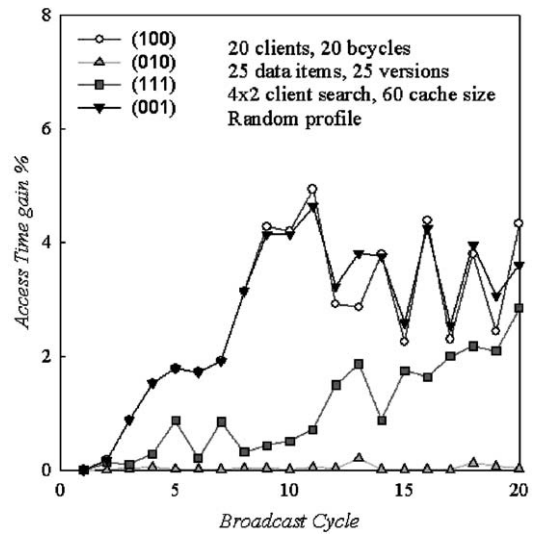


Fig. 15. Access time gain for different cache organizations in the case of the random profile.




Fig. 14. Hit rate for different cache organizations in the case of the binary profile.




Fig. 16. Hit rate for different cache organizations in the case of the random profile.

access time gain about 2–4%, (111) cache had it about 1–3%, and (010) had no gain in the access time at all.

We describe the next experiment with adaptive and non-adaptive (vertical) broadcast; these experiments are performed for all four cache organizations as well as without a cache. To make the figures more clear, we include the data for

(111) cache only. The other cache organization showed very similar behavior.

Fig. 17 (*Elements* = 10 × 1) shows the Access time vs. Broadcast cycle for the adaptive and non-adaptive broadcast with (111) type client cache in the case of the Linear profile of 20 clients for 20 broadcast cycles. Initially (broadcast cycle = 1), most of the clients have the Snapshot access

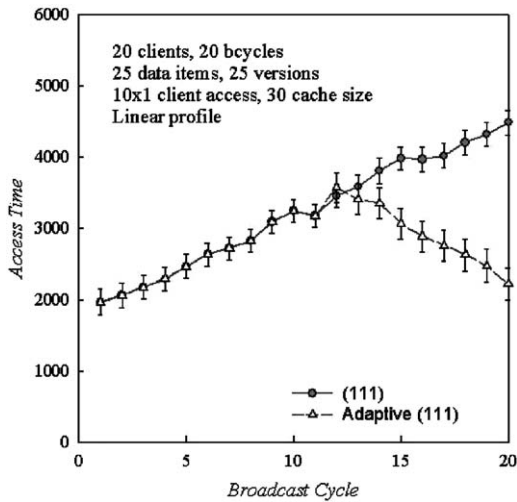


Fig. 17. Effect of the adaptive broadcast in the case of the linear profile.

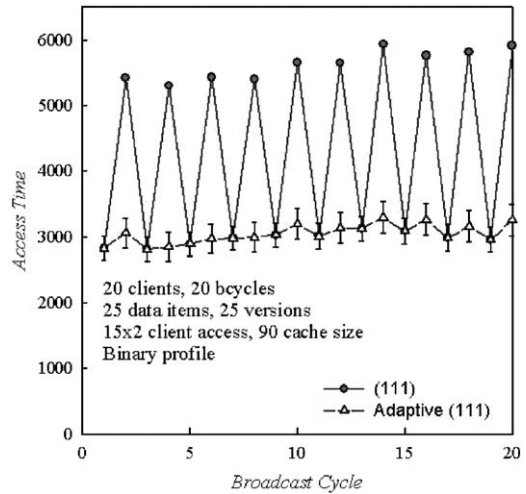


Fig. 18. Effect of the adaptive broadcast in the case of the binary profile.

pattern. In case of non-adaptive broadcasting server, which is chosen by default to be the Vertical, the access time is low because the broadcast organization (Vertical) and the client access pattern of the majority of clients (Snapshot) match. Later (higher broadcast cycles) the access time steadily goes up and reaches the maximum at the end of broadcast (broadcast cycle = 20). At this point all clients use the Historical access, which does not match the broadcast organization. In the case of adaptive broadcast, the broadcast organization in the beginning (broadcast cycle 1) is Vertical that satisfies most of the clients. Therefore the access time is small. Then for broadcast cycles close to 10, the number of clients with the Snapshot and the Historical access types is the same, so about a half of the clients are satisfied. As a result, the access time goes up. When broadcast cycle reaches 10, the adaptive broadcast switches from the Vertical to the Horizontal and the access time goes down.

Fig. 18 ( $Elements = 15 \times 2$ ) shows the Access time vs. Broadcast cycle in the case of the Binary profile of 20 clients for 20 broadcast cycles. As can be seen, non-adaptive broadcast produces bursts of low and high access time, in accordance with the matching and non-matching the client access pattern to the server broadcast organization. As




Fig. 19. Effect of the adaptive broadcast in the case of the random profile.

expected, the adaptive broadcast removes those bursts and makes the performance much smoother and closer to the best possible one (i.e., minimum access time).

Fig. 19 ( $Elements = 4 \times 2$ ) shows the Access time vs. Broadcast cycle in case of the Random profile of 20 clients for 20 broadcast cycles. On the average, the access time gain due to adaptive broadcast is about 15%. The gain depends on the

ratio of the clients with different interests. The farther from unity this ratio is, the greater the gain. We also performed experiments with different *Instability* parameters and found that the access time reduction is insensitive to that. However, the high values of *Instability* incurs higher overhead due to the frequent messages from the client to the server on the uplink channel.

The simulation data show that both adaptive broadcast and client side cache are good tools for the multiversion broadcast performance improvement in the case of “incoherent” (i.e., Vertical–Historical or Horizontal–Snapshot) server broadcast organization and client access pattern. Caching is a local method and relies only on local resources whereas adaptive broadcast is a global method and involves extra communication. Both techniques are compatible and hence can be applied at the same time, creating an even greater cumulative positive effect.

Finally, it should be noted that feedback messages are not uncommon in a typical broadcast push environment in which the server relies on such feedback messages to decide on the content of the broadcast. Our adaptive broadcast proposal can be implemented in conjunction with a document selection method so that the data items and their organization on the broadcast cater better to the needs of the majority of the clients, hence achieving even greater scalability.

## 8. Related work

The work in this paper falls within the solutions that use multicast/broadcast communication to support scalable dissemination of information. These solutions efficiently disseminate data to a large number of users by any combination of the following two schemes: *broadcast pull* and *broadcast push*. In broadcast pull, the clients make explicit requests for data. If multiple clients request the same data at approximately the same time, the server may aggregate these requests, and only broadcast the data once. Such a scheme makes an effective use of the low-wireless bandwidth and clearly improves user-perceived performance. Several preemptive and non-preemptive

scheduling algorithms have been proposed that attempt to achieve maximum aggregation both based on the exact match of the requests [1,11–13,34,36] and based on derivation dependency among requests [14,15].

In broadcast push, which is the context of our work, techniques for data organization, indexing for selective tuning and caching have been investigated. Among the first data organizations was the broadcast disks organization studied in Section 3.4 [9], while a scheme to generate non-uniform broadcasts that support range queries is described in [16]. Power conservative indexing methods for single-attribute and multi-attribute based queries appeared in [17–19,37].

The most commonly used caching policy LRU may lead to suboptimal performance in multicast environments and several algorithm has been investigated to improve it. For example, Landlord/Greedy-Dual-Size [20,21] attempts to improve LRU by taking into account the cost of downloading a page whereas the PT [22] and Gray [23] algorithms take in addition the popularity of a page. As part of our future work, we plan to evaluate such methods in the context of multiversion broadcast environments.

The concept of autoupfetching has been proposed in the context of mobile computing systems in an effort to improve cache performance to reduce access time without increasing power consumption [24,25]. These schemes dynamically predict what data to be prefetched based on the data access frequency and update frequency. These approaches share similar objective with our proposed multiversion autoupfetch scheme but our scheme uses a simple prediction which autoupfetches only new versions of the already cached data items. That is, in some sense, our autoupfetch behaves similar to a data recharging mechanism [26].

As mentioned in the introduction, a major goal of our work on multiversion broadcasts has been to achieve consistency and concurrency in broadcast environments. In [27], the authors discuss the tradeoffs between concurrency of data and performance issues when some of the broadcast data items are updated by the processes running on the server. However, the updates do not have



transactional semantics associated with them either at the server or at the clients. Realizing that serializability as the correctness criterion may be expensive, and perhaps unnecessary in such environments, and hence various protocols [2,5,28–30] attempt to cater to less demanding correctness requirements. All these protocols have been formally analyzed within a unified framework [31].

Finally, research on balancing the broadcast push of information and broadcast pull data delivery methods appeared in Refs. [1,32,33].

## 9. Conclusions

In this paper, we identify two basic multiversion organizations, namely Vertical and Horizontal broadcasts and propose an efficient compression scheme applicable to both. We also examine the applicability of these schemes in the context of both single disk and multiple disk broadcasts. We consider multiversion client data caching and introduce appropriate cache replacement techniques.

Our performance evaluation results showed that besides the size of a broadcast, the organization of the broadcast has an impact on performance, as different kinds of clients need different types of data. We recognized three kinds of clients' applications based on their access behavior: "historical" that access many versions of the same data, "snapshot" that access different data of the same version and "browsing" that access data and versions randomly.

Specifically, if the primary interest of clients is "historical" applications, the best way to broadcast is the Horizontal broadcast. If the primary interest of clients is "snapshot" applications, the best way to broadcast is the Vertical broadcast. In case of mixed environment it is possible to create adaptive broadcast with no extra cost due to flexibility of the broadcast format.

The suggested compression technique does not require extra time for client side decompression and works for both Vertical and Horizontal broadcasts. The auxiliary symbols overhead is small if the size of one data element significantly

exceeds a few bits. The effectiveness of a compressed broadcast depends on the repetitiveness of the data. The less frequent the data change, the better the gains are. But even in the worst case (completely random data), the compressed broadcast does not exhibit worse performance than the uncompressed broadcast.

The use of our compression scheme in the client cache exhibited similar advantages as in the case of the compressed broadcast, effectively increasing the size of the cache and consequently, the number of hits. However, the most interesting property exhibited by the client caching is that it can be used as a tool to ameliorate the negative effects due to any incompatibilities between a broadcast organization and a client data access behavior.

## Acknowledgements

This work was supported in part by the National Science Foundation award ANI-0123705 and in part by the European Union through grant IST-2001-32645.

## References

- [1] S. Acharya, M. Franklin, S. Zdonik, Balancing push and pull for data broadcast, ACM Sigmod Conference, 1997, pp. 183–194.
- [2] J. Shanmugasundaram et al., Efficient concurrency control for broadcast environments, ACM Sigmod Conference, 1999, pp. 85–96.
- [3] K.-L. Wu, P.S. Yu, M.-S. Chen, Energy-efficient mobile cache invalidation, *Distrib. Parallel Dat.* 6 (1998) 351–372.
- [4] J. Jing, A.H. Elmargamid, S. Helal, R. Alonso, Bit-sequences: an adaptive cache invalidation method in mobile client/server environment, *ACM/Baltzer Mobile Networks Appl.* 2 (2) (1997) 115–127.
- [5] E. Pitoura, P.K. Chrysanthis, Exploiting versions for handling updates in broadcast disks, 25th VLDB Conference, 1999, pp. 114–125.
- [6] C. Mohan, H. Pirahesh, R. Lorie, Efficient and flexible methods for transient versioning to avoid locking by read-only transactions, ACM Sigmod Conference, 1992, pp. 124–133.
- [7] E. Pitoura, P. Chrysanthis, Multiversion data broadcast, *IEEE Trans. Comput.* 51 (10) (2002) 1224–1230.
- [8] S.W. Golomb, "Run-length encodings", *IEEE. Trans. on Information Theory* 12 (July 1996) 399–401.

- [9] S. Acharya, R. Alonso, M. Franklin, S. Zdonik, Broadcast disks: Data management for asymmetric communication environments, ACM Sigmod Conference, 1995, pp. 199–210.
- [10] O. Shigiltchoff, P.K. Chrysanthis, E. Pitoura, Broadcast data organizations and client side cache, 23rd ICDCS Workshops, 2003, 420–426.
- [11] H.D. Dykeman, M. Ammar, J.W. Wong, Scheduling algorithms for videotex systems under broadcast delivery, IEEE Intl Conference on Communications, 1986, pp. 1847–1851.
- [12] N.H. Vaidya, S. Hameed, Scheduling data broadcast in asymmetric communication environments, ACM/Baltzer Wireless Networks 5 (3) (1999) 171–182.
- [13] J.W. Wong, Broadcast delivery, Proc. IEEE 76 (1988) 1566–1577.
- [14] M.A. Sharaf, P.K. Chrysanthis, Facilitating mobile decision making, 2nd ACM Mobicom Workshop on Mobile Commerce, 2002, pp. 45–53.
- [15] M.A. Sharaf, P.K. Chrysanthis, Semantic-based delivery of OLAP summary tables in wireless environments, 10th ACM Intl Conference on Information and Knowledge Management, 2002, pp. 84–92.
- [16] K.-L. Tan, J.X. Yu, Generating broadcast programs that support range queries, IEEE Trans. Knowledge Data Eng. 10 (4) (1998) 668–672.
- [17] T. Imielinski, S. Viswanathan, B.R. Badrinath, Energy efficient indexing on air, ACM Sigmod Conference, 1994, pp. 25–36.
- [18] Q. Hu, W.-C. Lee, D.L. Lee, Power conservative multi-attribute queries on data broadcast, 16th IEEE ICDE Conference, 2000, pp. 157–166.
- [19] R. Agrawal, P.K. Chrysanthis, Efficient data dissemination to mobile clients in e-commerce applications, 3rd Intl Workshop on E-Commerce and Web Information Systems, 2001, pp. 58–65.
- [20] P. Cao, S. Irani, Cost-aware WWW proxy caching algorithms, 1st USENIX Symposium on Internet Technologies and Systems, 1997, pp. 193–206.
- [21] N. Young, On-line file caching, 9th ACM-SIAM Symposium on Discrete Algorithms, 1998, pp. 82–91.
- [22] S. Acharya, M. Franklin, S. Zdonik, Prefetching from a broadcast disk, 12th IEEE ICDE Conference, 1996, pp. 276–285.
- [23] S. Khanna, V. Liberatore, On broadcast disk paging, SIAM J. Comput. 29 (5) (2000) 1683–1702.
- [24] G. Cao, Proactive power-aware cache management for mobile computing systems, IEEE Trans. Comput. 51 (6) (2002) 608–621.
- [25] H. Hu, J. Xu, D.L. Lee, Adaptive power-aware prefetching schemes for mobile broadcast environments, 4th Mobile Data Management Conference, CNCS Vol. 2574, Springer, 2003, pp. 374–380.
- [26] M. Cherniak, M.J. Franklin, S. Zdonik, Expressing user profiles for data rechatging, IEEE Pers. Commun.: Special Issue on Pervasive Computing 8 (4) (Aug. 2001) 32–38.
- [27] S. Acharya, M. Franklin, S. Zdonik, Disseminating updates on broadcast disks, 22nd VLDB Conference, 1996, pp. 354–365.
- [28] D. Barbará, Certification reports: supporting transactions in wireless systems, 17th IEEE ICDCS Conference, 1997, pp. 466–473.
- [29] E. Pitoura, P.K. Chrysanthis, Scalable processing of read-only transactions in broadcast push, 19th IEEE ICDCS Conference 1999, pp. 432–439.
- [30] V.C.S. Lee, S.H. Son, K. Lam, On the performance of transaction processing in broadcast environments, 1st Conference on Mobile Data Access (MDA'99), 1999.
- [31] E. Pitoura, P.K. Chrysanthis, K. Ramamritham, Characterizing the temporal and semantic coherency of broadcast-based data dissemination, 9th International Conference on Database Theory, 2003, pp. 410–424.
- [32] K. Stathatos, N. Roussopoulos, J.S. Baras, Adaptive data broadcast in hybrid networks, 23rd VLDB Conference, 1997, pp. 326–335.
- [33] M. Bhide, P. Deolasee, A. Katkar, A. Panchbudhe, K. Ramamritham, P.J. Shenoy, Adaptive push–pull: disseminating dynamic web data, IEEE Trans. Comput. 51 (6) (2002) 652–668.
- [34] S. Acharya, S. Muthukrishnan, Scheduling on-demand broadcasts: new matrices and algorithms, 4th Annual ACM/IEEE MobiCom Conference, 1998, pp. 43–54.
- [35] S. Acharya, M. Franklin, S. Zdonik, Dissemination-based data delivery using broadcast disks, IEEE Pers. Commun. 2 (6) (1995) 50–61.
- [36] D. Aksoy, M. Franklin, RxW: A scheduling approach for large-scale on-demand data broadcast, IEEE/ACM Trans. Network. 7 (6) (1999) 846–860.
- [37] T. Imielinski, et al., Data on air: organization and access, IEEE Trans. Knowledge Data Eng. 9 (3) (1997) 353–372.