

Data Broadcasting, Caching, and Replication in Mobile Computing

Panos K. Chrysanthis¹ and Evangelia Pitoura²

¹Department of Computer Science, University of Pittsburgh, Pittsburgh, PA, USA

²Department of Computer Science and Engineering, University of Ioannina, Ioannina, Greece

Synonyms

Data copy; Data dissemination; Push/pull delivery

Definition

Mobile computing devices (such as portable computers or cellular phones) have the ability to communicate while moving by being connected to the rest of the network through a wireless link. There are two general underlying infrastructures: single-hop and multi-hop ones. In *single-hop* infrastructures, each mobile device communicates with a stationary host, which corresponds to its point of attachment to the wired network. In *multi-hop* infrastructures, an ad-hoc wireless network is formed in which mobile hosts participate in routing messages among each other. In both infrastructures, the hosts between the source (or sources) and the requester of data (or data sink) form a dissemination tree. The hosts (mobile or stationary) that form the dissemination tree may store data and participate in computations towards achieving *in network* processing. Challenges include [1], (i) *intermittent connectivity*, which refers to both short and long periods of network unavailability, (ii) *scarcity of resources*, including storage and battery life, and (iii) *mobility* itself.

To handle these challenges, data items may be stored locally (cached or replicated) at the requester or at the intermediate nodes of the dissemination tree. Cache and replication aim

at increasing availability in the case of network disconnections or host failures as well as at handling intermittent connectivity. Mobility introduces additional challenges in maintaining cache and replica consistency and in replica placement protocols. Wireless data delivery in both infrastructures physically supports broadcasting. This broadcast facility has been used for providing a push-mode of data dissemination where a server broadcasts data to a large client population often without an explicit request from the clients. Issues addressed by related research include broadcast scheduling and organization (i.e., which items to broadcast and in which order), indexing broadcast data, and update propagation.

Historical Background

Mobile computing can be traced back to file systems and the need for disconnected operations in the late 1980s. With the rapid growth in mobile technologies and the cost effectiveness in deploying wireless networks in the 1990s, the goal of mobile computing was the support of AAA (anytime, anywhere and any-form) access to data by users from their portable computers, mobile phones and other devices with small displays and limited resources. These advances motivated research in data management in the early 1990s.

Foundations

Data Broadcasting

Many forms of wireless network infrastructures rely on broadcast technology to deliver data to large client populations. As opposed to point-to-point data delivery, broadcast delivery is scalable, since a single broadcast response can potentially satisfy many clients simultaneously. There are two basic modes of broadcast data delivery: pull-based and push-based. With *push-based* data delivery, the server sends data to clients without an explicit request. With *pull-based* or *on-demand* broadcast delivery, data are delivered only after a specific client request. In general, access to broadcast data is sequential with clients moni-

toring the broadcast channel and retrieving any data items of interest as they arrive. The smallest access unit of broadcast data is commonly called a *bucket* or *page*.

Scheduling and Organization

A central issue is determining the content of the broadcast or *broadcast scheduling*. Scheduling depends on whether we have on demand, push or hybrid delivery. In on-demand broadcast, there is an up-link channel available to clients to submit requests. The item to be broadcast next is chosen among those for which there are pending requests. Common heuristics for on-demand scheduling include First Come First Served and Longest Wait First [2]. The $R \times W$ strategy selects the data item with the maximal $R \times W$ value, where R is the number of pending requests for an item and W the amount of time that the oldest pending request for that item has spent waiting to be served [3]. More recent schemes extended $R \times W$ to consider the semantics of the requested data and applications such as subsumption properties in data cubes [4]. Push-based broadcast scheduling assumes a-priori knowledge of client access distributions and prepares an off-line schedule. Push-based data delivery is often periodic. In hybrid broadcast, the set of items is partitioned, so that some items are pushed, i.e., broadcast continuously, and the rest are pulled, i.e., broadcast only after being requested [5]. Commonly, the partition between push and pull data is based on popularity with the most popular items being pushed periodically and the rest delivered on demand. One problem is that for push items, there is no way to detect any changes in their popularity. One solution is to occasionally stop broadcasting some pushed items. This forces clients to send explicit requests for them, which can be used to estimate their popularity [6]. An alternative that avoids flooding of requests requires a percentage of the clients to submit an explicit request irrespective of whether or not a data item appears on the broadcast [7].

The organization of the broadcast content is often called *broadcast program*. In general, broadcast organizations can be classified as either *flat* where each item is broadcast exactly

once or *skewed* where an item may appear more than once. One can also distinguish between *clustered* organizations, where data items having the same or similar values at some attribute appear consecutively, and *non-clustered* ones, where there is no such correlation. In skewed organizations, the broadcast frequency of each item depends on its popularity. For achieving optimal access latency or response time, it was shown that (i) the relative number of appearances of items should be proportional to the square root of their access probabilities and (ii) successive broadcasts of the same item should be at equal distances [2]. It was also shown that the *Mean Aggregate Access* (MAD) policy that selects to broadcast next the item whose access probability \times the interval since its last broadcast is the highest achieves close to optimal response time [8]. Along these lines, a practical skewed push broadcast organization is that of *broadcast disks* [9]. Items are assigned to virtual disks with different “speeds” based on their popularity with popular items being assigned to fast disks. The spin speed of each disk is simulated by the frequency with which the items assigned to it are broadcast. For example, the fact that a disk D_1 is three times faster than a disk D_2 , means that items assigned to D_1 are broadcast three times as often as those assigned to D_2 . To achieve this, each disk is split into smaller equal-sized units called chunks, where the number of chunks per disk is inversely proportional to the relative frequency of the disk. The broadcast program is generated by broadcasting one chunk from each disk and cycling through all the chunks sequentially over all the disks.

Indexing

To reduce energy consumption, a mobile device may switch to *doze* or *sleep* mode when inactive. Thus, research in wireless broadcast also considers reducing the *tuning time* defined as the amount of time a mobile client remains active listening to the broadcast. This is achieved by including index entries in the broadcast so that by reading them, the client can determine when to tune in next to access the actual data of interest. Adding index entries increases the size of the

broadcast and thus may increase access time. The objective is to develop methods for allocating index entries together with data entries on the broadcast channel so that both access and tuning time are optimized. In $(1, m)$ indexing [10], an index for all data items is broadcast following every fraction $(1/m)$ of the broadcast data items. *Distributed indexing* [10] improves over this method by instead of replicating the whole index m times, each index segment describes only the data items that follow it. Following the same principles, different indexing schemes have been proposed that support different query types or offer different trade-offs between access and tuning time. Finally, instead of broadcasting an index, *hashing-based techniques* have also been applied.

Data Caching and Replication

A mobile computing device (such as a portable computer or cellular phone) is connected to the rest of the network through a wireless link. Wireless communication has a double impact on the mobile device since the limited bandwidth of wireless links increases the response times for accessing remote data from a mobile host and transmitting as well as receiving of data are high energy consumption operations. The principal goal of caching and replication is to store appropriate pieces of data locally at the mobile device so that it can operate on its own data, thus reducing the need for communication that consumes both energy and bandwidth. Several cost-based caching policies along the principles of *greedy-dual* ones have been proposed that consider energy cost.

In the case of broadcast push, the broadcast itself can be viewed as a “cache in the air”. Hence, in contrast to traditional policies, performance can be improved by clients caching those items that are accessed frequently by them but are not popular enough among all clients to be broadcast frequently. For instance, a cost-based cache replacement policy selects as a victim the page with the lowest p/x value, where p is the local access probability of the page and x its broadcast frequency [9]. Prefetching can also be performed with low overhead, since data items

are broadcast anyway. A simple prefetch heuristic evaluates the worth of each page on the broadcast to determine whether it is more valuable than some other page in cache and if so, it swaps the cache page with the broadcast one.

Replication is also deployed to support *disconnected operation* that refers to the autonomous operation of a mobile client, when network connectivity becomes either unavailable (for instance, due to physical constraints), or undesirable (for example, for reducing power consumption). Preloading or prefetching data to sustain a forthcoming disconnection is often termed *hoarding*. Optimistic approaches to consistency control are typically deployed that allow data to be accessed concurrently at multiple sites without a priori synchronization between the sites, potentially resulting in short term inconsistencies. At some point, operations performed at the mobile device must be synchronized with operations performed at other sites. Synchronization depends on the level at which correctness is sought. This can be roughly categorized as replica-level correctness and transaction-level correctness. At the *replica level*, correctness or coherency requirements are expressed per item in terms of the allowable divergence among the values of the copies of each item. At the *transaction level*, the strictest form of correctness is achieved through global serializability that requires the execution of all transactions running at mobile and stationary hosts to be equivalent to some serial execution of the same transactions. With regards to update propagation with *eager replication*, all copies of an item are synchronized within a single transaction, whereas with *lazy replication*, transactions for keeping replica coherent execute as separate, independent database transactions after the original transaction commits.

Common characteristics of protocols for consistency in mobile computing include:

- The propagation of updates performed at the mobile site follows in general lazy protocols.
- Reads are allowed at the local data, while updates of local data are tentative in the sense

that they need to be further validated before commitment.

- For integrating operations at the mobile hosts with transactions at other sites, in the case of replica-level consistency, copies of each item are reconciled following some conflict resolution protocol. At the transaction-level, local transactions are validated against some application or system level criterion. If the criterion is met, the transaction is committed. Otherwise, the execution of the transaction is either aborted, reconciled or compensated.

Representative approaches along these lines include *isolation-only transactions* in Coda, *mobile open-nested transactions* [11], *two-tier replications* [12], *two-layer transactions* [13] and *Bayou* [14].

When local copies are read-only, a central issue is the design of efficient protocols for disseminating server updates to mobile clients. A server is called *stateful*, if it maintains information about its clients and the content of their caches and *stateless* otherwise. A server may use broadcasting to efficiently propagate update reports to all of its clients. Such update reports vary on the type of information they convey to the clients, for instance, they may include just the identifiers of the updated items or the updated values themselves. They may also provide information for individual items or aggregate information for sets of items. Update propagation may be either synchronous or asynchronous. In *asynchronous* methods, update reports are broadcast as the updates are performed. In *synchronous* methods, the server broadcasts an update report periodically. A client must listen for the report first to decide whether its cache is valid or not. This adds some latency to query processing, however, each client needs only tune in periodically to read the report. The efficiency of update dissemination protocols for clients with different connectivity behavior, such as for workaholics (i.e., often connected clients) and sleepers (i.e., often disconnected clients), is evaluated in [15].

Finally, in the case of broadcast push-data delivery, clients may read items from different

broadcast programs. The *currency* of the set of data items read by each client can be characterized based on the current values of the corresponding items at the server and on the temporal discrepancy among the values of the items in the set [16]. A more strict notion of correctness may be achieved through transaction-level correctness by requiring the client read-only transactions to be serializable with the server transactions. Methods for doing so include: (i) an *invalidation* method [17], where the server broadcasts an invalidation report that includes the data items that have been updated since the broadcast of the previous report, and transactions that have read obsolete items are aborted, (ii) *serialization graph testing* (SGT) [17], where the server broadcasts control information related to conflicting operations, and (iii) *multiversion broadcast* [18], where multiple versions of each item are broadcast, so that client transactions always read a consistent database snapshot.

Key Applications

Data broadcasting, caching and replication techniques are part of the core of any application that requires data sharing and synchronization among mobile devices and data servers. Such applications include vehicle dispatching, object tracking, points of sale (e.g., ambulance and taxi services, FedEx/UPS), and collaborative applications (e.g., homecare, video gaming). They are also part of embedded or light versions of database management systems that extend enterprise applications to mobile devices. These include among others Sybase Inc. 's SQL Anywhere, IBM's DB2 Everywhere, Microsoft SQL Server Compact, Oracle9i Lite and SQL Anywhere Technologies' Ultralite.

Cross-References

- [Hash-Based Indexing](#)
- [MANET Databases](#)
- [Mobile Database](#)

- Replicated Database Concurrency Control
- Transaction Management

Recommended Reading

1. Pitoura E, Samaras G. Data management for mobile computing. Boston: Kluwer; 1998.
2. Dykeman HD, Ammar MH, Wong JW. Scheduling algorithms for videotex systems under broadcast delivery. In: Proceedings of the IEEE International Conference on Communications; 1986. p. 1847–51.
3. Aksoy D, Franklin MJ. RxW: a scheduling approach for large scale on-demand broadcast. *IEEE/ACM Trans Netw.* 1999;7(6):846–60.
4. Sharaf MA, Chrysanthis PK. On-demand data broadcasting for mobile decision making. *Mob Netw Appl.* 2004;9(6):703–14.
5. Acharya S, Franklin MJ, Zdonik SB. Balancing push and pull for data broadcast. In: Proceedings of the ACM SIGMOD International Conference on Management of Data; 1997. p. 183–94.
6. Stathatos K, Roussopoulos N, Baras JS. Adaptive data broadcast in hybrid networks. In: Proceedings of the 23th International Conference on Very Large Data Bases; 1997. p. 326–35.
7. Beaver J, Chrysanthis PK, Pruh K. To broadcast push or not and what? In: Proceedings of the 7th International Conference on Mobile Data Management; 2006. p. 40–5.
8. Su CJ, Tassiulas L, Tsotras VJ. Broadcast scheduling for information distribution. *Wirel Netw.* 1999;5(2):137–47.
9. Acharya S, Alonso R, Franklin MJ, Zdonik SB. Broadcast disks: data management for asymmetric communications environments. In: Proceedings of the ACM SIGMOD International Conference on Management of Data; 1995. p. 199–210.
10. Tomasz I, Viswanathan S, Badrinath BR. Data on air: organization and access. *IEEE Trans Knowl Data Eng.* 1997;9(3):353–72.
11. Chrysanthis PK. Transaction processing in a mobile computing environment. In: Proceedings of the IEEE Workshop on Advances in Parallel and Distributed Systems; 1993. p. 77–82.
12. Gray J, Helland P, Neil PO, Shasha D. The dangers of replication and a solution. In: Proceedings of the ACM SIGMOD International Conference on Management of Data; 1996. p. 173–82.
13. Pitoura E, Bhargava B. Data consistency in intermittently connected distributed systems. *IEEE Trans Knowl Data Eng.* 1999;11(6):896–915.
14. Petersen K, Spreitzer M, Theimer M, Demers AJ. Flexible update propagation for weakly consistent replication. In: Proceedings of the 16th ACM Symposium on Operating System Principles; 1997. p. 288–301.
15. Barbará D, Imielinski T. Sleepers and workaholics: caching strategies in mobile environments. *VLDB J.* 1995;4(4):567–602.
16. Pitoura E, Chrysanthis PK, Ramamritham K. Characterizing the temporal and semantic coherency of broadcast-based data dissemination. In: Proceedings of the 9th International Conference on Database Theory; 2003. p. 410–24.
17. Pitoura E, Chrysanthis PK. Scalable processing of read-only transactions in broadcast push. In: Proceedings of the 19th International Conference on Distributed Computing Systems; 1999. p. 432–9.
18. Pitoura E, Chrysanthis PK. Exploiting versions for handling updates in broadcast disks. In: Proceedings of the 25th International Conference on Very Large Data Bases; 1999. p. 114–25.

D

Data Center Energy Efficiency

Fred Chong
Computer Science, University of Chicago,
Chicago, IL, USA

Our increasingly data-centric society has required an exponentially growing capability in large-scale infrastructure to store and process data. Although advances in silicon technology have historically provided exponential gains in both cost and energy efficiency, two factors will make the energy-efficient design of data centers of critical importance in the future.

First, increased demand for datacenter infrastructure is on a faster exponential than Moore's law (Fig. 1).

Second, although "Moore's law" continues to offer exponentially more transistors on a chip (albeit at a slower exponential than before), increases in energy efficiency are dependent on "Dennard scaling." Dennard scaling involves decreases in the voltage at which chips operate over successive generations. Unfortunately, Dennard scaling ended several years ago as the operating voltage became too close to the on-off voltage boundary of a transistor (the "threshold" voltage). Lacking Dennard scaling, future chips may become limited by maximum power budgets, requiring some fraction of the chip to remain off at all times (termed "dark silicon").