

Mobile Computing and Databases—A Survey

Daniel Barbará

Abstract—The emergence of powerful portable computers, along with advances in wireless communication technologies, has made mobile computing a reality. Among the applications that are finding their way to the market of mobile computing—those that involve data management—hold a prominent position. In the past few years, there has been a tremendous surge of research in the area of data management in mobile computing. This research has produced interesting results in areas such as data dissemination over limited bandwidth channels, location-dependent querying of data, and advanced interfaces for mobile computers. This paper is an effort to survey these techniques and to classify this research in a few broad areas.

Index Terms—Mobile computing, databases, data dissemination, bandwidth, location-dependent queries, interfaces, transaction management.

1 INTRODUCTION

MOBILE computing has become a reality thanks to the convergence of two technologies: the appearance of powerful portable computers and the development of fast reliable networks. In the mobile wireless computing environment of the future [33] massive number of low powered computer machines will query databases over the wireless communication channels. Mobile clients will often be disconnected for prolonged periods of time due to the battery power saving measures; they will also frequently relocate between different cells and connect to different data servers at different times.

The mobile or nomadic computing environment no longer requires users to maintain a fixed and universally known position in the network and enables unrestricted mobility of the users. Mobility and portability will create an entire new class of applications and possibly new massive markets combining personal computing and consumer electronics.

In this paper, we survey the impact that mobile computing has had in the area of data management. We first analyze each of the distinctive features of mobile computing and how they impact the implementation of databases for mobile computers, creating new opportunities for research. Then for each of these problems we survey the papers that have appeared in the literature in recent years.

Fig. 1 displays the architecture of the general mobile environment. It consists of two distinct sets of entities: mobile units and fixed hosts. Some of the fixed hosts, called Mobile Support Stations, are augmented with a wireless interface to communicate with mobile units which are located within a radio coverage area called a *cell*. A cell could be a real cell as in cellular communication network or a wireless local area network which operates within the area of a building. In the former case the bandwidth will be severely limited

(supporting data rates on the order of 10 to 20 Kbits/sec.). In the latter, the bandwidth is much wider—up to 10 Mb/sec. Fixed hosts will communicate over the fixed network, while mobile units will communicate with other hosts (mobile or fixed) via a wireless channel.

Although a wireless network with mobile clients is essentially a distributed system, there are some characteristic features that make the system unique and a fertile area of research. These are:

- *Asymmetry in the communications*: The bandwidth in the downstream direction (servers-to-clients) is much greater than that in the upstream direction (clients-to-servers). Moreover, in some systems, the clients do not have the capacity to send messages to the servers at all. As we mentioned above, even the bandwidth in the downstream direction is limited: 10 to 20 Kbits/sec. in a cellular network and up to 10 Mb/sec. in a wireless LAN.
- *Frequent disconnections*: Mobile clients do not stay connected to the network continuously (as fixed hosts do), but rather users switch their units on and off regularly. Moreover, mobile clients can roam, disconnecting from a cell to connect to another.
- *Power limitations*: Some of the portable units are severely limited by the amount of energy they can use before the batteries have to be recharged.
- *Screen size*: Some of the portable units, such as the Personal Digital Assistants, have very small screens.

Each one of these features has an impact on how data can be effectively managed in a system with mobile clients. The communication asymmetry, along with the restriction in power that the mobile units have, make the model of broadcasting data to the clients (instead of waiting for the clients to request specific data items), an attractive proposition. This is called data dissemination and is the topic of Section 2 in this paper. The limited bandwidth and the pattern of frequent disconnections have a clear impact on how transaction management is implemented and how data consistency is guaranteed in this environment. This is the subject of Section 3. The roaming of clients through

• D. Barbará is with the Information and Software Engineering Department, George Mason University, Fairfax, VA 22030. E-mail: dbarbara@gmu.edu.

Manuscript received 20 May 1997; revised 20 Aug. 1998.

For information on obtaining reprints of this article, please send e-mail to: tkde@computer.org, and reference IEEECS Log Number 108303

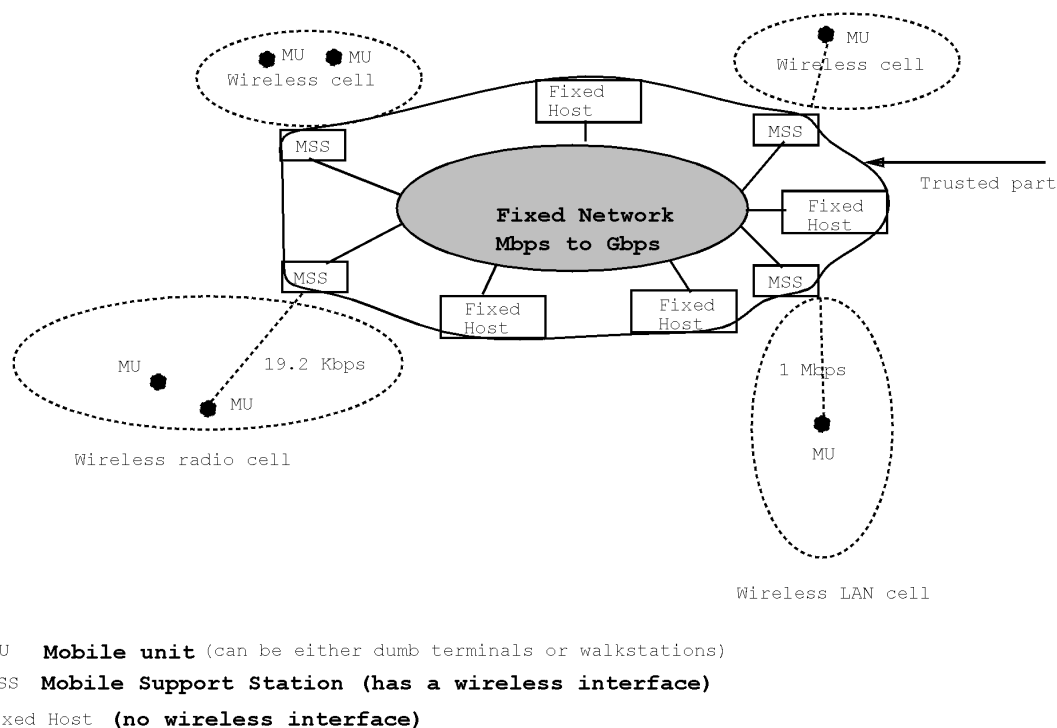


Fig. 1. A mobile environment.

different cells opens the possibility of answering queries in a way that is dependent on the current position of the client, i.e., implementing location dependent queries. We review the work in this area in Section 4. Finally, screen and power limitations have an impact on the kind of interfaces that can be implemented for data browsing and querying. We cover the work in this field in Section 5. Section 6 offers some conclusions about the future of this research area.

2 DATA DISSEMINATION

Data dissemination is the delivery of data from a set of producers to a larger set of clients. This is characterized by an inherent asymmetry in the communications: The bandwidth in the downstream direction (servers-to-clients) is much greater than in the upstream direction (clients-to-servers). (For a survey of data dissemination work see [24].) This model fits very well a mobile system, since mobile clients are usually unable to transmit data at a very high speed (sometimes this capacity is not even present). However, they can receive data at high rate.

The model of sending information to a client population, without waiting for specific requests is also called push-based. By pushing data, the servers avoid interruptions caused by requests and are allowed to propagate data whose existence would otherwise be ignored by clients. On the other hand, push-based systems have the problem of deciding about the relevance of the data to be broadcasted to clients. The usefulness of the system depends on the ability to predict the clients needs. A simple solution is to allow clients to subscribe to services while providing profiles of their interest [49]. The clients subscribe

to information providing queries that describe their interest. These queries can be viewed as continuous queries [62].

The server also needs to decide whether to send the data periodically or aperiodically. Periodic push has the advantage of allowing clients to disconnect for certain periods and still not miss out items (since they will be broadcasted again after the client is reconnected). Aperiodic dissemination on the other hand is a more effective way of using the bandwidth available. Periodic push has been used in the past in many systems [8], [69], [27], [19], [31].

Acharya et al. [5], [1] have proposed the use of a periodic dissemination architecture in the context of mobile systems. They call the architecture *Broadcast Disks*. Broadcast Disks are novel in two ways:

- They provide a multilevel mechanism that permits data items to be broadcast nonuniformly, so more bandwidth can be allocated to items according to their relative importance.
- Mechanisms for managing the storage in the clients are devised to tailor caching and prefetching to perform efficiently with the multilevel broadcast.

Using Broadcast Disks, one can construct a memory hierarchy in which the highest level contains a few items and broadcasts them with high frequency while subsequent levels contain more and more items and broadcast them with less and less frequency. In this way, one can establish a trade-off between access time for high-priority data and that of the low-priority items.

Fig. 2 shows an example in which two levels are provided. One can think of this broadcast as being provided by two disks: the first containing data item *A* and the second items *B1*, *B2*, and *B3*. The first disk rotates at a speed three

A B1 B2 B3 A B1 B2 B3 A B1 B2 B3

Fig. 2. An example of broadcast disks.

times as fast as the second. The characterization of broadcast programs was presented in [5]. A good program renders approximately fixed interarrival times for each item and allocates bandwidth to items according to their access probabilities. An effort in optimizing programs is described in [16].

Since no broadcast program can perfectly match the needs of an individual client, the need to compensate for the mismatches exists. This is done by doing intelligent caching and prefetching at the client site. Caching and prefetching policies for Broadcast Disks have been studied in [5], [1]. The mechanisms have been extended to incorporate the effect of updates in [2]. In this last case, the server must effectively communicate updates to the client by means of invalidation or propagation. Acharya et al. [2] studied the way of devising prefetching algorithms that work well under update scenarios.

Acharya et al. extend their previous work in [3] to study the integration of a pull-based and a push-based Broadcast Disk approach. In a pull-based operation, clients explicitly request items by sending messages to the server, which in turn sends the information back to the clients. Combining the two approaches can be beneficial for both server and clients. Their study assumes two independent data channels, front and backchannel: the frontchannel is used for push-based operation, while the back-channel serves as the media for the pull-based operation. The available bandwidth is shared between these two channels. The study concludes that for the extreme cases of little contention at the server or server overload, pure-pull and pure-push solutions respectively are the best options. For the other cases, the authors propose a hybrid solution, called Interleaved Push and Pull (IPP), which allows clients to use the backchannel to send pull requests for items that do not appear in the Broadcast channel, while the server supports Broadcast Disks plus interleaved responses to the pulls in the frontchannel. IPP suffers of the same bottleneck problems of any pull approach. To improve the scalability of IPP, the authors propose three different techniques:

- 1) Adjusting the pull bandwidth, at the expense of the push bandwidth. This introduces a trade-off between how fast the queue of pull requests is served versus how fast the broadcast delivery is executed. Lower values (30 percent of the available bandwidth) for the pull bandwidth do not produce the best results for lightly loaded systems, but they also do not produce the worse results for saturated systems, and thus they are good candidates for dynamic loads.
- 2) Providing a pull threshold. With this technique, clients only request their most expensive misses, i.e., those that will be broadcast in the longest future. The effect of reducing the number of pulls coming from clients is to delay the point at which the server queue saturates.

- 3) Chopping off the slowest part of the broadcast schedule successively larger pieces. This has the effect of increasing the available bandwidth for pull. This introduces the disadvantage that if there is still not enough bandwidth for pull, performance can degrade considerably since clients will not be able to pull the items that were not broadcast.

In [58], the authors discuss an algorithm to rapidly adapt the content of mobile client's caches based on the misses that result in on-demand data requests. The algorithm is developed as a way to effectively support data dissemination using Direct Broadcast Satellite systems. Their simulation results show that the heuristic of adapting caches based solely on information about the miss ratio works quite reasonably.

In [13], the authors present a novel way of sending invalidation messages over a limited bandwidth network. By using *Invalidation Reports* (IRs), the server can notify clients about changes in the items that are being cached by them. The IRs are very succinct ways of transmitting this information. For instance, one could send a list of identifiers for the items that have changed since the last Invalidation Report was sent. Depending on the technique used, IRs can cause "false negatives," making a client drop an item from its cache when in reality the item is still valid. In return for this anomalous behavior, one gets a shorter report that consumes less bandwidth. In [13], the authors study the trade-offs involved in using different types of Invalidation Reports. In [14], the issue of relaxing consistency of the caches is addressed. For instance if the mobile clients are caching stock prices, it may be perfectly acceptable to use values that are not completely up to date, as long as they are within 0.5 percent of the true prices. This can be accomplished by considering the cached values as *quasicopies* of the values in the server [6]. A quasicopy is a cached value that is allowed to deviate from the true copy in a controlled way. Using quasicopies, the IRs can be made even shorter, thereby saving extra bandwidth for their transmission. Also, in [14], the authors extend the IRs strategies to dynamically adjust to changes in the query and update rates and the disconnection patterns of the mobile clients. These new techniques, called adaptive IRs, take into account the on-demand queries that the clients make when they do not find the items in their local caches. In Wu et al. [70], the problem of discarding caches in mobile computers that have been disconnected for any period of time (sleepers) is addressed. In [13], the strategy a mobile unit that disconnects is to purge its cache after it comes back in operation. Wu et al. propose a mechanism to decide whether some items in the cache can still be used by the mobile unit. To do so, the mobile units need to contact the server when they come back from a disconnection and ask if the items in the cache are still usable. To decrease uplink traffic, the database is partitioned in groups and elements in the same group are cached together. In this way, the mobile unit only asks the server for the validity of specific groups. The server keeps track of the update history for elements in each group, and using an algorithm called Grouping with Cold Update-Set Retention (GCoRe), it determines whether or not all the items updated since the mobile computer

became disconnected have been included in the last IR. If the answer is yes, the cached group can be kept (since the mobile unit can hear about the updates in the IR). Otherwise, the cached group should be discarded (there is not enough information in the IR to determine which elements to keep). Using simulation, the authors determine that as the group size increases, the uplink traffic decreases but the downlink traffic increases. The total bandwidth usage first decreases and then increases suggesting an optimal operating point.

Liu and Maguire [30] use IRs to propose a mobility-aware dynamic database caching scheme for mobile computing. The IR techniques are combined with knowledge of the mobile user’s behavior to broadcast IRs only within the user’s mobility area. Using a performance model, the authors claim the scheme can reduce the system cost by more than 87 percent when compared to a fully replicated database system.

The previous Broadcast Disks protocols are not designed with real-time demands on mind. In real-time environments, minimizing the average latency time is no longer the main criteria for performance. Rather, guaranteeing that time constraints are met is the important concern. Also, in the previous Broadcast Disks protocols, failures are ignored. The basic premise for that is that errors in fetching data as it becomes available in the channel are deemed as tolerable, since the data will become available when it is rebroadcast. For real-time systems, however, this assumption may imply missing critical deadlines. To remedy these two problems, Bestavros [17] proposes a series of Broadcast Disk organizations that are suitable for real-time environments and the usage of an Adaptive Information Dispersal Algorithm (AIDA), which adds reliability to the broadcast.

First, Bestavros presents three different organizations: flat, rate monotonic and slotted rate monotonic which are suitable for real-time environments. In flat organizations, the simplest one of them, the server sends the union of the data items needed by its client in a periodic fashion. With a flat broadcast, there is the need of guaranteeing that the worst-case latency will be tolerable even for those items that have the strictest deadlines. In order to guarantee that, enough bandwidth should be allocated. This organization, is therefore, wasteful of bandwidth. To remedy this, the rate monotonic organization broadcasts each item at a rate inversely proportional to its time constraint. So, each item represents an independent disk “spinning” at a rate reciprocal to the timing constraint associated with the item. This is obviously an extreme organization that tries to optimize bandwidth at the expense of ease of broadcast programming. The slotted rate monotonic organization achieves a balance between the previous two organizations by coalescing data items together so they could be put into the same broadcast disk. This coalescence is achieved by partitioning the data set into bins. Items in the same bin share the same broadcast disk and therefore, have identical worst-case latencies.

Fault-tolerance is achieved by Bestavros using AIDA, an elaboration on the Information Dispersal Algorithm (IDA) of Rabin [53]. Information Dispersal takes a file F and divides

it in N independent pieces (dispersal) in such a way that recombining *any* m of such pieces ($n \leq N$) is sufficient to retrieve F (reconstruction). AIDA uses IDA to adaptively decide the amount of redundancy used in the broadcast. Namely, the number of pieces transmitted n is allowed to vary from m (no redundancy) to N (maximum redundancy). Fig. 3, taken from [17], illustrates how two files A and B are broadcast using IDA. File A is dispersed into 10 blocks, of which any five blocks are enough to reconstruct it. For file B the algorithm uses $N = 6, m = 3$. The pieces are broadcast periodically, using a program as shown in Fig. 4. There are two “periods” in that program: the broadcast period, which extends for eight units and, in the absence of errors is enough to recover both files; and the program data cycle which broadcasts all the pieces from all the dispersed files in the disk. Bestavros finds upper bounds for worst-case delays given the number of transmission errors.

If the data is just broadcast without the any form of directory, clients have to tune to the channel continuously until the required data items are downloaded. This forces clients to be in active mode for a long time, consuming a lot of power, which is already scarce. By providing a directory, the server enables the ability of tuning and that in turn allows clients to become active only when the data of interest is being broadcast. Clients can remain in *doze*




Fig. 3. Dispersal of two files.

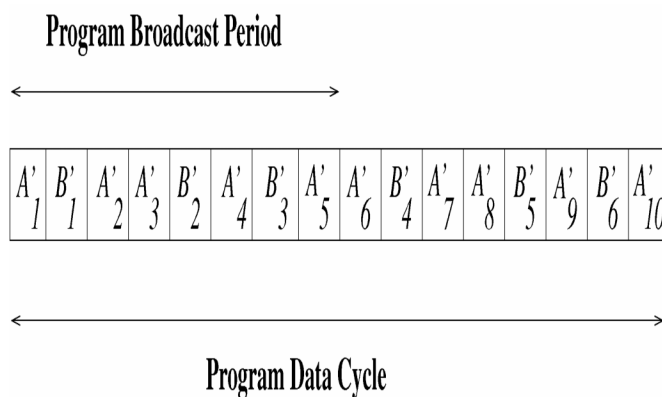


Fig. 4. Broadcast program.

mode most of the time, thereby saving energy. In [36], [60], [37], the authors state the concept of *Indexing on Air*, i.e., transmitting an index along with the data, so clients can tune only during the times they need to. One issue that arises is how the index is multiplexed with the data to make the latency and tuning time minimal. Sending an index with every data broadcast increases the latency since clients that miss the index have to wait for the next broadcast. Broadcasting the index m times during each data broadcast one can minimize the probability of missing the index. This is called $(1, m)$ indexing. The optimal value of m along with the analysis of the technique is provided in [36]. This technique can be further optimized by realizing that there is no need to replicate the entire index between data segments. It is enough to send the portion of the index relevant to the data segment that follows it. This is called *Distributed Indexing* [36]. Distributed indexing results in approximately the best latency and tuning time. It is also possible to have secondary indexes on air, which are described in [60]. All these techniques, however, do not take in consideration the access frequency of different data items, or data skew in the construction of the index. Reducing the average number of index probes can have a significant impact in power consumption, so using the skew to guide the structure of the index is a good idea. This is the motivation of Chen et al. in [20], where imbalanced indexes that favor most frequently accessed items are studied.

The indexing techniques can be used to disseminate data using *temporal addresses*; that is, the client is given the exact time when the relevant data item is going to be published. Temporal matching is performed by the clock circuitry in the client, consuming little energy. However, as a drawback, this method requires a high degree of clock synchronization, forcing the network to be deterministic. This assumption is reasonable for exclusively used channels such as Motorola's Satellite Networks (Embarc) where only one server controls the channel. In general, however this is not easily achieved (consider for instance and Ethernet where the disseminating traffic can be mixed up with on-demand traffic). For these cases, the techniques have to be slightly modified. One solution is to have clients tuning up ϵ slots ahead of time [60].

Alternatively, data dissemination can be achieved by using *multicast addresses* [35]. The server sends data to a group of clients using the same address. In IP, for instance, clients can use 32 bit addresses starting with 1111. Equipped with an Ethernet card, a client can match the predefined set of addresses and keep the CPU in doze mode until the data's multicast address matches the address specified by the client. Clients join multicast groups and filter the relevant data using the Ethernet card. Hashing can be efficiently used in combination with multicast addresses. The client computes the hash value of the data item of interest and sets the interface card to listen to all broadcast values on the multicast address corresponding to the hash value just computed. The hash function can be given to the client upon registration. A randomized algorithm that minimizes the expected latency is described in [35].

If the number of items exceeds the number of multicast addresses, the tuning time for accessing an item depends on the distribution of items among addresses. In [64], an algorithm that optimally partitions the data is provided.

Broadcast schedules to minimize latency have been studied in [63]. The results obtained there can be combined with schemes to reduce tuning time as well. The main contributions of that study are a rule that characterizes the minimal latency and three algorithms to schedule broadcasts. The rule shows that the latency is minimized when the frequency of an item is proportional to the square-root of its demand probability and inversely proportional to the square-root of its size. The first scheduling algorithm is an on-line algorithm that selects the item for which the expression $(T - R(i))^2 p_i / l_i$ is the largest. In the expression, T is the current time, $R(i)$ is the last time at which item i was transmitted most recently, p_i is the demand probability for item i and l_i the size of item i . This decision criteria can be computationally expensive. To decrease the computational demands, another on-line algorithm is proposed in [63]. In this algorithm, items are divided in buckets which are kept as cyclical queues. The algorithm chooses to broadcast the first item in the bucket for which the expression $(T - R(I_m))^2 q_m / d_m$ is the largest. In the expression I_m is the item at the front of the bucket m , while q_m and d_m are the average demand probability and average size for the items in bucket m . Simulations show that the buckets algorithm performs close to optimal, specially as the number of buckets is increased. Notice that the notion of buckets can be compared to the notion of Broadcast Disks [4]. The main differences, however are the facts that the buckets algorithm is on-line and the algorithms in [4] insists in placing the instances of items equally spaced in the disk. Finally an off-line algorithm is introduced in [63] by using an on-line algorithm to generate an a priori schedule of an specified size N .

Some recent applications of data dissemination include information dissemination in the Internet [71], [18], [32], and in private networks [28], [50], Advanced Traveler Information Systems [54], and dissemination using satellite networks [32], [22].

3 DATA CONSISTENCY

Consistency guarantees for data processed by mobile clients are an important area of research. These guarantees provide the basis for any collaborative work and transaction processing done with these systems.

In [61], the authors present a mechanism to provide individual applications with a view of the database that is consistent with their own actions. This is important since in their environment, clients can read or write data from any one of the available servers and these servers can contain inconsistent views of the database. For example, a mobile client of a distributed database could issue a write at one server and later on read the item at a different server, obtaining an inconsistent value if the two servers have not been synchronized between the two operations. The paper introduces *session* guarantees that alleviate this problem. A session is defined as a sequence of

read and write operations performed during the execution of an application by a client. The guarantees provided by their method are:

- *Read Your Writes*: Any read operation in the session must reflect the values established by previous writes in that session.
- *Monotonic Reads*: Successive reads reflect a nondecreasing set of writes.
- *Writes Follow Reads*: Writes are propagated after the reads on which they depend.
- *Monotonic Writes*: Writes are propagated after writes that logically preceded them.

These guarantees are enforced by having a session manager that serializes read and write operations and by assigning unique identifiers to each write. A server must be willing to return information about the unique write identifier and the set of write identifiers for writes that are relevant to a given read. The session manager maintains a read-set of identifiers for the writes relevant to session reads and a write-set of identifiers for those writes performed in the session. For instance, the Read Your Writes guarantee is enforced by following two steps. Whenever a write is accepted by a server, its identifier is added to the session's write-set. Before a read is processed in a server S at a time t , the session manager must check that the write-set is a subset of the sequence of writes received by S before t .

These session guarantees were developed to be deployed in the prototype implementation of the Bayou project, which aims to provide a collaboration framework between mobile computing users.

Krishnakumar and Jain [42] discuss a method to perform mobile sales transactions using site-transaction escrow methods [10], [11], [41], [51]. Escrow methods divide the total number of available instance of an item among the number of sites in the system. A transaction can successfully complete at a site if the number of instances it requires does not exceed the number available at that site. By using a reconfiguration protocol, such as the demarcation protocol [10], sites can negotiate an increase of the allotted fragment. In [42], a mobile client starts performing a transaction in one server, using escrow methods. If the client needs to move to another server during the life of the transaction (because it moved from one geographic area to another), the only information needed at the new server is the operation to be performed next. (Since the old server made its decisions based on escrowed resources.) At the end of transaction, however, a two-phase commit algorithm is needed to synchronize the two servers. This protocol is therefore very suitable for clients that do a fair amount of mobility during the life of the transaction.

In [66], Walborn and Chrysanthis generalize the usage of escrow techniques by exploiting object semantics to facilitate autonomous and disconnected operations in mobile database applications. The basic idea is to split large and complex objects into smaller fragments and cache a set of these fragments on a mobile host. By making these fragments the unit of consistency and defining consistency constraints to be enforced in the fragments, one can have concurrent operations on a set of mobile computing clients.

Using the demarcation protocol [10], one can ensure that local executions are serializable. The protocol does this by exploiting the object organization and constraints. When a mobile client requires access to an object, it sends a request to the database server invoking a *split* operation that selects part of the object and establishes some consistency conditions. This part of the object is then only accessible by transactions on the mobile client that requested it. The rest of the object remains available for other clients. Examples of objects that can be subjected to fragmentation are: aggregate items (e.g., plane seats), sets and stacks. For instance, sets can be split into a number of subsets and these subsets recombined in an arbitrary order to reconstruct the original set. Each mobile host can specify a range of elements as the selection criteria, with the ranges being disjoint.

Gray et al. [29] propose a two-tier replication algorithm that allows mobile applications to propose tentative transactions at the mobile client. These tentative transactions are performed over the data replicas stored at the clients while these clients are disconnected. When the mobile client reconnects, the tentative updates are reprocessed as base transactions in the node that stores the master copy of the data. The tentative transactions may fail when reprocessed (due to conflicts with other transactions). If that happens, the mobile client (and user) that originated the transaction is informed about the failure and why it happened. Processing the transaction in the master node may render different from the tentative results. This may be acceptable for some applications (checking account balances may change because other transactions may have affected the balance) and unacceptable for others (an item is out of stock and cannot be sold). This technique does a good job in solving a fundamental issue: standard ways of propagating updates to replicas (eager replication [26], [25], lazy replication [21], [56]) make deadlocks increase as the cube of the number of sites and as the fourth power of transaction size, thereby rendering the systems unscalable. On the other hand, the two-tier technique may result in an unacceptable number of failed transactions (after the clients reconnect) and user dissatisfaction.

In [9], the concept of *Certification Reports* is introduced as a way of supporting transactions in a mobile environment. The technique uses the broadcast channel to help mobile clients do some of the work of verifying if the transactions being run by them need to be aborted. Clients do that by listening to Certification Reports that contain lists of items that are in the read and write set of active transactions that have declared their intention to commit to the server during the previous broadcast period. If the client sees a conflict between the reports read and writesets and its own read and writesets, it aborts the transaction (this is inspired by optimistic concurrency control algorithms [43]). Otherwise, when the transaction is ready to commit, the client sends to the server the read and writesets along with the values to be written in the database. The server does the final verification to check whether two transactions running on different clients collide with each other. The server then acknowledges accepted and rejected transactions using the broadcast channel. One of the crucial advantages of this protocol is that it downloads most of the work of validating

transactions to the clients. Moreover, transactions which ultimately would be rejected get an early abortion at the client site and can be resubmitted promptly (as opposed to the situation created by the technique in [29]).

Lu and Satyanarayanan [47] propose a new transaction model called *isolation only transactions* (IOT), with which the system can provide strong consistency while not guaranteeing other traditional transaction properties, such as atomicity and durability. The execution model uses optimistic concurrency control [43]. As in the previous techniques, a transaction is executed in the mobile client and no partial execution is visible on the servers. When the transaction completes, it enters a *committed* or *pending* state. If the execution of the transaction does not contain any partitioned data access (i.e., the client maintains server connection for every item that has accessed), the transaction is committed and its results made visible on the servers. Otherwise, the transaction enters the pending state and waits to be validated later. If the validation succeeds, its results are reintegrated and committed in the servers. Otherwise, the transaction is resolved (manually or automatically). In order to do the validation, the technique applies Davidson's optimistic transaction model [23]. This method builds a precedence graph to represent interdependencies among transactions and detects cycles in the graph that are indicative of transactions that do not satisfy global serializability.

The last issue we address in this section is that of data replication. The ability to replicate data in mobile environments is essential since disconnected units can continue to process objects that are stored locally. The authors in [12] point out that even though the number of copies in a system can be large, the key issue is how to manage copies that are *updateable*. These copies are also called the *core* copies and their number is likely to be small for a simple reason: it becomes too expensive to update a large number of core copies. The number of read-only copies (or cached copies) can be large, but these copies may not be current. The core set can be managed using standard quorum algorithms [26] but very often the core set needs to be reconfigured. For instance, a person may be editing a document first from her home computer, then from a laptop on a train ride and then from the office computer. At each step a copy of the document is transferred from one computer to the next. At every step the copy in the computer being used currently is the core set. In [67], Wang and Pâris describe a protocol to facilitate the reconfiguration of core sets. The protocol relies in the introduction of nodes that act as *referees*. These nodes are responsible for the updating and storage of the core set description. Referee nodes do not participate in data accesses and do not hold replicas of the objects. Their sole mission is to record changes in the composition of the core set. Once a replica-holding node decides that it needs to change the core set, it sends the description of the new set to the referees and if a majority of them accept the change, that description becomes the new core. Referees can be highly replicated since they are only accessed when a new core is elected.

As it is pointed out in [12], the directory of objects is also a replicated object that has to be managed just like any other. The authors conclude that standard techniques

to manage copies are also adequate to manage the directory. However, a variation of the primary copy algorithm, called *Primary by Row* is proposed to manage directories. In a Primary by Row algorithm, each core copy has the ability to modify the portion of the directory that concerns the copy. For instance, in the example above, the core copy of the traveling document should be able to modify the row of the directory that points to the current location of the document.

A number of researchers are currently working in providing replication in weakly connected systems, mainly using the Distributed Shared Memory paradigm ([45]). Among them, the work in [59] focuses in providing support for very low-level objects of arbitrary sizes. The authors of [38] describe how to use *Lazy Release Consistency* [39] (which is based on the fact that real consistency is not needed until a synchronization situation is reached), to develop a programming model for large distributed systems with computing nodes that can be weakly connected.

4 LOCATION DEPENDENT QUERYING

The fact that clients in a mobile environment can change locations enables the possibility of answering queries in a way that is dependent on the current position of the client. This is the subject of this section.

Imielinski and Badrinath [33] presented the concept of queries with location constraints, i.e., constraints which involve location of the mobile computing users. For instance the query "find me a doctor near my house" has a unary constraint on location. The query "find X, Y, and Z such they are in the same building and Y is between X and Z" involves a ternary constraint (between) plus three unary constraints involving individual locations (in the same building). The main problem becomes how to minimize the communication cost to retrieve the necessary information to answer the query. Naive strategies result in too many messages and long latencies. On the other hand, building optimal plans can be shown to be NP-complete. The authors suggest using greedy heuristics based on the ID3 [32] algorithm to solve the problem.

Imielinski and Navas discuss a family of protocols that integrate Global Positioning System (GPS)¹ into IP to enable the creation of location dependent services [34]. Examples of these services are: multicasting messages selectively to specific geographical locations (such as a train station), providing services to clients within a certain geographical range from the server and providing information for mobile users when that information depends on the user's location.

Genesis and Advanced Traveler Information Systems (ATIS) [54], [55] are systems being developed to facilitate travel. Travelers are mobile users equipped with personal communication devices and portable computers. ATIS provides updated information on weather and road conditions, detours, construction zones, bus schedules, and parking. This information is available prior and during the trip. The

1. GPS is a radio-navigation satellite-based system developed and operated by the United States Department of Defense. It provides users with their three-dimensional geographical position, their velocity, and time.

users can as a variety of queries, such as recommendations of the most favorable routes from a starting location to a chosen destination. The term favorable may represent minimum travel time, shortest distance, etc. Genesis is an ATIS operational test project developed in Minnesota as a partnership between the University of Minnesota, the Minnesota Department of Transportation, and private sector companies. The Genesis system includes: data collection stations such as the Traffic Management Center (TMC) and Metropolitan Transportation Commission (MTC), the fixed-end ATIS database server which stores the data provided by the TMC and MTC along with the history of travel times, bus schedules and fares and other static information, a wireless communication service provider and mobile personal communication devices. Data is either disseminated by the server as a result of a predetermined transaction or sent at the request of a mobile client. User queries include requests for travel time, roads under construction in a given area, shortest path for a trip, route evaluation and emergency service requests. Events, such as traffic accidents and road hazards are monitored and trigger data dissemination to travelers that need to be warned. Typical triggers warn travelers about congestion and alternative routing.

Mobisaic [65] is a World Wide Web information system designed to serve mobile users. Mobisaic extends the WWW by allowing documents to refer and react to the current location of clients. It does so by supporting two new concepts: *Dynamic Uniform Resource Locators (URLs)* and *Active Documents*. A URL is dynamic if it contains at least a dynamic environment variable that gets substituted with the environment context (e.g., location). For instance a user gets a hyperlink to a document that describes the user's current location in a building by placing a dynamic URL in the corresponding WWW that contains the link. When a user selects a dynamic URL in a document, the client browser is responsible for resolving all references to dynamic environment variables within the URL. After this resolution, the browser has a standard URL which can be sent to the server to get the corresponding page. Active documents are HyperText Markup Language (HTML) documents that enable the client to automatically react to changes in a user's mobile computing context. Authors write active documents just like they write standard HTML documents, with the addition of a subscribe command that lists the dynamic environment variables that the client must subscribe when it loads the document. When the variables listed in the command change value, the client must take an action that reloads the current page (with new variables) or another page (marked by a new URL).

Spreitzer and Theimer [57] describe an architecture to manage location information. In this architecture, personal information is managed by User Agents, while a partially decentralized Location Query Service is used to facilitate location-based operations. There is a User Agent for each user. The agent collects and controls all personal information regarding its user. Applications can only get personal information through a user's agent. The sources of information collection for an agent include infrared-based

active badges [68], GPS, motion sensors and cameras, and human sources.

5 INTERFACES

The screen size of portable units has forced researchers to think about new interfaces for information retrieval.

In [48], Massari et al. present a query processing interface called Query By Icons (QBI), that addresses the issue of screen size along with the limitations in memory and battery power and the restricted communication bandwidth that are characteristic of wireless networks. The features of QBI are:

- An iconic visual language that allows users to compose queries by using a pointing device like a pen to manipulate icons.
- A semantic data model that captures most of the aspects of database structures.
- Metaquery tools that assist on the formulation of queries during periods of disconnection. This is done without having to access the actual data in the server. Data is accessed and transmitted back to the mobile client only when a complete query is answered.

The InfoPad is a project at the University of California, Berkeley that aims to build a light weight, portable multimedia notebook [44], [15]. The InfoPad main features are portability, continuous network connectivity through a high-bandwidth radio link, pen input with handwriting recognition, audio input with speech recognition and full-motion video playback with synchronized audio. One of the goals for the InfoPad is the development of a user interface for application areas such as groupware and mobile information browsing and retrieval [46]. This interface addresses issues such as how to deal with the uncertainty inherent in handwriting and speech recognition and how the pen and voice can be used as substitutes for the mouse and keyboard. Applications may use pen data from the InfoPad in three ways:

- First, regarding the pen as a mouse.
- Second, to get higher resolution data from the pen (an example of this can be a drawing program).
- Third, using handwriting recognition to treat the pen as a keyboard.

Alternatively, applications may use a continuous-word, speaker-independent speech recognizer to process audio data. To increase accuracy, the recognized can be customized to individual speakers.

Alonso and Mani [7], describe an implementation of a pen-based graphical database interface on a pen computer. Using a cellular phone, the interface connects to the remote databases and displays schema information. In response to gestures made by the user, the interface automatically generates queries and retrieves results. For instance, the system presents the user with the relations of a database. The user can pick attributes of interest by simply tapping on them with the pen. Furthermore, the user can ask to display all possible join paths between attributes with another tap of the pen or a *P* gesture. After that a join can be selected by tapping at the corresponding display and the resulting join

table will be displayed along with the corresponding SQL query that represents it. The Universal Relation concept [40] is used to aid in the automatic generation of queries based on the attributes chosen by the user.

6 CHALLENGES

Although the volume of work in this area has been large during the recent future, many challenges still remain. We point to some of them in this section:

- *Prototyping*: In the area of data dissemination, many theoretical studies, simulations, and focused implementations have been tried. However, a full scale prototype that encompasses all of the main ideas is still missing.
- *Bandwidth utilization*: A comprehensive study that analyzes the best way to divide the bandwidth between all the options (data broadcast, indexes, invalidation reports, answers to queries and uplink capacity) is needed. The study in [3] is a good step towards understanding the trade-offs, but more work is needed.
- *Transactional properties*: More work is needed in studying real cases where transactional properties are needed in mobile applications. Moreover, studying which properties can be effectively relaxed (or simply ignored) for real applications, as long as effective protocols to enforce the remaining properties is an area that could have a very practical impact in the field.
- *Optimization of location dependent query processing*: Very little has been done in finding ways to obtain quick, and perhaps approximate, answers to these queries.
- *Data visualization*: A key issue is to effectively use the scarce display space in mobile computers to present answers to queries.

7 CONCLUSIONS


Mobile computing has proven a fertile area of work for researchers in the areas of database and data management. The inherent limitations of mobile computing systems present a challenge to the traditional problems of database management.

As we can see, the amount of research in this area in the last few years have been staggering. However, some problems remain open for research. There is a need for better protocols in the area of data sharing and transaction management, better interfaces, clever algorithms that exploit locality to shape the answers to queries. Undoubtedly, we will continue to see a steady number of research contributions in the future.

REFERENCES

- [1] S. Acharya, M. Franklin, and S. Zdonik, "Dissemination-Based Data Delivery Using Broadcast Disks," *Personal Comm.*, vol. 2, no. 6, Dec. 1995.
- [2] S. Acharya, M. Franklin, and S. Zdonik, "Prefetching from A Broadcast Disk," *Proc. 12th Int'l Conf. Data Eng.*, New Orleans, 1996.
- [3] S. Acharya, M. Franklin, and S. Zdonik, "Balancing Push and Pull for Data Broadcast," *Proc. ACM SIGMOD Int'l Conf. Management of Data*, Phoenix, Ariz., pp. 183-194, May 1997.
- [4] S. Acharya, M. Franklin, S. Zdonik, and R. Alonso, "Disks on Air," *Proc. ACM SIGMOD*, 1994.
- [5] S. Acharya, M. Franklin, S. Zdonik, and R. Alonso, "Broadcast Disks: Data Management for Asymmetric Communication Environments," *Proc. ACM SIGMOD Int'l Conf. Management of Data*, San Jose, Calif., May 1995.
- [6] R. Alonso, D. Barbará, and H. Garcia-Molina, "Data Caching Issues in An Information Retrieval System," *ACM Trans. Database Systems*, vol. 15, no. 3, Sept. 1990.
- [7] R. Alonso and V.S. Mani, "A Pen-Based Database Interface for Mobile Computers," *Proc. First IEEE Workshop on Mobile Computing Systems and Applications*, Santa Cruz, Calif., Sept. 1994.
- [8] M. Ammar and J. Wong, "The Design of Teletex Broadcast Cycles," *Performance Evaluation*, vol. 5, 1985.
- [9] D. Barbará, "Certification Reports: Supporting Transactions in Wireless Systems," *Proc. 17th Int'l Conf. Distributed Computing Systems*, Baltimore, May 1997.
- [10] D. Barbará and H. Garcia-Molina, "The Demarcation Protocol: A Technique for Maintaining Arithmetic Constraints in Distributed Database Systems," *Proc. Int'l Conf. Extending Data Base Technology*, Vienna, 1992.
- [11] D. Barbará and H. Garcia-Molina, "The Demarcation Protocol: A Technique for Maintaining Constraints in Distributed Database Systems," *VLDB J.*, vol. 2, no. 3, 1994.
- [12] D. Barbará and H. Garcia-Molina, "Replicated Data Management in Mobile Environments: Anything New Under the Sun?" *Proc. IFIP WG10.3 Working Conf. Applications in Parallel and Distributed Computing*, Caracas, Venezuela, Apr. 1994.
- [13] D. Barbará and T. Imielinski, "Sleepers and Workaholics: Caching Strategies in Mobile Environments," *Proc. SIGMOD*, May 1994.
- [14] D. Barbará and T. Imielinski, "Sleepers and Workaholics: Caching Strategies in Mobile Environments (Extended Version)," *VLDB J.*, special issue on the best of *SIGMOD System-Oriented Papers*, vol. 4, no. 4, Oct. 1995.
- [15] B. Barringer, T. Burd, F. Burghadt, A. Burstein, A. Chandrakasan, R. Doering, S. Narayanaswamy, T. Pering, B. Richards, T. Truman, J. Rabaey, and R. Brodersen, "Infopad: A System Design for Portable Multimedia Access," *Proc. Calgary Wireless Conf.*, July 1994.
- [16] S. Baruah and A. Bestavros, "Pinwheel Scheduling for Fault-Tolerant Broadcast Disks in Real-Time Database Systems," *Proc. 13th Int'l Conf. Data Eng.*, London, Mar. 1997.
- [17] A. Bestavros, "AIDA-Based Real-Time Fault-Tolerant Broadcast Disks," *Proc. Real-Time Technology and Applications Symp.*, Boston, May 1996.
- [18] A. Bestavros and C. Cunha, "Server-Initiated Document Dissemination for the WWW," *IEEE Data Eng. Bull.*, vol. 19, no. 3, Sept. 1996.
- [19] T.F. Bowen, G. Gopal, G. Herman, T. Hickey, K.C. Lee, W.H. Mansfield, J. Raitz, and A. Weinrib, "The DATACYCLE Architecture," *Comm. ACM*, vol. 35, no. 12, pp. 71-81, Dec. 1992.
- [20] M. Chen, P.S. Yu, and K. Wu, "Indexed Sequential Data Broadcasting in Wireless Mobile Computing," *Proc. 17th Int'l Conf. Distributed Computing Systems*, Baltimore, May 1997.
- [21] P. Chundi, D.J. Rosenkratz, and S.S. Ravi, "Deferred Updates and Data Placement in Distributed Databases," *Proc. 12th Int'l Conf. Data Eng.*, New Orleans, 1996.
- [22] S. Dao and B. Perry, "Information Dissemination in Hybrid Satellite/Terrestrial Networks," *IEEE Data Eng. Bull.*, vol. 19, no. 3, Sept. 1996.
- [23] S.B. Davidson, "Optimism and Consistency in Partitioned Distributed Database Systems," *ACM Trans. Database Systems*, vol. 9, no. 3, pp. 456-481, Sept. 1984.
- [24] M. Franklin, "Special Issue on Data Dissemination," IEEE Technical Committee on Data Engineering, Sept. 1996.
- [25] H. Garcia-Molina and D. Barbará, "How to Assign Votes in A Distributed System," *J. ACM*, vol. 34, no. 4, pp. 841-860, Oct. 1985.
- [26] D.K. Gifford, "Weighted Voting for Replicated Data," *Proc. ACM SIGOPS SOSP Conf.*, Pacific Grove, Calif., Dec. 1979.
- [27] D.K. Gifford, "Polychannel Systems for Mass Digital Communication," *Comm. ACM*, vol. 33, no. 2, Feb. 1990.
- [28] D. Gance, "Multicast Support for Data Dissemination in Orbix-Talk," *IEEE Data Eng. Bull.*, vol. 19, no. 3, Sept. 1996.

- [29] J. Gray, P. Helland, P. O'Neil, and D. Shasha, "The Dangers of Replication and a Solution," *Proc. ACM SIGMOD Int'l Conf. Management of Data*, Montreal, June 1996.
- [30] G.Y. Liu and G.Q. Maguire Jr., "A Mobility-Aware Dynamic Database Caching Scheme for Wireless Mobile Computing and Communications," *Distributed and Parallel Databases*, special issue on databases and mobile computing, vol. 4, no. 3, pp. 271-288, July 1996.
- [31] G. Herman, G. Gopal, K. Lee, and A. Weinrib, "The Datacycle Architecture for Very Large High Throughput Database Systems," *Proc. ACM SIGMOD Conf.*, pp. 97-103, 1987.
- [32] Hughes Network Systems DirectPC, <http://www.direcpc.com>.
- [33] T. Imielinski and B.R. Badrinath, "Querying in Highly Mobile and Distributed Environments," *Proc. 18th Int'l Conf. Very Large Data Bases*, Vancouver, B.C., Canada, Aug. 1992.
- [34] T. Imielinski and J.C. Navas, "GPS-Based Addressing and Routing," Computer Science Dept., Rutgers Univ., Technical Report No. LCSR-TR-262, Mar. 1996.
- [35] T. Imielinski and S. Viswanathan, "Adaptive Wireless Information Systems," *Proc. SIG DataBase Systems Conf.*, Tokyo, Oct. 1994.
- [36] T. Imielinski, S. Viswanathan, and B. Badrinath, "Power Efficient Filtering of Data on Air," *Proc. Fourth Int'l Conf. Extending Database Technology*, Mar. 1994.
- [37] T. Imielinski, S. Viswanathan, and B. Badrinath, "Data on Air: Organization and Access," *IEEE Trans. Data and Knowledge Eng.*, 1997, to appear.
- [38] C. Ionitoiu and B. Andersen, "Replicated Objects with Lazy Consistency," *Proc. Workshop Mobility and Replication, ECOOP '95*, Aug. 1995.
- [39] P. Keleher, A.L. Cox, and W. Zwaenepoel, "Lazy Release Consistency for Software DSM," *Proc. 19th Symp. Computer Architecture*, May 1992.
- [40] H.F. Korth, G.M. Kuper, J. Feigenbaum, A. Van Gelder, and J.D. Ullman, "System U: A Database System Based on the Universal Relation Assumption," *ACM Trans. Database Systems*, vol. 9, no. 3, pp. 331-347, Sept. 1984.
- [41] N. Krishnakumar and A. Bernstein, "High Throughput Escrow Algorithms for Replicated Databases," *Proc. 18th Int'l Conf. Very Large Data Bases*, Vancouver, B.C., Canada, Aug. 1992.
- [42] R.M. Metcalfe and D.R. Boggs, "Ethernet: Distributed Packet Switching for Local Computer Networks," *J. ACM*, vol. 19, July 1976.
- [43] N. Krishnakumar and R. Jain, "Mobility Support for Sales and Inventory Applications," *Mobile Computing*, T. Imielinski and H. Korth, eds., Kluwer, pp. 571-594, 1995.
- [44] M.T. Le, S. Seshan, F. Burghardt, and J. Rabaey, "Software Architecture of the InfoPad System," *Proc. MOBIDATA Conf.*, Rutgers Univ., Oct. 1994.
- [45] K. Li and P. Hudak, "Memory Coherence in Shared Virtual Memory Systems," *Proc. Fifth ACM Symp. Principles of Distributed Computing*, Aug. 1986.
- [46] A.C. Long, S. Narayanaswamy, A. Burstein, R. Han, K. Lutz, B. Richards, S. Sheng, R.W. Brodersen, and J. Rabaey, "A Prototype User Interface for a Mobile Multimedia Terminal," *Proc. 1995 Computer Human Interface Conf.*, May 1995.
- [47] Q. Lu and M. Satyanarayanan, "Improving Data Consistency in Mobile Computing Using Isolation-Only Transactions," *Proc. Fifth IEEE HotOS Topics Workshop*, Orcas Island, May 1995.
- [48] A. Massari, S. Weissman, and P.K. Chrysanthis, "Supporting Mobile Database Access Through Query by Icons," *Distributed and Parallel Databases*, special issue on databases and mobile computing, vol. 4, no. 3, pp. 47-68, July 1996.
- [49] B. Oki, M. Pfluegl, A. Siegel, and D. Skeen, "The Information Bus—Architecture for Extensible Distributed Systems," *Proc. 14th SOSP Conf.*, Asheville, N.C., Dec. 1993.
- [50] B. Oki, M. Pfluegl, A. Siegel, and D. Skeen, "The Information Bus—An Architecture for Extensible Distributed Systems," *Proc. 14th Symp. Operating Systems Principles*, Asheville, N.C., Dec. 1993.
- [51] P.E. O'Neil, "The Escrow Transactional Model," *ACM TODS*, vol. 11, no. 4, pp. 405-430, Dec. 1986.
- [52] R. Quinlan, "Induction of Decision Trees," *Machine Learning*, pp. 81-106, 1986.
- [53] M.O. Rabin, "Efficient Dispersal of Information for Security, Load Balancing, and Fault Tolerance," *J. ACM*, vol. 36, no. 2, pp. 335-348, Apr. 1989.
- [54] S. Shekhar, A. Fetterer, and D. Lui, "Genesis: An Approach to Data Dissemination in Advanced Traveler Information Systems," *IEEE Data Eng. Bull.*, vol. 19, no. 3, Sept. 1996.
- [55] S. Shekhar and D. Lui, "Genesis and Advanced Traveler Information Systems," *Mobile Computing*, T. Imielinski and H. Korth, eds., Kluwer, pp. 699-720, 1996.
- [56] J. Sidell, P.M. Aoki, S. Barr, A. Sah, C. Staelin, M. Stonebraker, and A. Yu, "Data Replication in Mariposa," *Proc. 12th Int'l Conf. Data Eng.*, New Orleans, 1996.
- [57] M. Spreitzer and M. Theimer, "Providing Location Information in a Ubiquitous Computing Environment," *Mobile Computing*, T. Imielinski and H. Korth, eds., Kluwer, pp. 396-423, 1996.
- [58] K. Stathatos, N. Roussopoulos, and J.S. Baras, "Adaptive Data Broadcasting Using Air-Cache," *Proc. First Int'l Workshop Satellite-Based Information Services*, Rye, N.Y., Nov. 1996.
- [59] P.S. Knudsen, M.O. Svendsen, and P.T. Koch, "Replication of Low-Level Objects in a Distributed Shared Memory," *Proc. Workshop Mobility and Replication, ECOOP '95*, Aug. 1995.
- [60] T. Imielinski, S. Viswanathan, and B.R. Badrinath, "Energy Efficient Indexing on Air," *Proc. SIGMOD '94*, May 1994.
- [61] D. Terry, A.J. Demers, K. Petersen, M.J. Spreitzer, M.M. Theimer, and B.B. Welch, "Session Guarantees for Weakly Consistent Replicated Data," *Proc. Conf. Parallel and Distributed Computing*, Austin, Texas, Oct. 1994.
- [62] D. Terry, D. Goldberg, D. Nichols, and B. Oki, "Continuous Queries Over Append-Only Databases," *Proc. ACM SIGMOD Int'l Conf. Management of Data*, San Diego, June 1992.
- [63] N.H. Vaidya and S. Hameed, "Data Broadcast Scheduling: On-Line and Off-Line Algorithms," Dept. of Computer Science, Texas A&M Univ., Technical Report No. 96-017, July 1996.
- [64] S. Viswanathan and T. Imielinski, "Wireless Data Access Using Multicasting," *Proc. Seventh Int'l Conf. Management of Data (COMAD)*, Pune, India, Dec. 1995.
- [65] G.M. Voelker and B.N. Bershad, "Mobisaic: An Information System for a Mobile Wireless Computing Environment," *Mobile Computing*, T. Imielinski and H. Korth, eds., Kluwer, pp. 375-395, 1996.
- [66] G. Walborn and P.K. Chrysanthis, "Supporting Semantics-Based Transaction Processing in Mobile Database Applications," *Proc. 14th IEEE Symp. Reliable Distributed Systems*, Germany, Sept. 1995.
- [67] Q.R. Wang and J.F. Páris, "Managing Replicated Data Using Referees," *Proc. Workshop Mobility and Replication, ECOOP '95*, Aug. 1995.
- [68] R. Want, A. Hopper, V. Falcao, and J. Gibbons, "The Active Badge Location System," *Trans. Information Systems*, vol. 10, no. 1, Jan. 1992.
- [69] J. Wong, "Broadcast Delivery," *Proc. IEEE*, vol. 76, no. 12, Dec. 1988.
- [70] K. Wu, P.S. Yu, and M. Chen, "Energy-Efficient Caching for Wireless Mobile Computing," *Proc. 12th Int'l Conf. Data Eng.*, New Orleans, pp. 336-343, Feb. 1996.
- [71] T. Yan and H. Garcia-Molina, "SIFT—A Tool for Wide-Area Information Dissemination," *Proc. Usenix Technical Conf.*, 1995.



Daniel Barbará received his BSEE degree from Universidad Metropolitana in 1970 and the PhD degree in computer science from Princeton University in 1985. He is now an associate professor in the Department of Information and Software Engineering at George Mason University in Virginia. His current areas of research interest include data warehousing, mobile systems, and data mining.