

Stored Geocast

Christian Maihöfer, Walter Franz, Reinhold Eberhardt

DaimlerChrysler Research And Technology,
Communication Technology (RIC/TC)
P.O. Box 2360, 89013 Ulm, Germany
{christian.maihoefer|walter.franz|reinhold.eberhardt}@daimlerchrysler.com

Abstract. For many location based services like location based advertising, location based information services, or in particular, realizing a virtual traffic sign, which is bound to a location over time, we need a stored geocast solution. A stored geocast is a time stable geocast, which is delivered to all nodes that are inside a destination region within a certain period of time. The period of time starts at the time of sending plus network propagation delay and ends at a user defined time or is boundless.

This paper discusses the design space, the semantics, and three reasonable solutions for stored geocast in an ad hoc network. The first approach is an infrastructure-based server solution to store the messages. The second is an infrastructure-less approach that may be more suitable for an ad hoc network. A node inside the geocast destination region is elected to act temporarily as a server for geocast messages. The last approach works infrastructure-less, too, by complementing the exchange of neighbor information necessary for many geographic unicast routing protocols with geocast information.

1 Introduction

Geocast, i.e. the transmission of a message to some or all nodes within a geographical area, allows promising new services and applications. Of particular interest is geocast in the automotive domain. For example in the FleetNet project an ad hoc multi-hop radio network for inter-vehicle communications is developed [1]. Such a vehicle network allows safety-related and information-related applications like collection of traffic states, accident warning, wrong-way driver warning, icy road warning and so forth. Since these applications naturally address all (anonymous) vehicles in an area rather than a single vehicle by its fixed and known address, geocast is necessary to transmit the messages to the corresponding area.

Although the literature proposes several geocast approaches for infrastructure as well as ad hoc networks (see next section for details), we miss one important service for our applications, which we call stored geocast. The following example tries to clarify this. Assume that we want to inform following vehicles about an accident in front of them to avoid rear-end collisions. Currently, warning other vehicles is possible either by traffic signs or police presence or by a warning message over the radio, respectively. A radio warning message has the disadvantage that it is not limited to the vehicles of the area in which the danger

spot is located and possibly forgotten by a driver when he actually enters this area. Therefore, we aim at realizing a warning function which is similar to a conventional traffic sign or police presence, i.e. bound to a geographic location or area. All vehicles which enters such an area, possibly by additional considering their driving direction and other parameters, should receive the warning message.

Currently, geocast solutions for ad hoc networks provide only a means to send a message once, instead of periodically or on-demand every time a mobile node enters the geocast's message destination region. In this paper we discuss the design space, the semantics, and three approaches how such a stored geocast can be realized.

The remainder of this paper is structured as follows. In the next section background of geocast and related work is introduced. Section 3 discusses the semantics of stored geocast. In Section 4 the design space is defined before in Section 5 three approaches are described. Finally, the paper is concluded with a brief summary.

2 Background and Related Work

Mobile ad hoc routing protocols can be classified into topology- and position-based or geographic approaches [2]. Topology-based approaches use only information about existing neighborhood links rather than additional physical (geographical) position information of the participating nodes. Topology-based approaches can be further divided into table driven and source-initiated on-demand driven protocols [3]. Basically, table driven protocols attempt to maintain consistent and up-to-date routing information among all nodes, while source-initiated on-demand driven protocols create routes only when necessary to deliver a packet. Therefore, the former approach is also known as a proactive routing approach and the latter as a reactive approach.

In this paper, only geographic routing approaches are of interest since first, position information for vehicles are available by their navigation system, which promises more efficient routing schemes, and second, many applications explicitly address their (anonymous) communication peers by their position rather than by their identifier.

In contrast to topology-based routing protocols, geographic routing protocols usually refrain from setting up routes to forward packets, which decreases overhead. Instead, the forwarding decision of a node is based on the destination's position and the position of the forwarding node's neighbors (i.e. nodes with a one hop distance).

For geographical unicast routing protocols, three basic forwarding strategies can be identified: 1) greedy forwarding, 2) restricted directional flooding, and 3) hierarchical forwarding. With greedy forwarding a node forwards a packet to a neighbor that is located closer to the destination. If this forwarding strategy fails, since there may be situations in which there is no closer node to the destination than the forwarding node, recovery strategies have to deal with it. Nearest with Forward Progress (NFP) [4] is an example protocol using a greedy forwarding scheme.

The second approach, exemplified by the DREAM protocol [5], is similar to the first approach with the modification that a packet is forwarded to some neighbors rather than to just one neighbor. The third approach tries to improve scalability by forming a hierarchy of non-equal nodes. The Grid protocol [6] uses such a hierarchy for message forwarding.

Besides the unicast delivery described so far, the following approaches allow geocast addressing and routing. We refer to the *destination region* of a geocast packet as the geographical area to which a packet has to be delivered. Geocast protocols belong to one of the classes: 1) directed flooding, or 2) explicit route setup approaches without flooding.

Geocast directed flooding approaches are quite similar to the unicast directed flooding approaches. They define a forwarding zone, which comprises a subset of all network nodes. The forwarding zone includes at least the target area and a path between the sender and the target area. An intermediate node forwards a packet only if it belongs to the forwarding zone. If the target area is reached, they differ from unicast approaches, since they apply a flooding of the whole target area. A node broadcasts a received packet to all neighbors provided that this packet was not already received before and that the node belongs to the target area. Finally, a node accepts a packet and delivers it to its application if the own location is within the specified target area. Examples of geocast directed flooding protocols are Location Based Multicast (LBM) [7] and GeoGRID [8].

The second geocast scheme, explicit route setup without flooding, requires either a fixed network like the Internet, which is exemplified by the GeoNode approach [9]. Or, in the GeoTORA approach [10], for each geocast group a directed acyclic graph comprising all network nodes is maintained, which shows the routing direction to the destination. These acyclic graphs are initially created with a flooding scheme, too. However, their maintenance is achieved without flooding.

GeoNode is the only geocast approach that store messages for periodical delivery such that a stored geocast can be realized. Their assumption is that the network has a fixed cellular architecture with a GeoNode assigned to each cell. Routing is done in two steps, the first between sender and GeoNode and the second between GeoNode and destination region. GeoNodes are able to store the packets they receive for periodical delivery.

3 Semantics of Stored Geocast

Before we present approaches to realize stored geocast, we briefly discuss how to define the semantics, which is especially important if safety-related services have to be realized with a stored geocast solution.

We see two possible approaches to define the semantics. First we can consider the application side, define the semantics necessary for our intended applications and then try to find solutions for this semantics. Or we can define solutions for a broad range of geocast approaches and then analyze the solutions and define their semantics. As we believe that we cannot find all possible applications with their semantics and that likely their intended semantics would be too rich and restricting and would result in fat solution approaches, we follow the second

approach. This means, that the detailed semantics of stored geocast is defined together with the approaches in Section 5. Note that this is similar to many other research areas, for example reliable multicast. However, some of the characteristics are identical for all solutions.

First we have to dissociate stored geocast from reliability mechanisms. Although a stored geocast is bound to an area over time, our proposed solutions do not try to achieve reliability. For some applications, in particular safety-related applications, reliable stored geocast is desirable or even mandatory. However, this is a difficult task especially in ad hoc networks, where message transmission failures and network partitions are quite normal rather than unusual. Furthermore, in many cases reliability can be achieved simply by building on a reliable geocast protocol or using general reliability mechanisms on top of an unreliable stored geocast. Though this might not be the most efficient solution, it results in better structured and more modular mechanisms. Therefore, reliability mechanisms are not discussed in this paper.

An inherent question of the stored geocast semantics is the duration of storage and delivery availability. As discussed above, we provide an best effort service without guarantees, which means that we cannot provide guarantees about the duration availability, i.e. we cannot guarantee to reach the full lifetime. However, we assume that we have a mechanism in place to limit the lifetime to a user defined time.

Besides the natural definition of lifetime corresponding to physical clock time, it is possible to define lifetime based on some sort of hop count, similar to the IP approach or some sort of delivery count. For stored geocast it would make sense to limit the number of deliveries or the total number of hops (we will see that for some approaches a stored geocast has to hop in order to keep stored). Another approach would be to limit the lifetime by an opposing event, e.g. in our traffic scenario from the introduction a discard operation after a traffic jam has disappeared. As the exact definition lies outside the scope of this paper, we generically assume that a lifetime of some sort is added to a stored geocast message before it is sent. If this lifetime exceeds, the stored geocast message is discarded.

4 Design Space of Stored Geocast

In this section we discuss and structure the design space of stored geocast before we present geocast approaches in the following section. For a stored geocast solution we identify four building blocks: 1) the underlying geocast routing protocol, 2) the storage of geocast messages within their lifetime, 3) the hand over of stored geocast messages to other nodes, and 4) the delivery of geocast messages to their intended destination nodes.

The underlying geocast routing protocol is necessary for most approaches to deliver the first geocast message to its destination region and possibly for the delivery of all following geocast messages to new nodes entering the destination region later.

The second building block is the storage of geocast messages, which can be done either infrastructure-based by a central server or infrastructure-less, which means distributed on some or all nodes participating in the network.

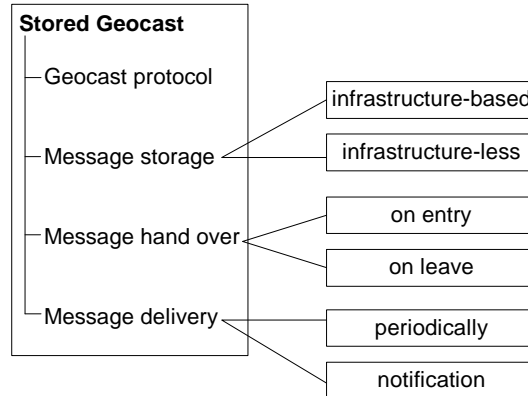


Fig. 1. Design space of stored geocast

With the third building block, the hand over of stored geocast messages, we refer to the problem that a node used for message storage may change its state so that it is no longer considered a suitable node and transfer the message to another, suitable node. For example, the principle of locality may make it desirable to store a geocast message only on nodes inside the destination region. If a storing node leaves the destination region, the stored message is then transferred to another node which is inside the destination region. Transferring a message to another node can be triggered when a new node enters the destination region or when a node inside the destination region is going to leave it.

Finally, the last building block is the delivery of a stored geocast message to new nodes inside the destination region. This can be done either by blind periodical resending of stored geocast messages or on demand, by a notification scheme, when a new node enters the destination region. Figure 1 summarizes the design space. Basically, these four building blocks have to be combined to realize stored geocast. In the next section we will discuss three reasonable combinations.

5 Stored Geocast Approaches

5.1 Infrastructure Approach

Overview: A server is used to store geocast messages. Hand over of messages is not necessary. Message delivery is done periodically or by notification.

Description: The geocast message is first unicasted to the geocast server provided by the infrastructure. Then the geocast server uses a geocast routing protocol to deliver the message to the destination region. After the first delivery, further deliveries can be done either periodically by the geocast server or by notification from moving nodes. Note that this approach has some similarities to the geocast proposal of [11].

If the server periodically delivers the geocast message, the delivery frequency has to depend on the maximum or average velocity of the network's nodes. For example, in a vehicular scenario with a maximum velocity of 200 km/h (= 55 m/s) and a circular destination region of a geocast message with a diameter of 1 km a message frequency of $1000 \text{ m} / 55 \text{ m/s} = 18$ seconds is required. With this message frequency, based on the maximum velocity, it can be ensured that every node is able to receive a geocast message, provided that no message loss occurs and provided that the nodes cross the region through the center of the stored geocast message's destination region.

If the message frequency is increased further, a certain number of message losses can be tolerated since the probability for receiving at least one of the messages increases. A second advantage of increased message frequency is that nodes receive the geocast message even if they do not cross the full diameter range of the geocast destination region. Assume that the message frequency is denoted as f , the maximum velocity as v and that the shortest crossing distance within the geocast region a node has to cover in order to receive the stored geocast message is denoted as c , the following holds:

$$f \geq \frac{c}{v}, 0 < c \leq s.d \quad (1)$$

s denotes the stored geocast message and $s.d$ the diameter of the destination region. If the message frequency is decreased below the obtained f , fast moving nodes and nodes crossing less than the assumed distance c within the geocast destination region may cross it without receiving a geocast message. Note that besides configuring a quite high message frequency it is possible to increase the destination region of a geocast message beyond its actually intended region. This is necessary if $c = 0$ is required, i.e. the delivery of a geocast message to a node which only touches the geocast destination region (see Figure 2). Assume that the increased destination diameter of the geocast message is denoted as $i.d$, then the following holds:

$$f \geq \frac{i.d - s.d + c}{v}, i.d > s.d, 0 \leq c \leq s.d \quad (2)$$

Besides the periodic delivery a notification from moving nodes can trigger the message delivery. However, as a moving node does not know about the defined destination region of geocast messages stored on the server, this requires moving nodes to periodically send their position to the geocast server. To realize an efficient location notification approach, it should depend on the distance between the current position and the position of the last report to the server rather than on time. The distance between two reports d has to be not greater than the minimum required crossing distance c within a geocast's destination region (which has to be not greater than the minimum diameter of geocast destination regions):

$$d \leq c, 0 < c \leq s.d \quad (3)$$




Fig. 2. Increased geocast destination diameter

Like in the periodic delivery scheme from above, increasing the destination region of geocast messages beyond their actually intended region can help to decrease the required message frequency:

$$d \leq i.d - s.d + c \quad , i.d > s.d, 0 \leq c \leq s.d \quad (4)$$

A consequence of the location notification approach is that the geocast server cannot be implemented stateless, since it has to remember requesting nodes and already delivered geocast messages to these nodes.

Coming back to our vehicular scenario from above, the notification approach requires that all vehicles report their location while moving. Assuming an average velocity of 60 km/h (≈ 17 m/s) and a minimum geocast crossing distance of 1 km, the location notification frequency of a random vehicle has to be one message every 1000 m or $1000 \text{ m} / 17 \text{ m/s} \approx 60$ s. Note that this message frequency may result in a significant overall network load, since all moving nodes have to send location notifications.

A reasonable optimization of the notification approach would be to suppress location notifications if a stored geocast message for the current location is received before, which means that another node has recently reported its presence in the same region. Another optimization would be to synchronize several nodes with similar movement patterns and to send just a single location notification for the synchronized group. Finally, the location information of a node could be sent to the server and additionally as a geocast message to the surrounding of the current location in order to suppress further location notifications of other nodes in the same region.

The decision about the most adequate scheme depends on the frequency of node movements. If nodes move frequently, the node penetration is high, or only few stored geocast messages with small destination regions are active, the periodic sending scheme may be more efficient, otherwise, the notification scheme. Note that with position calculations of dead reckoning approaches [12], especially in vehicular environments with its predefined routes, overhead of too frequent message delivery may be further decreased.

In summary, the infrastructure approach offers a simple and robust mechanism for stored geocast. One disadvantage, the large communication distance between the server and the destination region of geocast messages can be relaxed by distributed geocast servers close to the destination region of their stored messages. Note that a large distance results not only in high overhead but also in low robustness, especially in ad hoc networks where network partitioning and message loss may occur frequently and successful delivery of messages incorporating too many hops become unlikely.

Semantics: We assume that the server is replicated. In the absence of communication failures this approach guarantees that the geocast message is not lost and delivered to all intended nodes. However, note that communication failures are likely to occur due to the usually large distance between server and destination region. This can be relaxed by distributed geocast servers close to their destination regions.

5.2 Election Approach

Overview: A node in the destination region of a geocast message is elected to store geocast messages. Hand over of messages is done when this node leaves the destination region. Message delivery is done periodically or by notification.

Description: Instead of relying on a fixed server infrastructure, a dynamically elected node within the destination region of the geocast message is responsible for storing and delivering the message.

Basically, each node in the destination region is a candidate for the election process. However, to avoid frequent hand over, it is desirable to choose one that stays as long as possible in the destination region. Such a node is characterized by low velocity and closeness to the center of the destination region. For example, the unique tuple $\langle \text{velocity} * \text{center distance}, \text{nodeid} \rangle$ can be used in the election process. A suitable election algorithm is for example described in the GeoGRID approach [8].

Geocast message delivery is done as follows. The initial sender of a geocast message uses a geocast routing protocol to deliver the message for the first time. Inside the destination region, all nodes receive the geocast message and start the election process. The elected node stores the message and periodically or on request delivers the message as in the previous infrastructure-based approach. In case of periodical delivery, our calculations from the previous approach are effective for the election approach, too. In case of on request delivery, the location notification report is sent as a geocast message to a circular destination region with the actual position as the center.¹ In Figure 3 an example for the message request and delivery is given.

The configuration of important parameters like diameter of the location notification geocast and frequency of location notifications requires consideration.

¹ Note that this is feasible for a random walk scenario. In case of a directed walk like in a vehicular scenario, optimizations by sending a geocast location notification to the region in front of the vehicle may be worthwhile.




Fig. 3. Example of on request stored geocast delivery with the election approach

A basic observation is that the diameter of the location notification $n.d$ must be no smaller than the doubled maximum diameter of the geocast messages $s.d$:

$$n.d \geq 2 \cdot \max(\forall s \in \{\text{stored geocasts}\} : s.d) \quad (5)$$

With $n.d$ smaller than $2 \cdot s.d$ it would be possible to miss the elected node as depicted in Figure 4.

Besides the destination region's diameter of location notifications, the frequency of location notifications is of interest. The distance between two reports d has to be not greater than the minimum required crossing distance c within a geocast's destination region similar to the infrastructure approach possibly relaxed by an increased geocast destination region $i.d$:

$$d \leq c, 0 < c \leq s.d \quad (6)$$

$$d \leq i.d - s.d + c, i.d > s.d, 0 \leq c \leq s.d \quad (7)$$

Location notification suppression schemes as briefly discussed in the infrastructure approach are feasible, too. The geocast based location notification makes it quite simple to suppress an own location notification if another one for the same region has been received before.

Finally, if the elected node leaves the destination region, a new election round is started and the message is handed over to the new elected node. Fault tolerance can be increased by electing not only a single node but several ones which keep message replicas.

Semantics: In the absence of node and communication failures and given that at any time at least one node is within the destination region and that there




Fig. 4. Small notification diameter misses the elected node

are no network partitions inside the destination region, the election approach guarantees that the geocast message is not lost and delivered to all intended nodes. If there is not always at least one node within the geocast region, only the periodic sending scheme can give the same guarantee provided that a leaving elected node keeps sending geocast messages until another node is elected. The notification scheme fails, since the notification report addressed to the geocast region may fail to reach the elected node which is now outside the destination region. In the presence of a node failure of the elected node, the geocast message may get lost. However, robustness can be increased by replicated elected nodes.

5.3 Neighbor Approach

Overview: Each node stores all geocast packets destined for its location and keeps a table of all neighbor nodes and their location. If a node within a geocast destination region detects a new neighbor it delivers the geocast packet to it, i.e. hand over is done on entry and message delivery is done by notification.

Description: The initial stored geocast message is sent using a regular geocast routing protocol to the destination region. After the first delivery, geocast information is exchanged between neighbors inside the destination region.

Many location-based unicast routing protocols like NFP [4] or DREAM [5] proactively and periodically exchange neighbor information containing their location in order to forward a packet to a neighbor closer to the destination. The neighbor approach simply extends the exchanged neighbor information for the unicast routing with stored geocast information. The following alternative schemes to extend the exchanged neighbor information are reasonable: 1) with neighbor information all stored geocast messages relevant for this location are exchanged, 2) with neighbor information a list of already received stored geocast identifiers relevant for this location is exchanged, or 3) with neighbor information not only the current location but also the last (or some last) reported location information is exchanged.

The first scheme, blindly exchanging stored geocasts with all neighbors, is the most simple but also most bandwidth wasteful one. After receiving a stored geocast, a node has to check whether its location intersects with the geocast destination region. If so, the geocast is delivered to the higher protocol layer and stored for later exchange with neighbors. Otherwise, the geocast is discarded.

With the second scheme, filtering of geocast messages is done before they are exchanged. The exchanged list of geocast identifiers contains unique tuples, for example $\langle \text{initial geocast sender, sequence number} \rangle$ to identify a geocast message. If a node detects that it has stored a geocast message relevant for a neighbor node's location but unknown to the neighbor, the stored geocast is sent to it.

Instead of exchanging geocast messages or their unique identifiers, the third scheme derives a necessary geocast forwarding from the route a node has taken. If the route shows, that the neighbor has 'just recently' entered a geocast destination region, then the geocast message is delivered to the neighbor node. Note that it is difficult to define 'just recently' in more detail. Of course, if this is the first neighbor information exchange inside the destination region, the geocast packet has to be delivered. However, for a second or later neighbor information exchange there is no guarantee that the geocast was already delivered, since it is possible that all former neighbors were located outside the destination region or that two neighbors have entered the destination region simultaneously and hence both have no knowledge about existing stored geocasts. As a consequence, the concrete definition of 'just recently' is a compromise between robustness and overhead and lies outside the scope of our description.

As an optimization of all three schemes, a two round protocol can be introduced. In the first round, only information necessary for most location based unicast routing protocols is exchanged with neighbors. This includes the own identifier and location. If this results in detecting a new neighbor inside a geocast's destination region, a second information exchange round is triggered including the stored geocast information according to one of the three schemes from above.

Semantics: We assume the absence of node and communication failures. Given that at any time at least one node is within the destination region and that there are no network partitions inside the destination region possibly preventing a node from having neighbors, a stored geocast message is not lost and delivered to all intended nodes. Note that the frequency of exchanging neighbor information of the complemented unicast protocol must be adapted according to the obtained distance d from the previous approaches with notification schemes.

6 Summary

Stored geocast is a time stable geocast delivered to all nodes which are inside a destination region within a certain period of time. The period of time starts at the time of sending plus network propagation delay and ends at a defined time or is boundless. We have motivated the need for stored geocast solutions in ad hoc networks with a vehicular scenario, in which a stored geocast is able to realize a virtual traffic sign.

The design space of stored geocast comprises the four dimensions: geocast protocol, geocast storage, geocast hand over and geocast delivery. From the many possible solutions we have selected and described three reasonable ones.

The first approach is an infrastructure-based server solution to store the messages. The second is an infrastructure-less approach which elects a node inside the geocast destination region to act temporarily as a server. The last approach works infrastructure-less, too, by complementing the exchange of neighbor information necessary for many geographic unicast routing protocols with geocast information.

Future work of stored geocast will include performance evaluations of the protocols. Note that as usual for ad hoc protocols, we do not expect to identify an always superior protocol. Instead, the scenario will decide the comparison.

References

1. Franz, W., Eberhardt, R.: Fleetnet - internet on the road. In: Eight World Congress on Intelligent Transport Systems, Sydney, Australia (2001)
2. Mauve, M., Widmer, J., Hartenstein, H.: A survey on position-based routing in mobile ad hoc networks. *IEEE Network* **15** (2001) 30–39
3. Royer, E., Toh, C.: A review of current routing protocols for ad-hoc mobile wireless networks. *IEEE Personal Communications* **6** (1999) 46–55
4. Hou, T., Li, V.: Transmission range control in multihop packet radio networks. *IEEE Transactions on Communication* **34** (1986) 38–44
5. Basagni, S., Chlamtac, I., Syrotiuk, V.R., Woodward, B.A.: A distance routing effect algorithm for mobility (DREAM). In: Proceedings of the ACM/IEEE International Conference on Mobile Computing and Networking (MobiCom), Dallas, USA, ACM Press (1998) 76–84
6. Liao, W.H., Tseng, Y.C., Sheu, J.P.: GRID: A fully location-aware routing protocol for mobile ad hoc networks. *Telecommunication Systems* **18** (2001) 37–60
7. Ko, Y.B., Vaidya, N.H.: Geocasting in mobile ad hoc networks: Location-based multicast algorithms. In: Proceedings of the 2nd Workshop on Mobile Computing Systems and Applications (WMCSA 99). (1999)
8. Liao, W.H., Tseng, Y.C., Lo, K.L., Sheu, J.P.: GeoGRID: A geocasting protocol for mobile ad hoc networks based on GRID. *Journal of Internet Technology* **1** (2000) 23–32
9. Imielinski, T., Navas, J.: GPS-based addressing and routing. Internet Engineering Task Force, Network Working Group, Request for Comments, RFC 2009
URL: <ftp://ftp.isi.edu/in-notes/rfc2009.txt> (1996)
10. Ko, Y.B., H.Vaidya, N.: GeoTORA: A protocol for geocasting in mobile ad hoc networks. In: Proceedings of the 8th International Conference on Network Protocols (ICNP), Osaka, Japan (2000) 240–250
11. Imielinski, T., Navas, J.C.: GPS-based geographic addressing, routing, and resource discovery. *Communications of the ACM* **42** (1999) 86–92
12. Wolfson, O., Sistla, A.P., Chamberlain, S., Yesha, Y.: Updating and querying databases that track mobile units. *Distributed and Parallel Databases Journal* **7** (1999) 1–31