Mobicast: Just-in-Time Multicast for Sensor Networks under Spatiotemporal Constraints

Qingfeng Huang, Chenyang Lu, and Gruia-Catalin Roman

Department of Computer Science and Engineering
Washington University, Saint Louis, MO 63130.
{qingfeng,lu,roman}@cse.wustl.edu

Abstract. This paper is concerned with coordinated delivery of messages in sensor networks. The notion of multicast is re-examined in light of a new set of requirements that are specific to such networks. The result of this investigation is a new concept called *mobicast*. It entails the delivery of messages to large sets of nodes in a manner that satisfies a potentially dynamic set of spatiotemporal constraints. In order to demonstrate the feasibility of *mobicast*, we present a novel topology-aware protocol for sensor networks. Worst-case analysis shows that the protocol provides strong spatial and temporal delivery guarantees under a set of reasonable assumptions about the network. The design of the protocol relies on new notions of compactness for spatially distributed networks. By explicitly addressing the temporal domain associated with message delivery, *mobicast* is more general than geocast and makes it possible to save precious resources in sensor networks by exploiting its inherent just-in-time delivery semantics.

1 Introduction

Large-scale wireless sensor networks will be deployed in various physical environments to support a broad range of applications such as precision agriculture, smart highway, security, emergency response and disaster recovery systems [1]. These applications need to collect data from sensor networks, aggregate data from multiple sensors inside the network, and communicate aggregated information to end users over multi-hop ad hoc networks. Due to the need for high data fidelity and the severe energy constraint in sensor networks, in-network data aggregation has recently received significant attention [2–4]. While some forms of data aggregation can be performed on the end-to-end route from the source to the base station [2, 4], explicit group coordination among sensors in the locality of a monitored physical entity (e.g., an intruder) are needed by many applications. In the latter case, a group management protocol maintains a sensor group in the vicinity of a physical entity, and a multicast or unicast protocol provides the communication mechanism for data aggregation inside the group.

Local coordination is often subject to spatiotemporal constraints due to mobility in the physical environment. Environmental mobility, i.e., the movement of monitored physical entities, is common to many sensor network applications (e.g., personnel tracking in emergency sites, mobile robots in factories, and habitat monitoring of wildlife). To illustrate the kind of spatiotemporal constraints likely to be encountered in such applications, let us consider the deployment of acoustic sensors in a security area designed to track intruders. When there are no intruders, most sensors sleep and only periodically wake up to check for interesting events. A small number of sensors remain active to provide continuous vigilance and to activate other sensors when necessary. To track an intruder, sensors in its vicinity form a group to share their data and determine the location of the intruder through triangulation. Only the sensors within the vicinity of an intruder should contribute data for the triangulation operation. It is unnecessary and even incorrect to aggregate the data from sensors that are far away from the intruder because their data may have no correlation with the intruder's actual location. Hence the group is subject to a spatial constraint that requires it to be composed of sensors within a zone surrounding the moving intruder (e.g., a circle centered at the estimated location of the intruder). Meanwhile, the group is also subject to a timing constraint that requires it to move at the same speed as the intruder with sensors dynamically joining and leaving the group. Thus, sensors in the group must actively multicast the location of the intruder to other sensors that are likely to meet the moving zone within a certain deadline. The set of sensors to be notified depends on the moving speed of the intruder and the time it takes for a sensor to wake up and get ready to join the group. In addition, in order to conserve energy and maintain spatial locality as related to data aggregation, nodes should receive the multicast message as late as possible. We call this property "just-in-time" delivery.

We propose a novel class of multicast mechanisms that exhibit "just-in-time" temporal delivery semantics for disseminating data spatially in sensor networks. The distinctive trait of this new form of multicast, called *mobicast*, is the delivery of all nodes that happen to be in a prescribed region of space at a particular point of time. Spatial constraints are combined with temporal constraints by offering the application the ability to request the routing of a message to all points inside a delivery zone while allowing the latter to be defined as a function of time, thus having a continuously changing configuration. The first major challenge derives from the fact that early delivery may not be desirable as it leads to unnecessary energy consumption as the sensors become ready too far in advance relative to the required delivery time. It is the energy minimization constraint that rules out trivial solutions such as full network flooding. The second challenge arises from the fact that any protocol likely to succeed must factor in network topology and geometry. The strong "just-in-time" spatial delivery guarantee can be provided only if the protocol takes into account the spatial distribution of the sensors across the network. Sophisticated analysis will be required to ensure that the demanded guarantees are actually met. The use of spatiotemporal constraints in the specification, the focus on energy minimization, and the reliance on novel geometric analysis are the defining features of this research.

While *mobicast* is conceptually powerful, its implementation on sensor networks is fraught with difficulties. Key among them is the ability to ensure just-in-time delivery guarantees over a wide range of network topologies. The paper introduces two topological compactness metrics for spatially distributed networks designed to facilitate the analysis of information propagation behaviors across networks, and presents a protocol that uses these topological values for the network to meet the strong just-in-time delivery requirement of *mobicast*.

The remainder of the paper is organized as follows. We specify *mobicast* formally in Section 2. A protocol to achieve reliable *mobicast* in sensor networks is described in Section 3. An analysis of the protocol follows in Section 4. Discussion, related work and conclusions appear in sections 5, 6 and 7, respectively.

2 Problem Definition

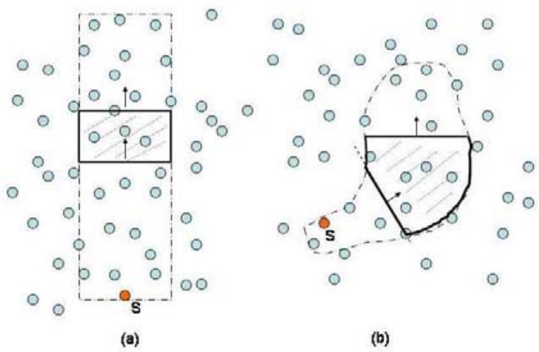

Fig. 1. Sample mobicast delivery zones

The ultimate goal of *mobicast* is to achieve just-in-time information dissemination to all nodes in some prescribed spatial area in the network. We use a "delivery zone," denoted as $Z[t]$, to represent the area where information D should be delivered at time t. As the *mobicast* delivery zone $Z[t]$ evolves over time, the set of recipients for D changes as well. Accordingly, we characterize a *mobicast* by the information D to be delivered and its associated delivery zone $Z[t]$ whose coverage changes over a period of time T:

$$\langle D, Z[t], T \rangle \tag{1}$$

Fig.1 shows two examples of *mobicast* with different kinds of delivery zones. Fig.1(a) depicts a rectangle-shaped zone (shaded) that moves from the source located at the bottom of the figure to the top. As the delivery zone moves, some nodes enter the zone and some others leave the zone. *Mobicast* may require that a node be delivered the message D at the time it gets in the zone, or before the time it moves out of the zone. Note that the shape and motion of a delivery zone are defined/specified by *mobicast* users (for their spatiotemporal delivery requirement of information D). A *mobicast* protocol then needs to achieve this spatiotemporal delivery requirements efficiently in various network topologies. Fig.1(b) shows a more general example where the delivery zone assumes an arbitrary shape, with both its shape and location evolving over time. This may be the case when the delivery requirements change in response to unexpected developments in the delivery zone.

The complexity of a *mobicast* protocol in general depends on the level of the delivery guarantee it wants to achieve. In this paper, we first consider the following strong delivery guarantee: once a node α is in a delivery zone $Z[t]$, it should receive the information D immediately. Let Ω be the set of all nodes in space, let $\boldsymbol{r}(j)$ be the location of node j, and let $D[j,t]$ denote the fact that j has been delivered the information D at time t. Let the time when the *mobicast* is initiated be zero. This *mobicast* delivery property can be formally stated as

$$\langle \forall j, t : j \in \Omega \wedge 0 \leq t \leq T :: \boldsymbol{r}(j) \in Z[t] \Longrightarrow D[j,t] \rangle^{1} \tag{2}$$

This statement can be interpreted as "During the *mobicast* session, all nodes inside zone Z at time t should have information D."

Unfortunately, delivery property (2) is practically impossible to realize in most wireless ad hoc networks. The reasons include:

- First, communication latency is often not negligible in wireless ad hoc networks. This is especially true in wireless sensor networks where sensor nodes might have a sleeping schedule in order to save energy. Note that (2) implies instantaneous delivery to all nodes at the initial delivery zone $Z[0]$. If $Z[0]$ contains a node other than the sender node, it is impossible for the node to receive information D instantly at time 0 when considering the communication latency.
- Second, a wireless ad hoc network may be partitioned. A delivery zone, specified by some geometric property alone, might cover nodes in multiple network partitions, which in turn renders the delivery impossible.
- Third, we did not put any restrictions on the evolving behavior of the delivery zone. One can imagine cases where a user-specified delivery zone evolves too fast such that its speed of change over space is faster than the maximum delivery speed a network can support.

As such, we are forced to weaken the ideal *mobicast* delivery property in the following practically-minded manner: *mobicast* satisfies property (2) only after some initialization time t_{init} on a connected network. That is

$$\langle \forall j, t : j \in \Omega \wedge t_{init} < t \leq T :: \boldsymbol{r}(j) \in Z[t] \Longrightarrow D[j,t] \rangle \tag{3}$$

[1] The three-part notation $\langle \mathbf{op}\, quantified\_variable : range :: expression \rangle$ used throughout the text is defined as follows: The variables from *quantified_variables* take on all possible values permitted by *range*. If *range* is missing, the first colon is omitted and the domain of the variables is restricted by context. Each such instantiation of the variables is substituted in *expression* producing a multiset of values to which **op** is applied, yielding the value of the three-part expression.

Thus, each *mobicast* session has two phases. The first, from time 0 to t_{init}, is an initialization phase in which no delivery guarantee is specified. The second phase, from time t_{init} to T, is a stable phase in which the strong spatiotemporal guarantee is required.

2.1 Three Optimization Concerns

Note that, because communication latency is a random variable, it is impossible for one to schedule the delivery of a message to a node at an exact time. In order to achieve the delivery property (3), one has to consider the worst case scenario and schedule the delivery of *mobicast* message ahead of time. Let $t_r(j)$ denote the time a node j first receives the *mobicast* message, $t_{in}(j)$ be the first instant of time j enters the delivery zone. We call the time difference $t_{in}(j) - t_r(j)$ the "slack time" of message delivery. Note that specification (3) implies that t_{in} is the deadline of message delivery, and the slack time measures how early the message is delivered to a node comparing to its deadline to be there.

One optimization concern for any *mobicast* protocol is to reduce the overall time interval between the reception of a message and its required delivery to the application, i.e., the slack time. Minimizing the average slack time t_{slack} for all nodes that were ever in the delivery zone leads to less energy consumption and better locality in spatial data aggregation.

Another optimization dimension for *mobicast* is to reduce the total number of retransmissions needed for each *mobicast* session while delivering the spatial and temporal guarantees. This direction is similar to that of all broadcast and multicast protocols for ad hoc networks.

The third optimization concern is to make the initialization phase as short as possible. In general, the length of the initialization time depends on the size of the delivery zone, the network connectivity pattern within the region, and the protocol execution behavior. While a *mobicast* protocol has no control over the former two factors, it can try to make t_{init} as short as possible by optimizing its execution strategy.

Next we consider the domain of sensor networks and present a *mobicast* protocol that satisfies property (3) in an efficient way.

3 Description of a *Mobicast* Protocol

As a proof of concept, we present a *mobicast* protocol for the case when the delivery zone is a convex polygon P that moves through space at constant velocity $\boldsymbol{v}$ for a duration T. For simplicity, we use an example where the convex polygon is a rectangle and whose shape does not change over time. While conceptually simple, this mobicast protocol is useful for coordination scenarios where the mobile event does not change its velocity and spatial confinement very often, and is very challenging to implement. Our effort in deriving the protocol yields a few insights and new concepts useful for the study of spatiotemporal information dissemination strategies in sensor networks. We will also discuss the potential implications of entertaining more general cases in later sections. Before presenting the protocol, we first describe its key assumptions regarding the network.

3.1 Sensor Network Model

The sensor network model for our protocol is as follows. The network does not have any partition, and all nodes are location-aware, i.e., they know their location $\boldsymbol{r}$ in space with reasonable accuracy. The maximum clock-drift among the sensors in the system is small enough to be negligible. All nodes support wireless communication and are able to act as routers for other nodes. Local wireless broadcast is reliable, i.e., once a local broadcast is executed, it will be heard by all its neighbors within latency τ_1.

3.2 A *Mobicast* Protocol

In order to describe the *mobicast* protocol more concisely, we introduce some terminology. The reader is reminded that the delivery zone is an area where the delivery of messages to the application takes place and is specified by the application itself. Our protocol also uses a "forwarding zone" $F[t]$ that is moving at some distance ahead of the delivery zone, as shown in Fig.2. We call the distance between the forwarding zone and its associated delivery zone the "headway distance" (of the forwarding zone). The shape of the forwarding zone is related to the shape of the delivery zone, and the topology of the underlying network. The choice of the headway distance and the size of the forwarding zone is such that it guarantees that all nodes entering the delivery zone will have received the *mobicast* message in advance, even if some of them are not directly connected (1-hop) to any nodes already in the delivery zone. In the meantime, the forwarding zone also serves to limit the retransmission to a bounded space while ensuring that all nodes that need to get the message will get it. We will discuss how the forwarding zone is determined in the next section. While nodes in a forwarding zone retransmit the *mobicast* message as soon as they receive it, the nodes in front of the forwarding zone enter a "hold-and-forward" state if they receive the *mobicast* message. They do not retransmit the message until becoming members of the forwarding zone. It is the action of the nodes in the hold-and-forward zone that ensures the "just-in-time" feature of the *mobicast* delivery policy while keeping the average slack time t_{slack} small. This behavior results in a virtual "hold-and-forward zone" in front of the forwarding zone, as also indicated in Fig.2.

When a request $\langle D, Z[t], T \rangle$ is presented to the *mobicast* service at time t_0, it constructs and broadcasts a *mobicast* message to all the neighbors. A *mobicast* packet $\tilde{m}$ contains the following information: a unique message identifier, a delivery zone descriptor, a forwarding zone descriptor, the session start time t_0, the session lifetime T, and the message data D. The unique message identifier is created from the combination of the location of the source and the time t_0 of the request. The delivery zone descriptor encodes the original location, the shape of the zone, and its velocity. The forwarding zone descriptor encodes the

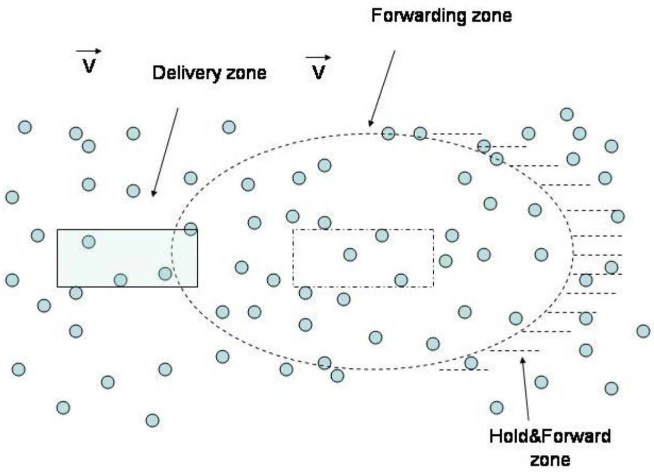

Fig. 2. Mobicast example

shape and the original location of the forwarding zone, which is computed using some knowledge about the network and the shape of the delivery zone. We will discuss in detail the computation of the forwarding zone in later sections.

The *mobicast* protocol is described in Fig.3. While not explicitly shown in the code, this *mobicast* protocol exhibits two phases in its spatial and temporal behavior. The first is an initialization phase, in which the nodes are trying to "catch-up" with the spatial and temporal demands of the *mobicast* . When a node in the path of the forwarding zone receives a message for the first time, it rebroadcasts the message as soon as possible. This phase continues until a stable forwarding zone that travels at a certain distance d_s ahead of the delivery zone is created.

The second phase is a cruising phase in which the forwarding zone moves at the same velocity as the delivery zone. The protocol enters this phase after the delivery zone and the forwarding zone reach the stable headway distance d_s. This cruising effect is achieved by having the nodes at the moving front of the forwarding zone retransmit the *mobicast* message in a controlled "hold-and-forward" fashion to make the forwarding zone move at the velocity $\boldsymbol{v}$. The initialization and the cruising phases together establish *mobicast* property (3) with t_{init} being the time required by the initialization phase.

In the next section we turn our attention to: how the forwarding zone and its stable headway distance are computed; what is the value of t_{init} given a specific *mobicast* request and the spatial properties of the underlying network; and how the protocol delivers on its guarantees.

Upon hearing a mobicast message $\tilde{m}$ at time t.

```
if (m̃ ) is new and t < t_0 + T
    if (I am in F[t]) then
        broadcast m̃ immediately ;                       // fast forward
        if (I am in Z[t]) then
            deliver the message data D to the application layer;
        else
            compute the earliest time t_in for me to enter the delivery zone;
            if t_in exists and t_in < t_0 + T
                schedule delivery of data D to the application layer at t_in;
            end if
        end if
    else
        compute the earliest time t' for me to enter the forwarding zone;
        if t' exists
            if t_0 ≤ t' ≤ t
                broadcast m̃ immediately ;               // catch-up!
            else if t < t' < t_0 + T
                schedule a broadcast of m̃ at t';        //hold and forward
            end if
        end if
    end if
end if
```

Fig. 3. A mobicast protocol

4 Analysis

A key element in the *mobicast* protocol (Fig.3) is the forwarding zone. As we mentioned earlier, the purpose of the forwarding zone is to ensure that all the nodes in a delivery zone receive the *mobicast* message, and that they receive the message before entering the delivery zone. The latter is guaranteed by sustaining a headway distance d_s between the forwarding zone and the delivery zone.

The shape of a forwarding zone depends on the following three factors: the shape of the delivery zone, the spatial distribution of the network nodes, and the topology of the network. Fig.4 shows a rectangle *mobicast* example to illustrate why this is the case. The source node S initiates a mobicast. For node A to be able to deliver the message when it becomes a member of the delivery zone, it should have received the message by that time. In scenario Fig.4(a), this means the message should have gone through G (in order for it to reach A). This implies that A and G should be in the forwarding zone together at some point in time before A can receive the message. On the other hand, if the network connectivity is "denser", as in Fig.4(b), it is obvious that the width of the forwarding zone can be relatively smaller. Furthermore, in Fig.4(a) the height of the forwarding zone has to be bigger than the height of the delivery zone to include D. Without being

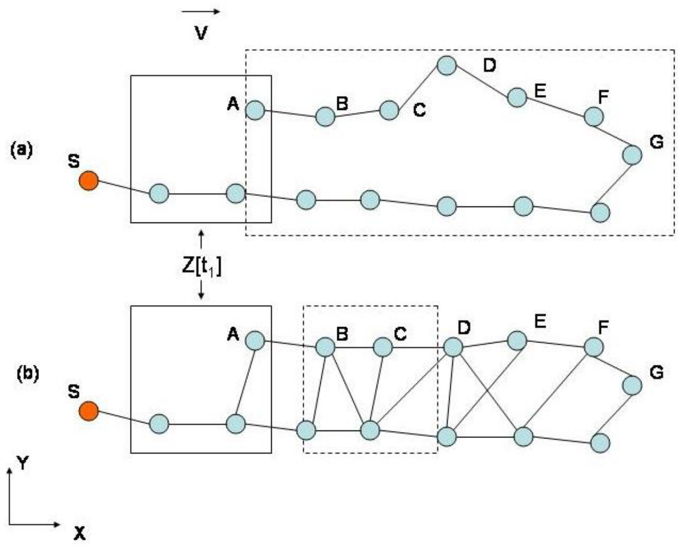

Fig. 4. Effect of network topology on the size of forwarding zone

so, nodes A, B, C would be effectively partitioned from the rest of the nodes in the network, as node D would not participate in forwarding (retransmission) as it was not in the forwarding zone. This is just one special example with an ad hoc choice of forwarding zone. The question we would like to answer is, in an arbitrary sensor network, how do we determine the forwarding zone for a specific delivery zone?

In the rest of this section we first discuss how to compute the forwarding zone, then show what headway distance is needed for ensuring the delivery guarantee. Finally, we show that our protocol provides the desired spatiotemporal guarantees given a proper choice for the forwarding zone and the headway distance.

4.1 Computing the Forwarding Zone

In order to compute the size of the forwarding zone for a specific delivery zone on an arbitrary network, we first introduce a compactness measure for the network, called "Δ-compactness."

Δ-Compactness. Given a geometric graph/network $G(V, E)$, Δ-compactness seeks to quantify the relation between the Euclidean distance and the *network spatial distance* among network nodes. The network spatial distance $\tilde{d}(i, j)$ between two nodes i and j is defined in the following manner. Let $d(e)$ denote the Euclidean distance of a network edge e. If a network path l contains an edge e, we say e is in l. We define the "edge-length" of path l to be the sum of the physical distances along its edges:

$$L(l) = \sum_{e \text{ in } l} d(e) \tag{4}$$

Let $M(i,j)$ be the set of shortest network paths between nodes i and j. The network spatial distance $\tilde{d}(i,j)$ is

$$\tilde{d}(i,j) = \min_{l \in M(i,j)} L(l) \tag{5}$$

The Δ-compactness of a geometric graph $G(V,E)$ is defined as the smallest Euclidean distance to network spatial distance ratio among the nodes:

$$\delta = \min_{i,j \in V} \frac{d(i,j)}{\tilde{d}(i,j)} \tag{6}$$

Theorem 1. *Let i,j be any two nodes in a network with Δ-compactness value δ. Let $E(i,j,\delta)$ be an ellipse using i,j as two foci and with eccentricity δ. There is at least one shortest path between i and j inside the ellipse $E(i,j,\delta)$.*

Proof: (We can prove this theorem by contradiction. Proof omitted due to page limit. Reader can find a proof of this theorem and other theorems presented in this paper in [5].)

This theorem is very useful for limiting the flooding region while guaranteeing fastest point to point message delivery in a geometric network. In our case, this metric helps us to decide the size of the forwarding zone.

δ-Cover. We also introduce a notion called the "δ-cover" of a polygon to simplify the mathematical description of the forwarding zone. The δ-cover of a convex polygon P is defined as the locus of all points p in the plane such that there exists two points q and r in the polygon P that satisfy the constraints

$$d(p,q) + d(p,r) \leq \frac{1}{\delta} d(q,r) \tag{7}$$

where $d(x,y)$ is the distance between points x and y.

Theorem 2. *Let i,j be two nodes in a network with Δ-compactness value δ. If i,j are inside a convex polygon P, then the δ-cover of P contains at least one shortest path between i and j.*

Given this theorem, we now have a way to determine the size of the forwarding zone for any convex-shaped delivery zone.

The Forwarding Zone. Given a *mobicast* delivery zone of convex shape P, if the *mobicast* is executed on a network with Δ-compactness value δ, then we choose the shape of the forwarding zone to be the δ-cover of P. We call the area of P in the forwarding zone the "core" of the forwarding zone. One may easily see the following corollary.

Corollary 1. *Let i, j be two nodes in the core of a forwarding zone in a network of Δ-compactness δ. Then the forwarding zone contains at least one shortest path between i and j.*

Note that this corollary provides the following guarantee: inside the δ-cover of a delivery zone, there is a shortest network path between any two nodes in the delivery zone. This also means that if one node in the delivery zone has a message for all other nodes in the delivery zone, it can choose only to flood the δ-cover of a delivery zone, i.e., forwarding zone, to ensure that they all get the message, with an additional property: the message is delivered through shortest paths.

Note also that, so far we are only concerned with guaranteed spatial delivery. In order for all nodes in the delivery zone to receive the multicast message on time, we need to have the forwarding zone moving ahead of the delivery zone.

4.2 Computing the Stable Headway Distance

The headway distance of the forwarding zone is a way to tell the protocol how far ahead to prepare the message delivery in order not to miss the delivery deadline due to some unexpected "twists and turns" on the related network path. One may imagine that networks with more "curved" network paths require longer headway distances than those that are more "direct." In order to capture this notion more precisely, we introduce another compactness metric for the network, called "Γ-compactness."

Γ-compactness. Γ-compactness quantifies the relation between the network distance (in terms of *hops*) and the Euclidean distance among the nodes in a geometric network. Let $h(i, j)$ be the minimum number of network hops between nodes i and j, and $d(i, j)$ be the Euclidean distance between them. We define the Γ-compactness of a geometric graph $G(V, E)$ to be the minimum ratio of the Euclidean distance to the network hop distance between any two nodes, i.e.,

$$\gamma = \min_{i,j \in V} \frac{d(i,j)}{h(i,j)} \tag{8}$$

Intuitively, if a network's Γ-compactness value is γ, then any two nodes in the network separated by a distance d must have a shortest path between them no greater than d/γ hops.

Theorem 3. *Let N be a network with a Γ-compactness value γ, and let τ_1 be its maximum 1-hop communication latency. The lower bound of the maximum message delivery speed over the space of N is $\frac{\gamma}{\tau_1}$.*

This theorem tells us that given a geometric network with Γ-compactness value γ, the delivery zone cannot move at a higher velocity than $\frac{\gamma}{\tau_1}$ if one wants delivery guarantee in all cases.

The Headway Distance. The stable headway distance d_s must be large enough to ensure that when the delivery zone reaches the current location of the core of the forwarding zone, all the nodes in the core have received the message, *i.e.*, $t_{in} > t_r$ is achieved for all nodes.

Theorem 4. *Let S_d be the maximum distance between the boundary points of the delivery zone, let v be the speed of the delivery zone, let τ_1 be the 1-hop maximum network latency of the network and let γ be its Γ-compactness. If we select $d_s = v\tau_1\lfloor\frac{S_d}{\gamma}\rfloor$, then all nodes in the core of the forwarding zone will have received the mobicast message when the delivery zone reaches them, assuming at least one node in the core has received the message.*

Given the headway distance d and the shape F of the forwarding zone, a node can easily determine the current location of the forwarding zone using velocity v, current time t, sending time t_0 and the source location r_0. Note that t_0 and r_0 can be obtained from the mobicast protocol message ID.

4.3 Length of the Initialization Phase

As we pointed out earlier, it is in the cruising phase that the *mobicast* protocol guarantees on-time delivery. In the initialization phase, the timing constraint of *mobicast* is realized in a best-effort way. It is possible that during the initialization phase, some nodes do not get the messages on-time. In general, the shorter the initialization phase, the more deliveries are on-time. The initialization phase continues until one node inside the core of the forwarding zone that is d_s ahead of the delivery zone receives the *mobicast* message. From discussions in the last section, we know that after this, the timing constraints of *mobicast* are always satisfied.

The time (t_{init}) taken by the *mobicast* protocol to enter the cruising phase is related to the stable headway distance needed, the delivery zone speed, and the maximum admissible spatial information propagation speed of the network. The upper bound of t_{init} that our mobicast protocol achieves is addressed by the following theorem.

Theorem 5. *Let d_s be the required headway stable distance between the forwarding zone and the delivery zone. Let w be the width of the delivery zone. Let v be the speed of the delivery zone and u be lower bound of the maximum message delivery speed achievable on the network. The mobicast protocol initialization time t_{init} is no greater than $\frac{(d_s+w)}{u-v}$*

The Spatiotemporal Guarantees of the Protocol. The spatiotemporal guarantee of the presented *mobicast* protocol is addressed by the following theorem:

Theorem 6. *If at any instant of time in a mobicast session, its (user-defined) delivery zone covers at least one node in the network, our mobicast protocol delivers property (3).*

We provide only a sketch of the proof of this theorem here.

Proof: If a delivery zone covers at least one node in the network at any instant of time, then whenever the last node in a delivery zone is leaving a delivery zone, there must be another node entering it. The same is true for the core of the forwarding zone, because it is of the same shape as the delivery zone and moves on the same path. So that if at one point in time, a node in the core of the forwarding zone has received the mobicast message, it will always be able to pass on to all others nodes on its path, because of theorem (4). From theorem (5), it is easy to see property (3) is satisfied. □

Note that if the network has a big "hole" such that the delivery zone may fall into it at some instant of time, i.e., the delivery zone covers no network node, then, our protocol does not provide the guaranteed spatiotemporal delivery. That is why theorem (6) requires the condition "at any instant of time in a *mobicast* session, its user defined delivery zone covers at least one node in the network".

5 Discussion

In the last section we introduced two network compactness metrics to help us choose the right forwarding zone and its headway distance from the delivery zone to achieve the *mobicast* delivery guarantee without unnecessary flooding. The higher the compactness, the smaller the forwarding zone and its headway distance. These compactness values must to be computed for supporting mobicast. Calculating them involves computing the shortest path and Euclidean distances of each pair of nodes in a given network. The all-pair shortest path of a graph $G(V, E)$ can be computed in $O(VE \log V)$ time by using Johnson's algorithm [6]. All-pair distance can be computed in $O(V^2)$ time. So we can compute the the $\mathit{\Gamma}$-compactness of the graph in $O(VE \log V)$ time. $\mathit{\Delta}$-compactness can also be computed in $O(VE \log V)$ time. Thus it is not feasible for individual sensor nodes to compute these values in a large network. In practice, one may have a central server collect all the location and connectivity information, do the computation and use one broadcast to inform all the nodes this value.

Note that the compactness metrics are defined for the whole network. Different areas in the network could have their regional compactness values. When those values are available to the corresponding nodes, the size of the forwarding zone can change from one area to another in the network. We expect that this adaptive behavior will reduce the overall retransmission overhead. Computing only regional compactness also is computationally less intensive. The tradeoff for doing this is one may not be able to support reliable *mobicast* with delivery zones larger than the size of the region used for the compactness computation. Note also that these compactness metrics are geared for worst-case analysis of a "communication unfriendly" network topology in any area of the network. They are chosen in this manner because the *mobicast* property as specified by (3) is an absolute guarantee. If one prefers a weaker, probabilistic delivery guarantee, weaker (e.g., average) compactness measures would be more appropriate.

For simplicity, our protocol carries out an "as soon as possible" flooding inside the forwarding zone. If nodes have accurate pictures regarding the locations of their one hop or two hop neighbors, one can reduce the number of necessary re-transmissions in a manner similar to techniques proposed for improving broadcast efficiency [7, 8]. In a probabilistic guarantee scenario, one may also use probabilistic retransmission-reduction techniques such as the one described in [9]. A review of these and other related methods can be found in [10].

Furthermore, in order to focus on the essential characteristics of *mobicast*, we assume that the local broadcast is reliable, i.e., any message broadcast by a node is to be heard by its neighbors in τ_1 time. Because of the possibility of "hidden nodes" and the high cost of coordination mechanisms to solve the hidden nodes problem, a more realistic choice would be to relax the reliability assumption about local broadcast and, in turn, weaken the delivery guarantee to a probabilistic one.

Finally, while the *mobicast* protocol we presented applies to cases where the delivery zone is a convex polygon P that moves through the space at constant velocity $\boldsymbol{v}$ for a duration T, *mobicast* in general applies to a much wider set of spatiotemporal constraints. The delivery zone can exhibit any evolving characteristics as long as it is sustainable by the underlying system. While they may all require similar ideas of forwarding zone and headway distance to maintain the spatiotemporal properties inherent in *mobicast*, a different type of delivery zone may require different protocol handling details. Classification of a useful set of mobicast delivery zone scenarios and design of the corresponding *mobicast* protocols are part of our plans for future work.

6 Related Work

Mobicast is a multicast mechanism that involves both the spatial and the temporal domains. The idea of disseminating information to nodes in a geographic area is not new. Navas and Imielinski proposed geographic multicast addressing and routing [11, 12], dubbed as "geocast," for the Internet. They argued that geocast was a more natural and economic alternative for building geographic service applications than the conventional IP address-based multicast addressing and routing. In a geocast protocol, the multicast group members are determined by their locations. The initiator of a geocast specifies an area for a message to be delivered, and the geocast protocol tries to deliver the message only to the nodes in that area. Ko and Vaidya investigated the problem of geocast in mobile ad hoc networks [13] and proposed to use a "forwarding zone" to decrease delivery overhead of geocast packets. Various other mechanisms [14–16] have been proposed to improve geocast efficiency and delivery accuracy in mobile ad hoc networks. Zhou and Singh proposed a content-based multicast [17] in which sensor event information is delivered to nodes in some geographic area that is determined by the velocity and type of the detected events. While different in style and approach, all these techniques assume the delivery zone to be fixed. They also assume the same information delivery semantics along the temporal domain, i.e., information is to be delivered “as soon as possible”. However, local coordination often requires just-in-time delivery in sensor networks.

Data aggregation is an important information processing step in sensor networks. Several techniques have been proposed to support data aggregation in sensor networks. For example, both directed diffusion [2, 18] and TAG [4] allow data to be aggregated on their route from the sources to a base station. No explicit local coordination is supported by these techniques. LEACH [3] organizes sensors into local clusters where each cluster head is responsible for aggregating the data from the whole cluster. However, there is no notion of mobility and the clusters do not move in space following a physical entity. In contrast, supporting local coordination for mobile physical entities is a primary goal of mobicast. Perhaps the EnviroTrack project [19] is closest in spirit to our work. EnviroTrack can dynamically create and maintain a group that tracks mobile entities in the environment. A transport layer protocol maintains connections between mobile groups. However, both Envirotrack and the other aforementioned projects do not provide any guarantees regarding spatiotemporal constraints.

7 Conclusion

In this paper we have presented the basic idea of *mobicast*, a new multicast paradigm for disseminating information to a set of nodes in a sensor network under spatiotemporal constraints. To demonstrate the feasibility of *mobicast*, we designed a protocol and explored its ability to deliver strong spatiotemporal guarantees. The key element in the protocol is a dynamic forwarding zone moving ahead of the delivery zone. Furthermore, we introduced two new notions of network compactness and proved several related theorems useful in the analysis of the information propagation over sensor networks. Using these results we were able to determine the shape of the forwarding zone and the headway distance needed for our protocol to ensure strong multicast delivery guarantees in space and time while keeping retransmission overhead and average slack time small. The powerful just-in-time spatial delivery semantics of *mobicast* serves to optimize resource utilization for multicast tasks in sensor networks and enables application programmers to address both spatial and temporal perspectives of communication and coordination explicitly, in a manner atypical of current multicast models.

References

1. D. Estrin, e.a.: Embedded everywhere: A research agenda for networked systems of embedded computers. National Academy Press (2001) Computer Science and Telecommunications Board (CSTB) Report.
2. Intanagonwiwat, C., Govindan, R., Estrin, D.: Directed diffusion: a scalable and robust communication paradigm for sensor networks. In: Mobile Computing and Networking. (2000) 56–67
3. Heinzelman, W.R., Chandrakasan, A., Balakrishnan, H.: Energy-efficient communication protocol for wireless microsensor networks. In: HICSS. (2000)

4. Madden, S., Franklin, M., Hellerstein, J., Hong, W.: Tag: a tiny aggregation service for ad-hoc sensor networks. (OSDI 2002, Boston MA)
5. Huang, Q., Lu, C., Roman, G.C.: Mobicast: Just-in-time multicast for sensor networks under spatiotemporal constraints. WUCS 42, Washington Unievrsity in Saint Louis (2002)
6. Cormen, T.H., Leiserson, C.E., Rivest, R.L.: Introduction to Algorithms. The MIT Press (1999)
7. Peng, W., Lu, X.: On the reduction of broadcast redundancy in mobile ad hoc networks. In: Proceedings of the ACM Symposium on Mobile Ad Hoc Networking and Computing (MOBIHOC). (2000)
8. Qayyum, A., Viennot, L., Laouiti, A.: Multipoint relaying: An efficient technique for flooding in mobile wireless networks. Technical Report Research Report RR-3898, INRIA (2000)
9. Ni, S.Y., Tseng, Y.C., Chen, Y.S., Sheu, J.P.: The broadcast storm problem in a mobile ad hoc network. In: Proceedings of the Fifth Annual ACM/IEEE International Conference on Mobile Computing and Networking. (1999) 152–162
10. Williams, B., Camp, T.: Comparison of broadcasting techniques for mobile ad hoc networks. In: Proceedings of the ACM International Symposium on Mobile Ad Hoc Networking and Computing (MOBIHOC). (2002) 194–205
11. Imielinski, T., Navas, J.C.: Gps-based addressing and routing. RFC2009, Computer Sciece, Rutgers University (1996)
12. Navas, J.C., Imielinski, T.: Geocast - geographic addressing and routing. In: Proceedings of the Third Annual International Conference on Mobile Computing and Networking (MobiCom '97. (1997) 66–76
13. Ko, Y., Vaidya, N.: Geocasting in mobile ad hoc networks: Location-based multicast algorithms (1998)
14. Stojmenovic, I.: Voronoi diagram and convex hull based geocasting and routing in wireless networks. TR TR-99-11, University of Ottawa (1999)
15. 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
16. Boleng, J., Camp, T., Tolety, V.: Mesh-based geocast routing protocols in an ad hoc network. In: Proceedings of the IEEE International Workshop on Parallel and Distributed Computing Issues in Wireless Networks and Mobile Computing (IPDPS). (2001) 184–193
17. Zhou, H., Singh, S.: Content based multicast (cbm) for ad hoc networks. Mobihoc 2000, Boston, MA (2000)
18. Intanagonwiwat, C., Estrin, D., Govindan, R., Heidemann, J.: Impact of network density on data aggregation in wireless sensor networks. International Conference on Distributed Computing Systems (ICDCS-22) (2001)
19. Blum, B., Nagaraddi, P., Wood, A., Abdelzaher, T., Son, S., Stankovic, J.: An entity maintenance and connection service for sensor networks. (The First International Conference on Mobile Systems, Applications, and Services (MobiSys), San Francisco, CA, May 2003)