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

Qingfeng Huang, Chenyang Lu, and Gruia-Catalin Roman

WUCS-02-42

December 18, 2002

Department of Computer Science and Engineering
Campus Box 1045
Washington University
One Brookings Drive
St. Louis, MO 63130-4899

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*. *Mobicast* 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.

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*. *Mobicast* 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 sensors, 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, 3, 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 the 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). We now give an example to illustrate the kind of spatiotemporal constraints in such applications. Assume we deploy acoustic sensors in a security area 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 activate other sensors when necessary. To track an intruder, sensors in the 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. Hence the group is subject to a spatial constraint that requires the group be always 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 the group to move at the same speed as the intruder while sensors dynamically join and leave 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 deadline. The set of sensors to be notified depends on the moving speed of the intruder and the time it takes a sensor to wake up and get ready to join the group. In addition, in order to conserve energy and maintain spatial locality of data aggregation, nodes should receive the multicast message as late as possible. We call this property the "just-in-time" delivery requirement.

We propose a novel class of multicast mechanisms that exhibits "just-in-time" temporal delivery semantics for disseminating data spatially in sensor networks. This multicast mechanism delivers information to all nodes in an application-specified "delivery zone". We call this multicast mechanism "*mobicast* " because its delivery zone may move, morph, and in general, evolve over time. For instance, the delivery zone can be an ellipse moving through space over time under certain velocity, or can be a pie-shaped area expanding over time. Application developers can easily encode their desired spatiotemporal constraints for information dissemination into the evolving behavior of the delivery zone. This unique characteristic of *mobicast* makes it a powerful mechanism for dynamic data aggregation in sensor networks. The dynamic group management protocol described above can be easily implemented by mobicasting a "join-group" message to a delivery zone that moves in front of an intruder at the same velocity. The wake-up just-in-time requirement for the sensors can be satisfied by using *mobicast* in this manner.

While *mobicast* is conceptually powerful, its implementation on sensor networks has many challenges. One of the key challenge we present and tackle in this paper is to ensure a strong just-in-time delivery guarantee on various sensor network topologies. We introduce two topological compactness metrics for spatially distributed networks in order to facilitate the analysis of information propagation behavior on these networks, and present a protocol that makes use of the knowledge of these topological values of the network and to achieve the strong just-in-time delivery requirement of *mobicast*.

The paper is organized as follows. We first present related work in section 2, then formally specify *mobicast* in section 3. A protocol to achieve reliable *mobicast* in sensor networks is introduced in section 4. An analysis of the protocol follows in section 5. Discussions and conclusions appear in sections 6 and 7, respectively.

2. Related Work

Mobicast is a multicast mechanism that involves both spatial and temporal domain. The idea of disseminating information to nodes in a geographic area is not new. Navas and Imielinski proposed geographic multicast addressing and routing [5, 6], dubbed as "geocast," for the Internet. They elegantly argued that geocast is 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 geocasting in mobile ad hoc networks [7] and proposed to use a "forwarding zone" to decrease delivery overhead of geocasting packets. Various other mechanisms [8, 9, 10] have been proposed to improve geocast efficiency and delivery accuracy in mobile ad hoc networks. Zhou and Singh proposed a content-based multicast [11] 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, 12] 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 and 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 [13] 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 did not provide any guarantees on the spatiotemporal constraints.

3. Problem Definition

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. This kind of *mobicast* delivery zone appears to be applicable to the object tracking example we discussed in the previous section. 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.

Ideally, one may expect that once a node α is in a delivery zone $Z[t]$, it should receive the information D immediately. To capture formally this notion some notation needs to be introduced first. Let Ω be the set of all nodes in space, let $\vec{r}(j)$ 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. This idealized *mobicast* delivery property can be formally stated as

$$\langle \forall j, t : j \in \Omega \wedge t \le t_0 + T :: \vec{r}(j) \in Z[t] \Longrightarrow D[j,t]\rangle^1 \tag{2}$$

[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.

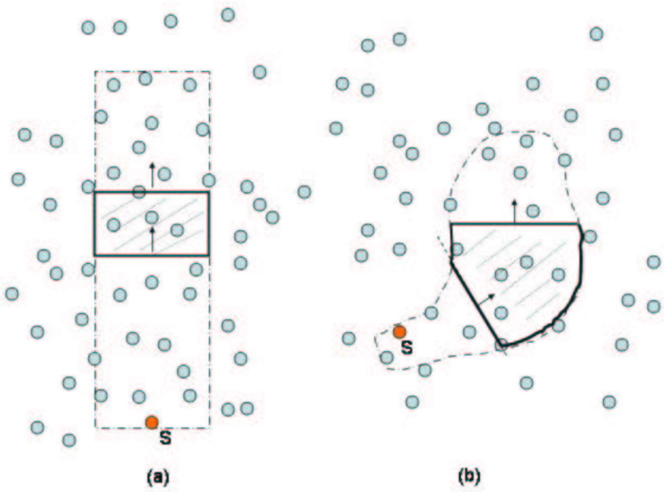

Figure 1: Sample mobicast delivery zones

where the session constant t_0 is the time when the *mobicast* was initiated, t_0+T is when the *mobicast* session expires.

Unfortunately, this delivery property (2) is practically impossible to realize in most wireless ad hoc networks. 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 have partitions. 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. This is yet another case one cannot achieve a successful *mobicast*.

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

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

The delivery property (3) requires a node to be delivered the *mobicast* message at or before the time it enters the delivery zone. While this property might appear to be not implementable, we will present a protocol that can actually meets property (3) by delivering messages ahead of schedule.

Note that we differentiate the message delivery time (to the application) from the message "reception time" at a node. A node's *mobicast* message reception time refers to the time when the *mobicast* message is received by the node. Let $t_r(j)$ denote the time a node j first receives the *mobicast* message, $t_d(j)$ be the time j delivers the message to the applications listening for the *mobicast* , and $t_{in}(j)$ be the first instant of time j enters the delivery zone. *Mobicast* requires one to guarantee $t_d(j) = t_{in}(j)$ assuming negligible latency in handing the message from the *mobicast* protocol layer to the application layer.

One optimization concern for any *mobicast* protocol is to reduce the overall time interval between the reception of a message and its delivery to the application. We call this time interval for a node "waiting time". Minimizing the average waiting time $t_{average-waiting}$ for all nodes that were ever in the delivery zone means less energy consumption and better locality in spatial data aggregation. The ideal case involves reducing $t_{average-waiting}$ to zero, i.e., a node only receives the *mobicast* message (from its neighbors) precisely at the time it enters the delivery zone. Yet, this may not always be possible for all spatial and connectivity configurations of a given network.

Another optimization dimension for *mobicast* is to reduce the total number of retransmissions needed for the 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.

Note that there are two reasons for t_{init} to exist. First, we allow users to specify a delivery zone whose initial size is greater than zero and may contain other nodes. Second, one cannot look ahead in a protocol before a *mobicast* request is presented. (Note that if we require an admissible delivery zone to be one that expands from size zero with some bounded changing rate, then we do not need t_{init}.) 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 given the other two factors.

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

4. 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 the space at constant velocity $\vec{v}$ 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 coordination 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 that 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 the key assumptions it makes regarding the network.

4.1. Sensor Network System Model

The sensor system model for our protocol is as follows. All nodes are location-aware, i.e., they know their location $\vec{r}$ 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.

4.2. 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 shapes of the forwarding zone is related to the shape of the delivery zone, and the topological patterns 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 in the past delivery zone. In the meantime, the forwarding zone also serves to limit the retransmission to a bounded space while ensuring all the nodes that need to get the message will get the message. We will discuss how it is determined in the next section. While nodes in a

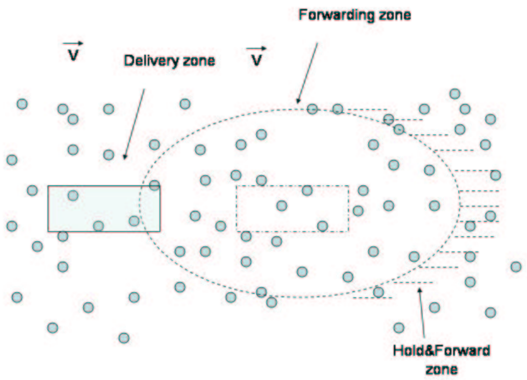

Figure 2: Mobicast example

forwarding zone retransmit the *mobicast* message as soon as they receive it, the nodes in front of the forwarding zone enter "hold-and-forward" state if they receive the *mobicast* message. They do not retransmit the message until becoming a member 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 $t_{average-waiting}$ 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 moving velocity. The forwarding zone descriptor encodes the 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.

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

```
if (m~ ) is new and t < t_0 + T
    if (I am in F[t]) then
        broadcast m~ 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 routing zone;
        if t' exists
            if t_0 <= t' <= t
                broadcast m~ immediately ;          // catch-up!
            else if t < t' < t_0 + T
                schedule a broadcast of m~ at t';          //hold and forward
            end if
        end if
    end if
end if
```

Figure 3: A mobicast protocol

In this phase 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 reaches 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 $\vec{v}$. 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 guarantee.

5. Analysis

A key element in the *mobicast* protocol (Fig.3) is the forwarding zone and its headway distance d_s from the delivery zone. As we mentioned earlier, the purpose of the forwarding zone and its headway distance is to ensure all the nodes that will be in a delivery zone to receive the *mobicast* message before entering the delivery zone, while keeping the total number of nodes participating in each mobicast session at a minimum.

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 the 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 members of the delivery zone, it should have received the message by that time. In scenario Fig.4(a), that means the message should have gone through G (in order for it to reach A). This implies 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 the width of forwarding zone can be relatively smaller. Furthermore, in Fig.4(a) the height of the forwarding zone has to be

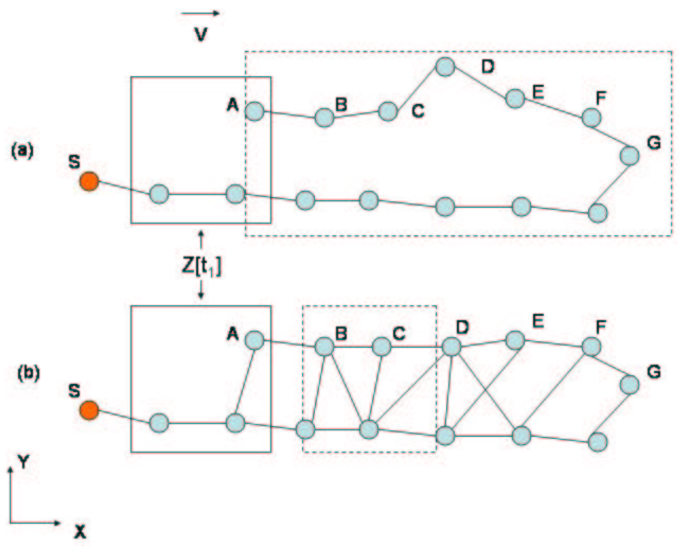

Figure 4: Spatial and connectivity configuration of the network influence the size of forwarding zone

bigger than the height of the delivery zone to include D. Without being so, nodes A, B, C will be effectively partitioned from the rest of the nodes in the network, as node D will not participate in the routing (retransmission) in the scheme. 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 how much headway distance is needed to ensure the delivery guarantee. Finally, we show that our protocol provides the desired spatiotemporal guarantees given the proper forwarding zone size and headway distance.

5.1. Computing the forwarding zone

In order to describe how the forwarding zone can be determined for a specific delivery zone on an arbitrary network, we first introduce a compactness measure for the network, called "Δ-compactness".

5.1.1. Δ-compactness . Δ-compactness aims to quantify the relation between the direct 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 "E-length" of path l to be the sum of the physical distances of all its edges:

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

Let $M(i,j)$ be the set of shortest network path between nodes i and j. The spatial network 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 spatial network distance ratio among the nodes:

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

Theorem 5.1. *Let i,j be any two nodes in a Δ-compact network. 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 prove this theorem by contradiction.

Assume the theorem is not true. That means there is at least one pair of nodes i and j, whose shortest paths are all have at least one vertex outside the ellipse $E(i,j,\Delta)$. Using the fact that for all points k on the ellipse, $d(i,k) + d(j,k) = d(i,j)/\Delta$, it is easy to prove in this case

$$\tilde{d}(i,j) > \frac{d(i,j)}{\Delta}$$

that is

$$\Delta > \frac{d(i,j)}{\tilde{d}(i,j)}$$

this directly contradicts the definition of Δ-compactness (6). □

This theorem is very useful for limiting the flooding region while guaranteeing point to point message delivery in a geometric network. In our case, this metric help us to decide the shape and size of the forwarding zone, which is related to a notion called "Δ-cover".

5.1.2. Δ-cover. We introduce the notion "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 that satisfy the following constraint: there exists two points q and r in the polygon P, such that,

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

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

Theorem 5.2. *Let i,j be two nodes in a Δ-compact network, and in a convex polygon P. Then the Δ-cover of P contains at least one shortest path between i and j.*

Proof (The proof is similar to that of theorem (5.1). Omitted.)

5.1.3. 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.

COROLLARY 5.1. *Let i, j be two nodes in the core of a forwarding zone on a Δ-compact network. Then the forwarding zone contains at least one shortest path between i and j.*

Proof This is a direct result from theorem (5.2) and the construction of the forwarding zone. □

5.2. 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 to not miss the delivery deadline due to some unexpected "twists and turns" on the related network path. One may imagine a network with more "indirect" network paths requires a longer headway distance than one that is more "direct". In order to capture this notion more precisely, we introduced another compactness metric for the network, called "Γ-compactness".

5.2.1. Γ-compactness. Γ-compactness quantifies the relation between the logical network distance and the Euclidean distance among the nodes in a geometric network. The logical network distance between two nodes is measured in terms of the minimum number of hops between them. 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 logical network distance between any two nodes in the network, 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 if any two nodes in the network are at a distance d has a shortest path of no greater than d/γ hops.

THEOREM 5.3. *Let N be a network that Γ-compactness value γ, and let τ_1 be its maximum 1-hop communication latency. Then the lower bound of the maximum message delivery speed over the space on N is $\frac{\gamma}{\tau_1}$.*

Proof Let $d(i, j)$ be the distance between two arbitrary nodes i and j in the network. We know that the shortest network path h between the two nodes is bounded by

$$h(i, j) \leq \frac{d(i, j)}{\gamma} \tag{9}$$

We also know that a message sent from one node to another node h-hops away takes no longer than $h\tau_1$ if each intermediate node forwards the message immediately after receiving it. That is, let t be the time it actually takes for the message to go from i to j, in this case we have

$$t \leq h(i, j)\tau_1$$

From this we know that the average speed v of this information propagation over distance $d(i, j)$ is

$$v = \frac{d(i, j)}{t} \geq \frac{d(i, j)}{h\tau_1} \geq \frac{\gamma}{\tau_1} \tag{10}$$

Note that the bound $\frac{\gamma}{\tau_1}$ is not dependent on $d(i,j)$. This inequality (10) is true for any two nodes in any network with Γ-compactness value γ, when all nodes in the network relay the message as soon as possible. That means $\frac{\gamma}{\tau_1}$ a lower bound of the maximum spatial message delivery speed on networks with Γ-compactness value γ. □

Theorem (5.3) that given a geometric network, there is a clear limit to what spatiotemporal information dissemination can achieve. For instance, given a geometric network with Γ-compactness value γ, the delivery zone cannot move at a higher velocity than $\frac{\gamma}{\tau_1}$ in all areas.

5.2.2. The Headway Distance. The stable headway distance d_s shall be large enough to ensure that when the delivery zone reaches where the core of the forwarding zone is, all the nodes in the core have received the message, so $t_{in} > t_r$ is achieved for all nodes.

THEOREM 5.4. *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 γ be its Γ-compactness . If we let $d_s = v\tau_1\lfloor\frac{S_d}{\gamma}\rfloor$, then all the nodes in the core of the forwarding zone will have received the the mobicast message when the delivery zone reaches where it is, assuming there is at least one node in the core that has received the message.*

Proof Let i denote the node in the core that already has the message. Then its distance to all other nodes in the core is less than S_d, because S_d is the maximum size of the delivery zone, as well as that of the core. That means the longest of the shortest network paths from i to all other nodes in the core of is less than $\lfloor\frac{S_d}{\gamma}\rfloor$ hops, which needs no greater than $\tau = \lfloor\frac{S_d}{\gamma}\rfloor\tau_1$ time for a message to traverse if all nodes help to forward the message as soon as possible. So we can conclude after τ, all nodes in the core of the forwarding zone will get the message, because all nodes forward *mobicast* messages in an ASAP fashion after entering the forwarding zone, and there is always a shortest path exists inside the forwarding zone for any two nodes inside its core.

Because the speed of the delivery zone is v, a distance $d_s = v\tau_1\lfloor\frac{S_d}{\gamma}\rfloor$ takes exactly τ time to traverse.

Hence, it is true that all the nodes in the core of the forwarding zone will have received the the mobicast message when the delivery zone reaches where it is, assuming there is 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 find the current forward 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 mobicast protocol message ID.

5.3. The Length of Initialization Phase

As we pointed out earlier, it is in the cruising phase that the *mobicast* protocol guarantees in-time delivery. In the initialization phase, the timing constraint of *mobicast* is realized in a best-effort way. It is possible that in the initialization phase, some nodes do not get the messages in-time.In general, the shorter the initialization phase, the more guaranteed deliveries. 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 theorem (5.4), we know that after this, the timing constraint of *mobicast* is always satisfied.

The time (t_{init}) it takes for the *mobicast* protocol to enter the cruising phase is related to the stable distance needed, the delivery zone speed, and the maximum admissible spatial information propagation speed of the network.

THEOREM 5.5. *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 $\frac{(d_s+w)}{u-v}$*

Proof In the protocol, the nodes in the forwarding zone and between the forwarding zone and the delivery zone all retransmit the message asap the first time receiving it. So the protocol achieves a maximum message propagation speed v_{max} in this phase. This message propagation speed relative to the delivery zone is $v_{max} - v$. Meanwhile, the end-to-end distance between the delivery zone and the core of the forwarding zone is $d_s + w$, which can be covered by a message running at $v_{max} - v$ in $t = \frac{d_s+w}{v_{max}-v}$ time. When a message from the delivery zone reaches the core of the forwarding zone d_s ahead, by definition the initialization phase is over. Hence we have

$$t_{init} \leq t = \frac{d_s + w}{v_{max} - v} \leq \frac{(d_s + w)}{u - v} \tag{11}$$

in the above we also used $u < v_{max}$, which by definition is true. □

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

THEOREM 5.6. *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) with*

$$t_{init} \leq \frac{v\tau_1 \lfloor \frac{S_d}{\gamma} + w}{\frac{\gamma}{\tau_1} - v}$$

.

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 delivery zone has received the mobicast message, it will always be able to pass on to all others nodes on the path, because our protocol and the way we choose the forwarding zone guarantees if two nodes ever appear together in the same core of the forwarding zone, one having the message means another will get the message.

Then using theorems (5.4) and (5.5), it is easy to see property (3) is satisfied. □

6. Discussion

In the last section we introduced two network compactness metrics to help us to choose the right forwarding zone and its headway distance from the delivery zone to achieve the *mobicast* delivery guarantee without unnecessary flooding. These metrics should be computed when the network is initially deployed. Computing compactness 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 [14]. All-pair distance can be computed in $O(V^2)$ time. So we can compute the the Γ-compactness of the graph in $O(VE \log V)$ time.

Δ-compactness can also be computed in $O(VE \log V)$ time. Thus it is not feasible for individual sensor nodes to compute this value of the 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 when at different areas of the network. We expect this adaptive behavior will reduce the overall retransmission overhead. Computing only regional compactness also is less computationally intensive. The tradeoff for doing so is one may not be able to support reliable *mobicast* with relatively large delivery zones. Note also that these compactness metrics are geared for worst-case analysis of "communication unfriendly" network topology in any area of the network. They are chosen in this manner because *mobicast* property as specified by (3) is an absolute guarantee. If one prefers a weaker, probabilistic delivery guarantee, then corresponding weaker (e.g., average) compactness measures would be a better choice.

While we chose the shape of the forwarding zone to be a Δ-cover of the shape of the delivery zone in previous section, it was for pure analysis purpose. Computing an exact Δ-cover for an arbitrary polygon P can be difficult. Yet one can always choose some approximation techniques such as using the Δ-cover of P's bounding box, which is computationally much simpler, yet still has the required property. The tradeoff for doing so is that the resulting forwarding zone is bigger than necessary, thus may entail more re-transmissions for the same delivery goal. We should note that in our protocol the forwarding zone only needs to be computed once by the sender. The nodes that receive the mobicast message only need to translate the forwarding zone by their distances from the sender. This choice may not be the best as in some cases regional compactness can be much higher than the global compactness, and choosing the size of forwarding zone according to regional compactness appear to reduce the number of necessary routing nodes. Yet, in this approach has to ask to routing nodes to compute the forwarding zone themselves, which means relatively expensive in terms of computation cost while cheaper in communication cost.

For simplicity in presentation, our protocol essentially carries out an "as soon as possible" flooding inside the forwarding zone. If the nodes have accurate pictures about the locations of their one hop or two hop neighbors, then one can reduce the number of necessary re-transmissions by using this knowledge, in a manner similar to techniques proposed for improving broadcast efficiency [15, 16]. In a probabilistic guarantee scenario, one may also use probabilistic retransmission-reduction techniques such as the one proposed in [17]. A review of these and other related methods can be found in [18].

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

Finally, while the *mobicast* protocol we presented applies to the cases where the delivery zone is a convex polygon P that moves through the space at constant velocity $\vec{v}$ 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*, different type of delivery zone may require different protocol handling detail. Classification of a useful set of mobicast delivery zone scenarios and design the corresponding *mobicast* protocols are also important aspects of future work.

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 proved its ability to deliver a 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 a strong multicast delivery guarantee in space and time. The powerful just-in-time spatial delivery semantics of *mobicast* serves to optimize resource utilization for many 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] Imielinski, T., Navas, J.C.: Gps-based addressing and routing. RFC2009, Computer Sciece, Rutgers University (1996)

[6] 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

[7] Ko, Y., Vaidya, N.: Geocasting in mobile ad hoc networks: Location-based multicast algorithms (1998)

[8] Stojmenovic, I.: Voronoi diagram and convex hull based geocasting and routing in wireless networks. TR TR-99-11, University of Ottawa (1999)

[9] 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

[10] 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

[11] Zhou, H., Singh, S.: Content based multicast (cbm) for ad hoc networks. Mobihoc 2000, Boston, MA (2000)

[12] 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)

[13] Abdelazaher, T., Blum, B., Evans, D., George, J., George, S., Gu, L., He, T., Huang, C., Nagaraddi, P., Son, S., Sorokin, P., Stankovic, J., Wood, A.: Envirotrack: An environmental programming model for tracking applications in distributed sen-sor networks. (Submitted to ICDCS 2003)

[14] Cormen, T.H., Leiserson, C.E., Rivest, R.L.: Introduction to Algorithms. The MIT Press (1999)

[15] 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)

[16] 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)

[17] 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 (MobiCom'99). (1999) 152–162

[18] 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