TRADEOFF BETWEEN ENERGY AND LATENCY FOR CONVERGECAST

Mahesh Arumugam Sandeep S. Kulkarni

Department of Computer Science and Engineering
Michigan State University, East Lansing, MI 48824
Email: `{arumugam, sandeep}@cse.msu.edu`

ABSTRACT

We focus on the problem of energy-efficient convergecast in sensor networks. This problem identifies the energy-latency tradeoff during convergecast. Whenever a group of sensors communicate an event of interest, the latency involved in delivering such messages to the base station should be minimized. Since the sensors are constrained by limited power and are mostly idle, it is important that the sensors conserve energy. We show how time division multiple access (TDMA) can be effectively used to provide energy-efficient convergecast. This solution allows the sensors to save energy when the network is idle and to switch to active mode when the network observes an event. Furthermore, for a typical application where the event probability is less than $10 - 15\%$, our solution improves the network lifetime by approximately 3 fold.

1 INTRODUCTION

The ability to reliably communicate an event of interest to the base station or the outside-world is an essential function in sensor networks. For example, in applications such as A Line in the Sand (LITeS) [1, 2], where sensors detect, classify, track, and visualize intruders along an area, and habitat monitoring [3], where sensors monitor the activities of a habitat, the sensors are required to communicate their observed values to the base station. Such many-to-one (or source-to-sink) communication is often referred as *convergecast*. One of the important requirements in [1–3] is that the latency involved during such communication should be minimized.

One of the important constraints in sensor networks is limited power. To deal with this constraint, as part of an energy management scheme, sensors are allowed to turn their radio off or switch to low-power mode, where the amount of idle listening and overhearing are reduced. For example, time division multiple access (TDMA) algorithms proposed in [4–6] can be effectively used to reduce the amount of idle listening by allowing a sensor to turn its radio off in the time slots not assigned to itself and its neighbors.

It is easy to observe that there is a potential conflict between energy-management and low-latency convergecast. Specifically, if the radio is always on then the latency may be reduced. Similarly, if the sensors execute as part of an energy-management scheme then the latency may increase.

With this motivation, in this paper, we focus on the problem of energy-efficient convergecast while ensuring that the latency is within the application requirements. Specifically, we propose a convergecast algorithm using a TDMA based MAC protocol (e.g., [4–6]). With the help of simulations, we study the performance of the proposed convergecast algorithm and show that it provides a better energy-latency tradeoff.

Organization of the paper. In Section 2, we state the assumptions made in this paper and discuss the network model. In Section 3, we present our convergecast algorithm. Subsequently, in Section 4, we discuss the simulation results of the proposed algorithm. Finally, in Section 5, we state the related work, and in Section 6, we make concluding remarks.

2 MODEL AND ASSUMPTIONS

The assumptions made in this paper are in two categories; existence of base station and TDMA slot assignment.

Base station or sink. We assume that there exists a base station that is responsible for data gathering. Also, it is responsible for setting up the TDMA schedule, initiating tree construction process for routing, and sending commands/tasks to the network. In case of large scale networks (e.g., extreme scaling experiment [2]), the network is partitioned into multiple sections with one or more higher-tier nodes. These higher-tier nodes act as the base stations for the corresponding sections.

TDMA slot assignment. Once the network is deployed, the base station assigns TDMA slots (pre-computed or determined dynamically using a slot assignment algorithm [4–6]) to each sensor. Each sensor turns its radio off in the slots not assigned to itself and its neighbors.

3 CONVERGECAST ALGORITHM

Suppose each sensor listens to the medium always. Whenever a sensor (say, j) observes an event, it sends a convergecast message (say, m_c) to the base station. Now, consider a sensor (say, k) that is on the path between j and the base station. We observe that until k receives m_c from j, k spends time on *idle listening* and, hence, most of the energy is wasted. To provide a better energy-latency tradeoff, we propose a TDMA based algorithm for convergecast in sensor networks.

In this algorithm, each sensor listens to the medium in the slots assigned to its neighbors at certain distance. Whenever a sensor observes an event, it sends a convergecast message in its TDMA slot. The sensors may use a routing algorithm (e.g., logical grid routing protocol or LGRP [7]) to forward the message to the base station for visualization and monitoring. Since multiple sensors observe an event, a TDMA based algorithm for reliably communicating all the messages to the base station increases the latency considerably.

To improve the latency in convergecast, in this algorithm, each sensor operates in one of the following two modes; *TDMA mode* or *active mode*. Initially, sensors execute in the TDMA mode (to conserve energy and reduce the amount of idle listening). Whenever a sensor receives a convergecast message, it forwards the message in its TDMA slot, and switches to active mode after a timeout. Furthermore, the sensor sets another timer, called *TDMA timer*; when this timer expires, the sensor returns to TDMA mode. In the active mode, the sensor listens to the medium always and forwards messages according to a CSMA protocol (e.g., ReliableComm [8]). Thus, the convergecast algorithm is shown in Figure 1.

```
$j$ executes in TDMA mode
  listen to the TDMA slots of neighbors;
  if $j$ receives a convergecast message $m_c$
    forward $m_c$ in the next TDMA slot;
    setup wakeup-timer
      when timer fires, switch to active mode;
    setup TDMA-timer
      when timer fires, switch to TDMA mode;
$j$ executes in active mode
  listen to the medium always;
  if $j$ receives a message $m$
    forward $m$ using a CSMA based algorithm;
```

Figure 1: Convergecast algorithm

4 SIMULATION RESULTS

We simulated our algorithm in the framework based on Prowler [9], a probabilistic wireless network simulator for MICA motes [10]. First, we discuss the TDMA algorithm, routing and the reliable CSMA protocols used in our simulations. Next, we discuss the simulation model and the results.

Self-stabilizing TDMA (SS-TDMA). We note that our algorithm does not depend on a specific TDMA algorithm. We have chosen SS-TDMA [4] due to its applicability in [1]. Moreover, SS-TDMA can be customized for convergecast.

In [4], the sensors are arranged in a 2 dimensional grid and the base station is located at $\langle 0, 0\rangle$. Sensors have the notion of communication and interference range. Communication range is the distance up to which a sensor can communicate with high certainty. Interference range (say, y) is the distance up to which a sensor can communicate, although the probability of such a communication may be low. Now, the initial slot of a sensor located at $\langle i, j\rangle$ is $(P-1)i+(P-(y+1))j$, where $P=(y+1)^2+1$ is the TDMA period. If x_a is the initial slot of sensor a then a can transmit in slots: $x_a + c * P$, where $c \geq 0$.

Logical grid routing protocol (LGRP). In LGRP [7], the sensors are arranged in a logical grid. Each sensor classifies its neighbors within H hops as low or high neighbors; sensor k located within H hops of sensor j is classified as j's low neighbor if k's distance to the base station is less than that of j's distance. Otherwise, k is classified as j's high neighbor. Each sensor also maintains a variable, *inversion count*. The inversion count of the base station is 0. If j chooses one of its low neighbors as its parent then it sets its inversion count to that of its parent. Otherwise, j sets its inversion count to inversion count of its parent + 1. If the inversion count exceeds a certain threshold then the tree may be corrupted and, hence, j sets its parent to $null$. Sensor j later rejoins the tree if it finds a neighbor that provides a better inversion count.

ReliableComm (RComm). ReliableComm [8] is a CSMA based reliable communication protocol, designed to improve per-hop and end-to-end reliability in presence of fading, collisions and congestion. It maintains a queue, where a given message is removed from the queue when the sensor receives an *implicit acknowledgment* from its parent. If it fails to receive the acknowledgment within a timeout, it retransmits the message. The number of retransmissions is bounded by a threshold. Since the base station does not require to forward the messages, it explicitly acknowledges the receipt of a message.

Simulation model and parameters. In our simulations, we assume that the base station is located at one corner of the grid and the sensors that observe an event are located diagonally opposite to the base station. Furthermore, we require at least 50% of the messages to be delivered to the base station. This value is based on the reliability requirements of the LITeS experiment [1].

We assume that the sensors can communicate with high reliability among their neighbors. Also, we assume that the signal from a sensor may reach sensors within distance 6 although the probability of successful communication is very low. However, the actual radio propagation is based on the distance-fading model, where the strength of a signal from a sensor is inversely proportional to the square root of the distance. The parameters used in our simulations are listed in Table 1.

Table 1: Simulation parameters

Parameter	Value
Network/event.	
Network size	7x7 grid
Sub-grid sending event messages	1x1 - 5x5
Convergecast algorithm.	
Wakeup timeout	1 s
TDMA timeout	20 s
SS-TDMA/LGRP/RComm.	
Time slot interval	50 ms
Interference range	6
H	2
Maximum number of retransmissions	3
Retransmission timeout	55 ms

Results. We compare the performance of our algorithm using LGRP and ReliableComm (i.e., SS-TDMA+RComm+LGRP) with (i) SS-TDMA+LGRP, where sensors execute in TDMA mode always, and (ii) RComm+LGRP, where sensors execute in active mode always.

Latency. In Figure 2, we show the time required to convergecast event messages to the base station. With SS-TDMA+RComm+LGRP and RComm+LGRP, only 50% of the event messages are delivered. The latency of the proposed convergecast algorithm follows closely that of the RComm+LGRP scheme. This meets the requirement of the LITeS experiment, where the desired latency for reporting an event is less than 13 seconds (with 50% reliability) [1].

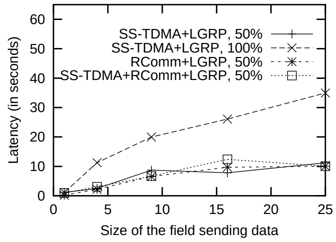

Figure 2: Latency

With SS-TDMA+LGRP, all the messages are delivered to the base station. This is due to the fact that SS-TDMA ensures reliable communication. Moreover, from Figure 2, we observe that the latency for delivering 50% of the messages with SS-TDMA+LGRP is almost equal to the other two schemes. This is due to the fact that SS-TDMA is customized for convergecast. Specifically, the slots assigned to sensors closer to the base station are after the slots assigned to sensors farther from the base station. Hence, whenever a sensor receives a message, it can forward it within 50 ms (= timeslot interval). Thus, the latency is close to RComm+LGRP scheme.

Suppose TDMA is customized for broadcast. Whenever a sensor receives a message, it has to wait at most P slots (=2.5 seconds, in our simulations) before it can forward, where P is the TDMA period. This delay is added at each hop and for each message. Hence, the latency may increase. With SS-TDMA+RComm+LGRP, a sensor switches to RComm once it forwards the first message. Hence, it offers better latency.

Active radio time (ART). In Figure 3, we show the average ART of the sensors. As expected, SS-TDMA+LGRP is energy-efficient, since the sensors listen to the medium only in the slots assigned to their neighbors. In the remaining slots, the sensors conserve energy by turning their radio off. Thus, ART is approximately 40% of the time required for convergecast.

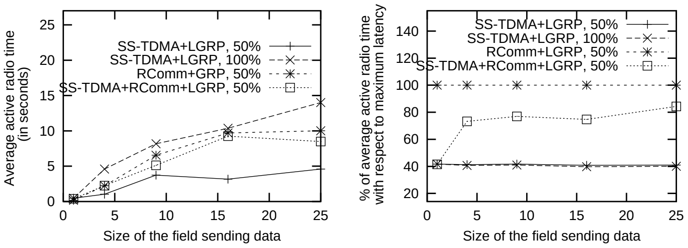

(a) (b)

Figure 3: (a) Average ART and (b) % of average ART with respect to latency. Note that the scale is different in each figure.

With SS-TDMA+RComm+LGRP, the sensors conserve energy until the actual communication starts. From Figure 3(b), the sensors remain in active mode for 76.89% of the time (i.e., 5.13 seconds in 6.67 seconds), when a sub-grid of 3x3 sensors send messages. With RComm+LGRP, the sensors spend 100% of the time in active mode. Hence, it is not energy-efficient.

Message communication. In Figure 4, we show the number of transmissions/receptions during convergecast. The number of transmissions is in order of 500 when a sub-grid of 3x3 sensors send messages with 50% reliability for different schemes. Similarly, the number of receptions for different schemes is approximately the same. In case of SS-TDMA+LGRP, 100% reliability is achieved with increased communication.

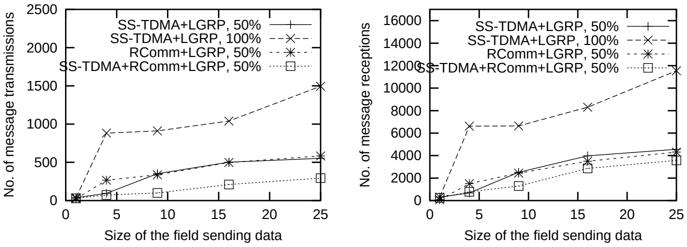

(a) (b)

Figure 4: (a) No. of transmissions and (b) No. of receptions. Note that the scale is different in each figure.

Network lifetime. In Figure 5, we show the analytical estimate of network lifetime with respect to the probability of occurrence of an event at any instant for SS-TDMA+RComm+LGRP. If the probability is 0, the sensors execute in the TDMA mode always. The sensors can turn their radio off in the slots not assigned to itself and its neighbors. Thus, the network lifetime improves by 3.8 times (when the interference range $= 6$ and $H = 2$). When the probability of occurrence of an event increases, ART also increases. As a result, the network lifetime decreases. When this probability is close to 1, the sensors start to operate in the active mode more frequently and, hence, energy conservation is negligible.

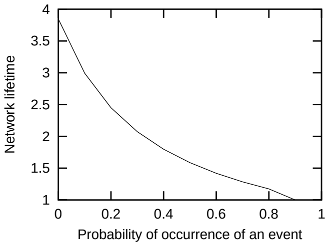

Figure 5: Network lifetime

In a typical application (e.g., LITeS [1, 2]), the probability of occurrence of an event at any instant is less than $10 - 15\%$. From Figure 5, SS-TDMA+RComm+LGRP improves the lifetime of the sensor network by approximately 3 fold.

5 RELATED WORK

Related work that deals with convergecast include [11–14]. In [11], a TDMA based convergecast is investigated. Specifically, the paper proposes a tree construction algorithm for convergecast. Once the tree is constructed, it assigns schedules to the sensors for collision-free communication. By contrast, in our convergecast algorithm, we use an existing TDMA algorithm. Moreover, our solution improves the network lifetime by reducing the amount of idle listening.

In [12], a randomized convergecast algorithm is proposed. This paper identifies the lower bound on the running time for an arbitrary network. Moreover, this paper studies the energy-latency tradeoff. In [13], for an offline problem, dynamic programming based approximation solution is proposed, where the energy dissipation of sensors in the data aggregation tree is minimized. For a real-time scenario, this paper proposes an online protocol for data aggregation. Unlike [12, 13], our solution uses a TDMA based algorithm to conserve energy and to reliably switch to active mode when the network observes events of interest. Furthermore, we show that the network lifetime improves by 3 fold for a typical application.

In [14], a randomized algorithm for convergecast is proposed for ad hoc networks. One of the assumptions in [14] is that the nodes have collision detection capability. By contrast, we do not assume that collisions are detectable. Collision detection may not be possible since the sensors have limited communication capabilities and limited power.

6 CONCLUSION AND FUTURE WORK

In this paper, we presented a convergecast algorithm for sensor networks. This algorithm is designed using a TDMA protocol for energy-conservation and reliable communication. In this algorithm, each sensor is allowed to save energy whenever the network remains idle. And, whenever a sensor receives convergecast messages, the algorithm switches to a CSMA based protocol after forwarding the first message. Thus, it reduces the amount of idle-listening and improves the latency.

We studied the performance of our algorithm and showed that it meets the requirements of a typical application (e.g., LITeS [1, 2]). We showed that ART is within 75% of the time required for convergecast (cf. Section 4). Moreover, for a typical application where the probability of occurrence of an event is less than $10-15\%$, we argued that our solution improves the network lifetime by approximately 3 fold.

There are several questions raised by this work. For one, an interesting question is to analyze the energy-latency tradeoff in large scale networks. Towards this end, we can use a TDMA algorithm that allows the network to be partitioned into several sections [2]. Another issue is to study the effect of multiple base stations and their locations on energy-latency tradeoff.

Acknowledgments. We thank Hongwei Zhang and Young-Ri Choi for providing the Prowler version of their ReliableComm and LGRP programs respectively.
This work was partially sponsored by NSF CCR-0092724, DARPA OSURS01-C-1901, ONR N00014-01-1-0744, NSF EIA-0130724, and Michigan State University grants.

REFERENCES

[1] A. Arora, P. Dutta, S. Bapat, V. Kulathumani, H. Zhang, V. Naik, V. Mittal, H. Cao, M. Demirbas, M. Gouda, Y-R. Choi, T. Herman, S. S. Kulkarni, U. Arumugam, M. Nesterenko, A. Vora, and M. Miyashita. A line in the sand: A wireless sensor network for target detection, classification, and tracking. *Computer Networks (Elsevier)*, 46(5):605–634, 2004.

[2] The Ohio State University NEST team. ExScal: Extreme scaling in sensor networks for target detection, classification, and tracking, 2004. DARPA, `http://www.cse.ohio-state.edu/exscal`.

[3] R. Szewczyk, A. Mainwaring, J. Anderson, and D. Culler. An analysis of a large scale habitat monitoring application. *Conference on Embedded Networked Sensor Systems*, 2004.

[4] S. S. Kulkarni and M. Arumugam. SS-TDMA: A self-stabilizing MAC for sensor networks. In *Sensor Network Operations*. IEEE Press, 2005, to appear. Also, available at: `http://www.cse.msu.edu/~arumugam/research/papers/ieee-sn-operations05`.

[5] T. Herman and S. Tixeuil. A distributed TDMA slot assignment algorithm for wireless sensor networks. *Workshop on Algorithmic Aspects of Wireless Sensor Networks*, 2004.

[6] W. B. Heinzelman, A. P. Chandrakasan, and H. Balakrishnan. An application-specific protocol architecture for wireless microsensor networks. *IEEE Transactions on Wireless Communications*, 2002.

[7] Y-R. Choi, M. G. Gouda, H. Zhang, and A. Arora. Routing on a logical grid in sensor networks. Technical Report TR-04-49, The University of Texas at Austin, Department of Computer Sciences, 2004.

[8] H. Zhang, A. Arora, Y-R. Choi, and M. G. Gouda. Reliable bursty convergecast in multi-hop wireless sensor networks. Technical Report OSU-CISRC-7/01-TR42, The Ohio State University, July 2004.

[9] G. Simon, P. Volgyesi, M. Maroti, and A. Ledeczi. Simulation-based optimization of communication protocols for large-scale wireless sensors networks. *IEEE Aerospace Conference*, 2003.

[10] J. Hill and D. E. Culler. Mica: A wireless platform for deeply embedded networks. *IEEE Micro*, 2002.

[11] V. Annamalai, S. K. S. Gupta, and L. Schwiebert. On tree-based convergecasting in wireless sensor networks. *Wireless Communications and Networking Conference*, 2003.

[12] Q. Huang and Y. Zhang. Radial coordination for convergecast in wireless sensor networks. *Workshop on Embedded Networked Sensors*, 2004.

[13] Y. Yu, B. Krishnamachari, and V. K. Prasanna. Energy-latency tradeoffs for data gathering in wireless sensor networks. *IEEE Infocom*, 2004.

[14] A. Kesselman and D. Kowalski. Fast distributed algorithm for convergecast in ad hoc geometric radio networks. *Conference on Wireless On demand Network Systems and Services*, 2005.