Available at www.ComputerScienceWeb.com
POWERED BY SCIENCE @ DIRECT®

Ad Hoc Networks 1 (2003) 317–331

www.elsevier.com/locate/adhoc

Aggregation in sensor networks: an energy–accuracy trade-off

Athanassios Boulis [1], Saurabh Ganeriwal *, Mani B. Srivastava [2]

Networked and Embedded Systems Lab, EE Department, University of California at Los Angeles, 56-125B Eng. IV Bldg., Box 951594, Los Angeles, CA 90095-1594, USA

Abstract

Wireless ad hoc sensor networks (WASNs) are in need of the study of useful applications that will help the researchers view them as distributed physically coupled systems, a collective that estimates the physical environment, and not just energy-limited ad hoc networks. We develop this perspective using a large and interesting class of WASN applications called *aggregation applications*. In particular, we consider the challenging periodic aggregation problem where the WASN provides the user with periodic estimates of the environment, as opposed to simpler and previously studied snapshot aggregation problems. In periodic aggregation our approach allows the spatial–temporal correlation among values sensed at the various nodes to be exploited towards energy-efficient estimation of the aggregated value of interest. Our approach also creates a system level energy vs. accuracy knob whereby the more the estimation error that the user can tolerate, the less is the energy consumed. We present a distributed estimation algorithm that can be applied to explore the energy–accuracy subspace for a subclass of periodic aggregation problems, and present extensive simulation results that validate our approach. The resulting algorithm, apart from being more flexible in the energy–accuracy subspace and more robust, can also bring considerable energy savings for a typical accuracy requirement (fivefold decrease in energy consumption for 5% estimation error) compared to repeated snapshot aggregations.

Keywords: Sensor networks; Aggregation application; Distributed estimation; Energy vs. accuracy trade-off

1. Introduction

The technological advances in embedded computers, sensors, and radios have led to the emergence of wireless ad hoc sensor networks (WASNs) as a new class of system with uses in diverse and useful applications. Indeed, the early papers in the area [1–4] talk about the vision of cheap self-organizing ad hoc networks that are able to perform a higher level sensing task through the collaboration of a large number of cheaper and resource constrained wireless sensor nodes. Leveraging numerous sensing devices placed close to the actual physical phenomena, the information that such networks can provide is more accurate and richer than the information provided by a system of few, expensive, state-of-the-art sensing devices. Since WASNs operate largely unattended, often in environments where the access cost of deploying or maintaining nodes is high, a key problem in designing WASNs is how to prolong their useful lifetime by conserving energy. Consequently, a large fraction of research in WASNs has

* Corresponding author. Tel.: +1-310-825-7707.
E-mail addresses: boulis@ee.ucla.edu (A. Boulis), saurabh@ee.ucla.edu (S. Ganeriwal), mbs@ee.ucla.edu (M.B. Srivastava).
[1] Tel.: +1-310-825-7707.
[2] Tel.: +1-310-267-2098.

1570-8705/$ - see front matter
doi:10.1016/S1570-8705(03)00009-X

been dedicated to aspects of the energy-efficiency problem.

The original vision and promise of WASNs was that multiple nodes collectively perform the sensing task requested by the users and communicate the results to the users. However, most of the research so far has simply viewed WASNs as just another kind of wireless ad hoc networks, albeit one composed of nodes that are more energy-constrained and whose data sources are sensors. So, for example, much work has focused on issues such as energy-efficient MAC and ad hoc routing protocols to realize the needed point-to-point and point-to-multipoint communication patterns in WASNs. But, little has been done to develop an understanding of a WASN as a *collective* or an *aggregate* where sensor nodes collaborate to jointly estimate the desired answer about the sensed environment. In part this is because not many actual applications useful to the end-user have been studied. The only notable exception is the target-tracking problem, which has drawn attention from several research groups. Otherwise, the applications that have been examined are usually "toy" scenarios used to showcase the abilities of protocols and programming frameworks (e.g., [5]), or very specific applications examined for the sake of some energy-saving technique (e.g., [6]).

In this research we have made a first attempt at exploring and understanding the performance of a WASN as a collective that performs a sensing task. We examine a general class of WASN applications that we call *aggregation applications* where the desired answer depends on the sensed value at multiple nodes. In particular, we explore the energy vs. accuracy subspace, i.e. how much energy savings can one get by relaxing some accuracy requirements and vice versa. We propose an algorithm that exploits this trade-off and jointly considers networking and signal processing issues to create a distributed estimation mechanism.

1.1. Aggregation applications

Many of the examples and simple applications presented in WASN research are based around some kind of aggregation function. The most popular and simple examples of aggregation functions are "maximum" and "average". That is, a user may be interested in knowing the max (or average) of a value in the WASN or in some restricted area of the WASN. If this function needs to be performed once, we refer to it as "snapshot aggregation". If the user needs an update in periodic intervals we refer to it as "periodic aggregation".

The snapshot aggregation problem is trivial for a single static user. The user sends a request to flood the sensor network (or the area of interest). Upon reception of a request message a node sets the sender of the message as its parent, leading towards the user. This way an aggregation tree is formed with the user node at its root. Data are flowing along the aggregation tree towards the user while being aggregated at intermediate nodes. For instance, in the max function a node receiving multiple values (i.e., its own local reading and values sent by other nodes) finds their maximum and sends it to its parent. For more details on snapshot aggregation the reader can refer to [7,8].

More generally, in aggregation applications, the user seeks a condensed view of the physical environment the WASN is monitoring, or a condensed view of the network's state. To achieve this, the values from all the nodes (i.e., sensor readings or node state values) are aggregated to a size-bounded vector describing this condensed view. Furthermore, several important properties hold for the aggregation process: (i) multiple local values can be combined to an aggregated description with a single pass, (ii) multiple aggregated descriptions and multiple local values can be combined to an aggregated description with a single pass. These properties permit the aggregation process to be done easily within the network, without the need for multiple passes of the data. A counter-example is the calculation of median, as it requires two passes of the data. The bound on the aggregated description (i.e., vector) is $O(1)$. In order to include more specific cases of applications (like some referenced in related work) the bound can be relaxed to $O(N)$, where N is the number of nodes in the network.

Some examples more advanced than "max" and "average" include: (i) approximate contours of