Computer Communications 45 (2014) 47–58

Contents lists available at ScienceDirect

Computer Communications

journal homepage: www.elsevier.com/locate/comcom

Distributed energy-efficient estimation in spatially correlated wireless sensor networks ☆

Iordanis Koutsopoulos [a,]*, Maria Halkidi [b]

[a] Athens University of Economics and Business and CERTH, Greece
[b] University of Piraeus, Greece

ARTICLE INFO

Article history:
Received 20 May 2013
Received in revised form 14 March 2014
Accepted 15 March 2014
Available online 27 March 2014

Keywords:
Wireless sensor networks
Distributed optimization
Estimation
Energy efficiency

ABSTRACT

We study the problem of maximizing lifetime in sensor networks that are deployed for estimating an unknown parameter or process. Sensors take measurements and transport them in multi-hop fashion to a fusion center (FC) for Maximum Likelihood (ML) estimation of this process. A prime distinguishing feature of wireless sensor networks is the *spatial correlation* among measurements of nearby sensors. We investigate the way spatial correlation shapes the tradeoff between estimation quality and energy efficiency, according to which more measurement data improve estimation quality, but they are more energy-costly to transport. If the effect of spatial correlation is understood, then a given estimation quality can be achieved with minimum data redundancy and energy consumption. We study the dynamic control of the sensor sampling rates and the routes to the FC. Sensor attributes such as spatial correlation, measurement accuracy and energy reserve as well as the quality of wireless links, collectively affect the sampling rates and routes to the FC. Further, due to spatial correlation, sensor sampling rates alter the joint probability density functions (p.d.f's) for sensor readings and thus they affect the average estimation error. We show that the optimization problem can be decomposed into smaller ones, where each sensor autonomously takes its sampling rate and next-hop forwarding decisions, and we propose an iterative, low-overhead primal–dual algorithm. Our work yields interesting insights on the fundamental tradeoff between network lifetime and estimation quality, it provides a clear intuition on when spatial correlation is beneficial for the tradeoff above, and it provides a solution with distributed sensor coordination.

1. Introduction

The prime reasons for deploying a wireless sensor network are the estimation of an unknown parameter (such as temperature, vibration, air quality, road traffic intensity), or the tracking of a process (e.g. person/particle movement), or the detection and localization of an event (e.g. fire) in the deployment area. Such tasks are mapped onto concrete objectives such as low estimation error or small probabilities of false alarm and missed detection. These objectives are different from conventional ones like throughput maximization or delay minimization, for which wireless networks are deployed.

The salient distinguishing feature of sensor networks is the *spatial correlation* among readings of co-located sensors. Nearby sensors are likely to perceive an event or phenomenon in a similar manner, and thus their measurements are correlated. In this work, we investigate the way spatial correlation shapes the fundamental tradeoff between estimation quality and energy efficiency. If neighboring sensor readings are correlated, sensors should reduce their sampling rates so as to eliminate possible redundancies and thus reduce the amount of consumed energy for data transmission. Interestingly enough though, in our study we find out that this intuition is not always true.

We study the problem of maximizing lifetime in a wireless sensor network that is deployed for estimating an unknown parameter or process with given precision. Lifetime is defined as the time elapsed until the first sensor battery empties. Sensors take measurements and transfer them in multi-hop fashion to a fusion center (FC) which performs Maximum Likelihood (ML) estimation. The intuitive tradeoff is that more measurements improve estimation quality, but they need more energy to transport, and thus they lead to reduced network lifetime. If the impact of spatial correlation is understood, a given estimation accuracy can be achieved with less data redundancy and energy consumption. In order to maximize lifetime, the network must control the *number of endogenous measurements per unit of time* for each sensor

☆ Parts of the material in this paper were presented in [1], in IEEE GLOBECOM 2009 conference and in [2] in WiOpt 2010. This work was performed while I. Koutsopoulos was with University of Thessaly, Greece.

* Corresponding author. Tel.: +30 210 8203933; fax: +30 210 6081271.
E-mail address: jordan@aueb.gr (I. Koutsopoulos).

http://dx.doi.org/10.1016/j.comcom.2014.03.012
0140-3664/

(namely the *sensor sampling rate*) and the *data traffic flow across the network* toward the fusion center (FC). In order to maximize network lifetime, sensor energy consumption rates must be as balanced as possible. On the other hand, sensor sampling rates alter the joint p.d.f of measurements at the FC and thus they affect the expected estimation error.

For a simple motivating example, consider three sensors that are identical in terms of energy consumption and transmit directly to a FC. Assume that 30 measurements total are needed at the FC for given estimation error. Assume that the objective is that sensor energy consumption are as balanced as possible, so that batteries empty with similar rate. As we will see later, balanced energy consumption maximizes network lifetime. If sensor readings are uncorrelated, each sensor should submit 10 measurements. Now, suppose sensors are spatially correlated. Intuition says that each pair of correlated sensors should collectively send fewer than 20 measurements so as to avoid redundancy. If measurements of sensors 1 and 2 are more correlated than those of sensors 2 and 3, this should be reflected on the number of measurements of 1 and 2 versus that of 2 and 3. Consider the extreme case that sensors 1 and 2 are fully correlated (i.e. their measurements are statistically identical up to linear scaling), while sensor 3 is uncorrelated. Sensor 3 sends 10 measurements again. However, since sensors 1 and 2 are fully correlated, only one sensor needs to send 10 measurements (because the other 10 measurements needed to match the 30 required ones would be redundant, due to sensor full correlation). A good choice that balances energy consumption would then be to have sensors 1 and 2 send 5 measurements each. Interestingly, this intuitive benefit of spatial correlation in reducing the amount of data needed to achieve a certain estimation accuracy does not always hold, as we will see later. Determining the precise impact of spatial correlation on sensor measurement rates and the tradeoff between estimation quality and energy efficiency is the main contribution of this work.

1.1. Related work

In information theory, spatially correlated sensor sources are jointly encoded (compressed), so that the minimal necessary amount of bits is transmitted for a given allowable level of distortion. The problem becomes one of deciding on the source coding rates, that is the number of bits that each source should transmit [3, p. 312–315]. For multiple Gaussian sources, it turns out that bits are allotted in an inverse water-filling fashion on the source variances or on the eigenvalues of the covariance matrix, depending on whether sources are uncorrelated or not. For uncorrelated sources, in essence this means that bits are allocated only for describing sources with large enough variance.

Various aspects of estimation objectives have been considered in the literature [4–7]. In [4], the authors consider the class of policies where the number of transmitted bits (and corresponding quantization levels) for each sensor are computed so as to achieve given estimation error with minimum total energy consumption. The work in [6] studies the impact of transmit power on estimation error. The problem of sensor power allocation for minimum estimation error subject to a power constraint, and the one of power minimization subject to an estimation error constraint both possess the water-filling structure: at the optimal solution, only a subset of sensors transmit and the rest are off. In [7], the joint problem of quantization and power allocation is considered, with the objective to minimize total transmit power, while ensuring a given estimation error performance. Here, the estimation error depends on the quantization noise and the sensor measurement noise. Two common assumptions in the works above are that each sensor transmits at most one reading to the fusion center (FC), and that sensor readings are spatially uncorrelated.

In [8], spatial correlation is captured by conditional entropy, and a framework for joint data compression is proposed, that minimizes total cost of data aggregation to sinks. In [9], a similar model is proposed for minimum-cost joint source coding and routing under no contention. A similar problem is studied in [10], in which the authors present an approximation of the joint entropy of sources that captures compressed information and a bit-hop metric that models the cost of joint routing and compression. They deduce that a static clustering scheme results in near-optimal performance. In [11], a distributed optimization method is proposed for joint source coding, routing and access control, so that total energy is minimized. Spatial correlation in sensor networks has also been studied in various contexts [12].

The work [13] exploits spatial correlation for designing efficient medium access control protocols with low contention overhead under a distortion constraint. The authors create a tessellation of the area through vector quantization and they select one representative transmitting sensor from each region. The authors in [14] study sensor activation for maximizing a generic utility function that captures network objectives for rechargeable sensors. Spatial correlation is taken into account through correlated discharge and recharge processes. Neighboring sensors that perceive a local phenomenon in a similar manner have their batteries emptied almost at the same times. In [15], the authors study the power minimization problem of transmitting spatially correlated sensor data to a FC subject to an accuracy constraint. Finally, in the recent work [16], the problem of maximizing a metric of quality of monitoring is considered, through sensing rate and route decision, under flow conservation and energy constraints at each sensor node.

A seminal work on lifetime maximization in wireless sensor networks is [17], where routing of sensor data with *given sampling rates* was addressed. Network lifetime was defined as the time elapsed until the first sensor battery empties. The work [18] considers the problem of sleep-wake sensor scheduling and next-hop packet forwarding for maximizing network lifetime subject to an expected end-to-end packet-delivery delay. The work in [19] finds a maximum-lifetime aggregation tree among shortest-path ones. Routing and sensor placement to maximize network lifetime were studied in [20], where the distortion depends on sensor positions relative to cluster-heads. Rate-distortion theory was used for formulating the lifetime-distortion tradeoff in [21] and for quantifying the impact of spatial correlation on total energy consumption and lifetime in a cluster-based transmission model [22]. In the work [23] which is included in [24], lifetime is defined as the time until which the network fulfills the estimation error constraint. The authors find the quantization level, the number of measurements per sensor and the multi-hop routes to the fusion center that maximize lifetime in the sense above. The work [25] studies cooperative routing for detecting a random field with a proposed link metric that optimally aggregates data. The joint control of data queues and battery energy queues is studied in [26], where the goal is to maximize the long-term sensing rate under energy replenishment regimes. Finally, comprehensive survey on research issues about distributed inference in wireless sensor networks may be found in [27]. Finally our prior work [2] addresses the problem of lifetime maximization in wireless sensor networks, where sensors are allowed to perform local aggregation of incoming traffic flows.

1.2. Our contribution

In this paper, we study lifetime maximization in multi-hop wireless sensor networks with given estimation error specification. The first key novelty lies is a distributed optimization algorithm that encompasses control of sensor sampling rates and network flows in order to optimize network lifetime under a given estimation error constraint. The second novelty is the explicit