Computers, Environment and Urban Systems 33 (2009) 79–89

Contents lists available at ScienceDirect

Computers, Environment and Urban Systems

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

Optimizing query processing using selectivity-awareness in Wireless Sensor Networks

Muhammad Umer *, Lars Kulik, Egemen Tanin

National ICT Australia, Department of Computer Science and Software Engineering, University of Melbourne, Victoria 3010, Australia

ARTICLE INFO

Keywords:
Wireless Sensor Networks
Monitoring
Query optimization
Selectivity
Localized Steiner tree approximation

ABSTRACT

Monitoring queries are fundamental for Wireless Sensor Networks (WSNs) that collect data for physical phenomena. In this work we address three key characteristics of monitoring queries. First, a monitoring query can be selective, i.e., it requests readings only from parts of a WSN. Second, a monitoring query can be continuous, i.e., it draws sensor readings for long periods of time. Finally, since physical phenomena are spatially correlated, a monitoring query selects spatially co-located nodes. In our earlier work, we proposed the Pocket Driven Trajectories (PDT) algorithm; a selectivity-aware data collection technique that tailors data collection paths for a monitoring query based on the spatial layout of selected nodes. In this work, we extend the basic PDT algorithm with an adaptive behavior. We show that the enhanced PDT algorithm is ideal for real world WSNs due to its two major strengths; the PDT algorithm is local, i.e., it does not require any global information about node locations or network connectivity. Furthermore, the PDT algorithm efficiently adapts its data collection paths over the lifetime of a query as changes in the spatial layout of selected nodes occur. Using extensive simulations, we show that in terms of energy efficiency the PDT algorithm clearly outperforms well-known WSN data collection algorithms.

1. Introduction

Recent advances in hardware miniaturization and wireless communications have given rise to a new class of networked systems known as Wireless Sensor Networks (WSNs). These systems are empowered by a large number of self-contained sensor nodes, each smaller than a human palm, providing an unprecedented capability to sense and monitor the physical world. The low cost of the sensor nodes establishes an economy of scale that makes it possible to deploy these nodes in large numbers over vast areas. WSNs are finding applications in diverse fields such as volcanic activity monitoring, tsunami detection, and ecological monitoring (NASA Volcano Sensorweb, 2007; Secure CITI, 2006; Tolle et al., 2005). As in other large computing environments, the success of a WSN is tied to its sustainability, i.e., the ability of the system to maintain itself over time. Typically, the main constraint for a sensor node is its power. In a large WSN deployment, it is often not viable to change the batteries of the nodes. Therefore, minimizing power consumption is a prime objective common to all WSN operations, including sensing, data processing, and communication.

Monitoring of physical phenomena is one of the main areas of application for WSNs. In such applications, a database-oriented view of WSNs has proven to be useful. According to this view, a WSN is considered as a distributed data source where the sensed values generated by a set of WSN nodes form the rows of a relation split across all nodes in the set (Madden, Franklin, Hellerstein, & Hong, 2003; Yao and Gehrke, 2003). The database-oriented view motivates the design of WSN data acquisition regimes targeted at two fundamental aims. First, similar to classical database systems, a WSN database should provide SQL-like abstractions so that nodes can be easily programmed for data sensing and collection. Second, the data collection process should minimize the overall energy expenditure. Research into current sensor hardware has shown that the energy consumption of a sensor node is a function of its communication workload (Pottie & Kaiser, 2000). For a data acquisition system, this insight motivates the optimization of the communication process either by involving a smaller number of nodes, or by establishing efficient communication paths between the sensor nodes and the base station. In this paper, we consider these two techniques in tandem, and argue that certain distinct features of monitoring queries can be exploited to achieve higher levels of energy efficiency.

1.1. Characteristics of monitoring queries

We define a monitoring query as a continuous data collection task that requests sensed values from nodes fulfilling selection criteria based on certain physical conditions. For instance, queries

* Corresponding author.
E-mail addresses: mumer@csse.unimelb.edu.au (M. Umer), lars@csse.unimelb.edu.au (L. Kulik), egemen@csse.unimelb.edu.au (E. Tanin).

0198-9715/$ - see front matter © 2009 Elsevier Ltd. All rights reserved.
doi:10.1016/j.compenvurbsys.2009.01.006

that monitor herd movements in a farm would report the sensed values only from the nodes that have recently sensed animal movements. The following are key aspects of monitoring queries:

- *Monitoring queries are selective*: typically, a WSN can cover an area much larger than the area of interest at any given point in time. For instance, in the herd monitoring example presented above, although nodes are present all over the farm, the herd may only be found within a limited area. We argue that for energy efficient optimization of such queries, the data collection task should be *selectivity-aware*.
- *Monitoring queries are continuous*: a monitoring task, by design, is expected to query readings from sensor nodes over an extended period of time. The mainstream WSN database systems have realized the need for continuous queries and provide SQL clauses to define such queries. For instance, the Cougar database system provides DURATION and EVERY clauses (Yao & Gehrke, 2003), which specify the lifetime of a query and the rate of the query answers, respectively.
- *Monitoring queries select spatially correlated nodes*: physical phenomena are characterized by their spatial correlation; hence, when monitoring a physical phenomenon, sensor nodes at proximal locations tend to have similar values. Therefore, this spatial correlation, coupled with the notion of selectivity, results in clustered or pocketed node participation. For instance, if a node is selected by a query based on its sensed temperature value, there is a high probability that neighboring nodes will also be selected by the same query.

1.2. Our contributions

In this work, we present an adaptive location-based data collection algorithm, Pocket Driven Trajectories (PDT), which optimizes the monitoring queries by exploiting their key characteristics. The PDT algorithm first discovers the set of pockets of selected nodes for a given query. It then aligns a data collection tree to the spatially optimal path connecting these pockets. This path minimizes the use of non-selected nodes in the data collection tree and consequently optimizes the overall energy efficiency of the data collection process. The PDT algorithm is *localized*; i.e., no computation in the algorithm requires any global information such as node connectivity or locations. Moreover, the PDT algorithm is *adaptive*; i.e., it continuously adapts the data collection tree to changing node participation. The PDT algorithm does incur set-up and tree adjustment costs; however, due to the continuous nature of the monitoring queries, these costs can be amortized over the lifetime of the query.

A basic version of the PDT algorithm was presented in one of our earlier works (Umer, Kulik, & Tanin, 2006). The basic version establishes the strength of selectivity-awareness in the optimization of monitoring queries. However, it does not adapt to changing node participation and environmental conditions. The basic version is thus limited in its potential applications. In this paper, we extend and improve our previous work by introducing adaptability into the algorithm and analyze the performance for various dynamic situations.

1.3. Organization

The remainder of this paper is organized as follows. In Section 2 we review related research and position our work in the context of the WSN literature. Section 3 presents the formal underpinnings of our algorithm, including a discussion of its adaptability and costs. Section 4 details the experiments and discusses the results of the simulation-based analysis of our algorithm. Finally, Section 5 presents the conclusions and discusses possible future directions.

2. Related work

Data collection has been an active research topic in the WSN community. As shown in Fig. 1, we identify three major trends in WSN data collection research: *Systems*, *Testbeds*, and *Algorithmic*. Systems research has been mainly led by data collection systems such as Directed diffusion (Intanagonwiwat, Govindan, Estrin, Heidemann, & Silva, 2003), TinyDB (Madden et al., 2003), and Cougar (Yao & Gehrke, 2002). The main motivation behind these systems is to simplify the users' access to a WSN by providing abstractions at the communication, routing, and node programming levels. Although these systems are based on a number of data collection algorithms, new applications and evolving requirements have generated a constant need for better algorithms. This need has thus motivated the algorithmic research in the WSN domain. Since our work also falls under the category of algorithmic research, we discuss this category in detail below.

2.1. In-network data aggregation

In-network data aggregation was one of the first data collection optimization approaches proposed for WSNs (Bonfils & Bonnet, 2003; Considine, Li, Kollios, & Byers, 2004; Madden, Franklin, Hellerstein, & Hong, 2002; Nath, Gibbons, Seshan, Anderson, 2004). In-network aggregation algorithms exploit the fact that a sensor node consumes less energy for information processing than for information communication. These techniques establish multi-hop data collection paths in a network where a node first aggregates the incoming packets of the nodes in communication range and communicates only the aggregated information to the next node. Some in-network aggregation techniques, such as TAG (Tiny AGgregation) (Madden et al., 2002, 2005), randomly create the data collection path and use it to compute all aggregate queries. In contrast, several recent approaches propose to tailor the data collection paths specifically to a group of similar queries (Bonfils and Bonnet, 2003; Pattem et al., 2004; Sharaf, Beaver, Labrinidis, & Chrysanthis, 2004). A data collection path can be tree-based or multi-path (graph) based (Bawa, Gionis, Garcia-Molina, Motwani, 2004; Nath et al., 2004). Interested readers are referred to (Umer et al., 2006) for a detailed review of major in-network data aggregation schemes.

2.2. Suppression-based data collection

Similar to data aggregation, suppression-based methods also attempt to reduce the amount of data required to be reported to the base station. As a spatial suppression method, clustered in-network aggregation exploits the spatial correlation of sensor readings to preserve energy (Pattem et al., 2004; Xu, Heidemann, & Estrin, 2001; Yoon & Shahabi, 2005). Data suppression can also be achieved by approximating physical phenomena using statistical models (Chu, Deshpande, Hellerstein, & Hong, 2006; Deshpande, Guestrin, Madden, Hellerstein, & Homg, 2004). For instance, the BBQ framework uses a statistical model along with live data acquisition so that a large number of queries can be answered locally by the base station (Deshpande, Guestrin, Madden, Hellerstein, & Homg, 2004). Similarly, *Ken* (Chu, Deshpande, Hellerstein, & Hong, 2006) proposed a dynamic probabilistic model for WSN data.