Poster Abstract: Towards Adaptive Wireless Sensor Networks

Silvia Santini
Institute for Pervasive Computing, ETH Zurich
8092 Zurich, Switzerland
santinis@inf.ethz.ch

1. INTRODUCTION

Being deployed close to the real world, wireless sensor networks allow to monitor the evolution in time and space of a variety of real world phenomena. The sensor nodes, covering a given region of interest, can perform at typically regular time intervals a distributed sampling of some measurable physical quantity. At each time instant, the information residing in the network can thus be visualized as a snapshot of this physical phenomenon. Reporting these snapshots to an interested user requires the sensor nodes to relay their sensor readings to a sink node and thus imply a significant communication overhead. Since radio communication is known to be the dominant factor of energy consumption in wireless sensor networks [7], efficient data gathering strategies are needed to conserve the typical poor network resources. One common approach to reduce the amount of data that need to be delivered to the user is to select, among all data produced by the sensor network, a subset of sensor readings such that the original observation data can be reconstructed within some user-defined accuracy. Exploiting spatio-temporal correlation among data, for example, it is possible to identify a subset of sensor readings from which the remaining measurements can be predicted within a given minimal accuracy. Readings which can be predicted from already delivered data do not need to be reported to the sink, thus reducing communication. Prediction can be performed in both time and space for example on the basis of some predefined model, as proposed in a series of recent papers [2, 3, 5]. While model-based prediction techniques have proven to be effective for reducing communication in wireless sensor networks, they not only suffer from performance losses when the model becomes outdated, but are also not well-suited to follow fine grained changes in sensor readings. To overcome these drawbacks we refer to the classic adaptive filter theory and propose an alternative solution for performing predictions over data streams which do not require any a-priori knowledge about the phenomena of interest [4]. Adaptive filters are typically used in environments where signals with unknown and non-stationary statistics are involved and appear for this reason particularly suited to be used in highly dynamic systems like sensor networks.

We present a prediction-based strategy for performing efficient data gathering in sensor networks that exploits the Least-Mean-Square (LMS) adaptive algorithm. The LMS is an adaptive algorithm with very low computational overhead and memory footprint that – despite its simplicity – provides excellent performance. We show that our strategy can significantly reduce the number of readings that a sensor node is required to report to a sink node, while ensuring that the user can reconstruct the original observation data within a pre-specified minimal accuracy e_{max}.

The work presented in this paper constitutes a preliminary step toward the realization of a fully adaptive wireless sensor network, which we envision to be able to autonomously adapt to environmental and network dynamics and correspondingly allocate network resources. The LMS-based prediction scheme can indeed be hierarchically extended to perform a *joint* prediction over a block of readings from neighboring nodes. Since a joint prediction can capture spatial correlation among neighboring sensors, a further reduction in communication can be achieved.

In the next section we briefly explain our adaptive, quality-based strategy for data reduction in wireless sensor networks, while in Section 3 we report our experimental results and point out some open issues and directions for further research.

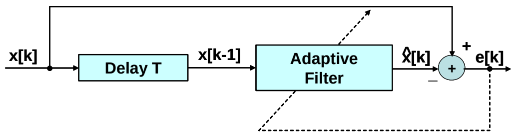

Figure 1: Adaptive predictor.

2. OUR APPROACH

Figure 1 shows the general structure of an adaptive filter used as a predictor. The basic principle consists in computing an estimation $\hat{x}[k]$ of the current input value $x[k]$ as a weighted sum of the previous N samples:

$$\hat{x}[k] = \sum_{i=1}^{N} w_i[k] * x[k-i] \qquad (1)$$

where $w_i, i \in \{1, 2, \ldots, N\}$ are the so-called *filter weights*. The actual prediction error is thus easily computed as $e[k] = \hat{x}[k] - x[k]$ and fed back to the adaptation algorithm, which accordingly updates the filter weights. In order to minimize the average error power $E\{e^2[k]\}$, the LMS algorithm modifies at each time step the filter weights according to the following update equation:

$$\underline{w}[k+1] = \underline{w}[k] + \mu \underline{x}[k]e[k] \quad (2)$$

where the *step-size* μ is a parameter that controls the convergence speed of the algorithm and can be easily estimated from the input signal power [4, 6]. In equation 2 above, $\underline{w}[k]$ and $\underline{x}[k]$ represent the vectors: $[w_1[k], w_2k, \ldots, w_N[k]]^T$ and $[x[k-1], x[k-2], \ldots, x[k-N]]^T$. From equations 1 and 2 it can be seen that, provided the number of filter weights N to be sufficiently small, the LMS algorithm has a very low computational load, since it only needs $2N$ operations for computing the prediction and $2N+1$ operations for the weights updating.

In our scheme, the LMS algorithm runs on both a sensor and a sink node. As long as the prediction error exceeds a user-defined threshold e_{max}, the node works in *normal* mode, i.e., it keeps computing the prediction at each time step, transmitting its readings to the sink node and feeding back the prediction error to the weight adaption algorithm. As long as the node remains in *normal* mode, the sink lets the prediction filter run over the received sensor readings in order to update the filter weights coherently with the node[1]. If, while working in *normal* mode, the sensor node observes the prediction error remaining below the threshold e_{max} for N consecutive readings, it switches to *stand-alone* mode, i.e., it just discards the reading and inputs the predicted value $\hat{x}[k]$ to the filter instead of the actual value $x[k]$. The sink, receiving no data from the node, implicitly assumes the predicted readings being a good enough approximation of the real readings and keeps thus running the prediction filter on these values, as well as the node does. If at a given time step the prediction error, continuously monitored by the sensor node, exceeds the threshold e_{max}, the node switches again to *normal* mode. We will show in the next section, that implementing this simple dual prediction scheme with the adaptive LMS algorithm we are able to achieve significant communication reduction.

3. EVALUATION AND OPEN ISSUES

In this section we present the results obtained when applying our LMS dual prediction scheme to real world data, as well as some open issues which are object of our current research.

We tested our LMS-based data reduction strategy on a set of real world data publicly available at [1]. Once every 31 seconds, humidity, temperature, light and voltage values where collected from 54 Mica2Dot sensor nodes deployed in the Intel Berkeley Research Lab between February 28 and April 5, 2004. To report our results, we picked 4 of these 54 motes, namely motes 1, 11, 13, and 49 which were distributed in different sectors of the deployment area. We applied our scheme to the data reported by the temperature sensors of these four motes between March 6 and 9.

In Figure 2 we report the percentage of sensor readings that the four selected motes need to report as the error budget e_{max} increases, for a filter length $N = 4$. As it can be seen, a minimal accuracy of $0.5°C$ can be guaranteed while transmitting only about 10% of the collected sensor readings. This significant data reduction is due to the optimal tracking capability of the LMS algorithm. We would like to point out that in our experiments no significative changes in the performances are observed when varying the number of filter weights fron $N = 4$ to $N = 10$. Moreover, using these small values of N allow to keep extremely low the computational overhead and memory footprint of the algorithm.

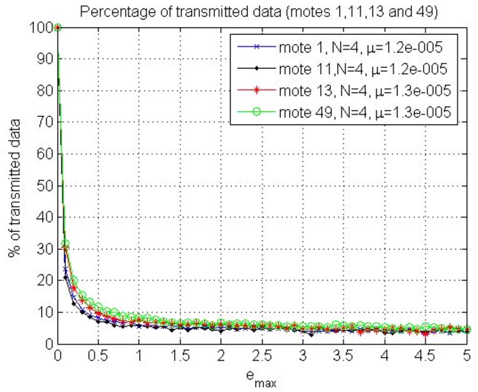

Figure 2: Percentage of transmitted data by motes 1, 11, 13, and 49, for a filter length $N = 4$.

In order to relax the assumption of loss-free communication, we are currently investigating an alternative approach, where, based on the actual packet loss ratio p of the link between source and sink (known using either an a priori estimate or in situ measurements), a modified *equivalent* error budget $e^{eq}_{max} = f(p, e_{max}) \leq e_{max}$ is computed and used instead of e_{max}, such that, on average, accuracy e_{max} can be achieved even in case of message loss. This solution will of course cause some deterioration with respect to the performances presented above, but will also allow to apply and test our approach in a more realistic framework.

Moreover, we are currently working on an extension of our prediction scheme to perform prediction jointly in both the time and space domain.

[1] Since wireless transmission in sensor networks is known to be unreliable, it is likely to happen that some of the node's reported readings never reach the sink node, thus causing a misalignment between the prediction filter at the node and the correspondent filter at the sink. For simplicity, we assume for the remainder of this section a loss-free communication link between the nodes and the sink. See section 3 for a discussion of a mechanism that allow to relax this assumption.

4. REFERENCES

[1] Intel lab data: *http://db.lcs.mit.edu/labdata/labdata.html*.

[2] A. Deshpande, J. H. C. Guestrin, S.R. Madden, and W. Hong. Model-driven data acquisition in sensor networks. In *VLDB'04*, Toronto, Canada, 2004.

[3] S. Goel and T. Imielinski. Prediction-based monitoring in sensor networks: Taking lessons from mpeg. *ACM Computer Communication Review*, 31(5), 2001.

[4] S. Haykin. *Least-Mean-Square Adaptive Filters*. Wiley-Interscience, New York, 2003.

[5] A. Jain, E. Chang, and Y.-F. Wang. Adaptive stream resource management using kalman filters. In *ACM SIGMOD/PODS Conference (SIGMOD '04)*, Paris, France, June 2004.

[6] G. Moschytz and M. Hofbauer. *Adaptive Filter*. Springer Verlag, Berlin, 2000.

[7] J. Polastre, R. Szewczyk, and D. Culler. Telos: Enabling ultra-low power research. In *IPSN/SPOTS'05*, Los Angeles, USA, April 2005.