A FLEXIBLE DISTRIBUTED WAVELET COMPRESSION ALGORITHM FOR WIRELESS SENSOR NETWORKS USING LIFTING

Alexandre Ciancio and Antonio Ortega

University of Southern California
Signal and Image Processing Institute
ciancio@usc.edu, ortega@sipi.usc.edu

ABSTRACT

We address the problem of compression for wireless sensor networks, where each of the sensors has limited power, and acquires data that should be sent to a remote central node. The final goal is to have a reconstructed version of the sampled field at the central node, with the sensors spending as little energy as possible. We propose a distributed wavelet algorithm, based on the lifting scheme, as a means to decorrelate data at the nodes by exchanging information between neighboring sensors. A key result of our work is that by using a locally adaptive distributed transform it is possible to optimize overall power consumption by operating at the right trade-off point between local processing and transmission costs.

1. INTRODUCTION

One of the first applications using sensor networks involved deploying acoustic sensors at the ocean bottom to detect and keep track of submarines. Nowadays, with the advance of technology, disposable sensors with processing capabilities can be deployed in a number of environments to perform tasks such as target tracking (e.g. vehicles, chemical agents, or personnel), traffic control, environment monitoring and surveillance [1]. Over the past few years, the decreasing cost for wireless microsensor networks has brought a lot of attention to the development of a variety of distributed algorithms to enhance the performance of such networks [2, 3, 4, 5].

Assume that a number of power-constrained sensors are spread over an area, acquiring data. The data can consist, for instance, of temperature measurements. A central node is supposed to provide an estimate of the temperature in each point of that area based on transmissions from the sensors (see Fig. 1).

A simple and naive design would be for each sensor to just transmit a quantized version of its own measurement to the central node. However, this approach would not be exploiting the fact that measurements originated from spatially close sensors are likely to be correlated, and energy would be wasted with the transmission of redundant data to the central node.

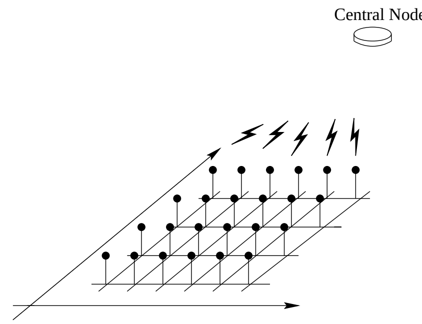

Fig. 1. A number of sensors spread over an area send correlated data to a remote central node.

As an alternative, since data is correlated, it would be reasonable to try to use some sort of transform in an attempt to decorrelate the information from sensors, and, therefore, represent the measurements using fewer bits. However, sensors have access only to their own data, and to compute a transform in the distributed network scenario, we need to define a distributed transform. Such a transform requires inter-sensor communication, so nodes would access the data necessary to compute the tranform coefficients. This means that extra power would be consumed in the form of local processing and additional transmission of information to other sensors. To illustrate this trade-off, consider the effect of choosing trasnforms of different sizes. In general, larger transforms will tend to provide better decorrelation, but at the expense of added communication cost between sensors. For example, using a block transform of size N would mean that N sensors would need to exchange information, with an average communication distance greater than for, say, an $N/2$ size transform.

In our work, we propose a distributed wavelet transform, based on the lifting scheme. The lifting factorization pro-

This work was supported in part by CAPES, Brazil.

vides a convenient representation of the transform as it assumes in place computation (each sensor represents a single memory location), and explicitly breakes down the transform into elementary operations that can be easily evaluated in terms of communication costs. The key idea is that the overall performance of the system depends mainly on the local processing cost and on the communication cost, which is related to (i) the correlation of the signal, (ii) the distances between sensors and (iii) the distances between sensors and the central node. Our algorithm takes those costs into account, and has flexibility to choose the right level of transformation for each environment. By operating at the best trade-off point between processing and transmission cost, given the network requirements, our system would be capable of reducing the overall energy consumption. This would be achieved by assigning different tasks to different groups of sensors. For instance, some nodes might be much closer to the central node than others. Those nodes would be good candidates to perform direct transmission instead of a distributed transform, depending on where the trade-off point is. Some systems could have a very low energy constraint. For those systems, diferent levels of transform decomposition could perform better. The management of coefficients close to the block boundaries can be addressed by use of efficient parallelization algorithms for the lifting scheme, as proposed in [6].

It is interesting to mention that for the ideal decorrelation case, a Karhunen-Loève transform could be computed. However, even if the sensors are grouped in small sized blocks, the KLT would still require one of the sensors to have knowledge about *all* the measurements inside a block, which could potentialy increase the communication costs. Therefore, altough the KLT is flexible in the sense that different block sizes could be chosen, the number of inter-sensor communication could still be too high. Analysis of distributed approximations to the KLT can be found in [3].

In [2, 7] decorrelation is also achieved by means of a distributed wavelet algorithm. However, the dependency between inter-sensor transmission costs and their distance is not considered, and, depending on the number of stages of the wavelet decomposition, measurements from a sensor that is far apart might be needed, elevating power consumption to an unacceptable level. Also, in the performance analysis for that algorithm, the very restrictive assumption that all the information of interest is located in the low-pass subband is made.

One idea, proposed in [8], was based in the use of coset codes to decorrelate data, and did not require inter-node communication. However, although no extra energy is spent with local communications, overall performance might not be as good as a transform-based method, due to the limitations of DPCM. One possible drawback of such a system would be that if the actual correlation berween two given sensors is below that the one used in designing the corresponding Wyner-Ziv encoders, then the resulting decoding will have higher distortion, which can then propagate the error th neighboring sensors.

The paper is organized as follows. In Section 2 the lifting algorithm and its use in the proposed algorithm are described. In Section 3 we discuss how comparisons between methods were made, addressing energy dissipation for both transmission and local processing. Section 4 presents some preliminary results. In Section 5 we present our conclusion and discuss future work.

2. THE PROPOSED ALGORITHM

The lifting scheme is an alternative method to compute biorthogonal wavelets (Fig.2). It allows a faster implementation of the wavelet transform, along with a full in-place calculation of the coefficients [9].

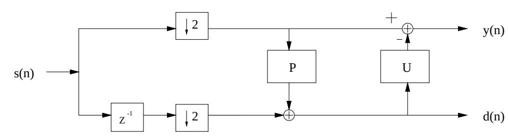

Fig. 2. Lifting implementation of wavelet transform. P corresponds to the prediction stage, and U to the update stage.

We propose the use the lifting scheme to generate the 5/3 wavelet coefficients at each of the sensors. Even-numbered sensors would correspond to the even samples and odd-numbered sensors to the odd samples. In-place computation would reduce the memory requirements for the sensors. Also, as mentioned before, another very attractive property of the lifting scheme in this distributed network scenario is the use of efficient parallelization algorithms. Implementations as the one proposed in [6] enable the partial computation of coefficients at the boundaries. In this perspective, the sensors would be divided into groups and increasing levels of the wavelet decomposition would be computed, as long as the extra energy spent in the calculations and local transmissions were compensated by the increase in decorrelation. At this point, the partial coefficients would be transmitted to the central node. Since each group would be essentialy independent of the others, the system could be setup such that each of them could be assigned different tasks (such as different levels decompositions, or even direct transmission), depending on the system configuration, optimizing the overall performance.

In our implementation, each sensor would only need data from its neighbors at a given scale. That means the lifting scheme for the wavelet can be performed in a two-step distributed way, as seen in Fig. 3.

During the first step, the odd-numbered sensors receive the data measurements from their even-numbered neighbors,

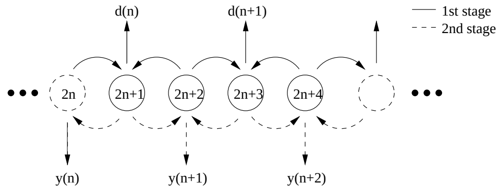

Fig. 3. Two-step implementation of wavelet transform using lifting.

and compute the correspondent detail coefficient. On the second step, these coefficients are sent to the even-numbered sensors, and to the central node. The even sensors use them (along with their own measurement) to generate the smooth coefficients, which are then transmitted to the central node. A more detailed view of the process, including the predictor and update coefficients for the 5/3 wavelet can be seen in Fig 4. In this figure, each of the branches of the graph would have an associated transmission cost.

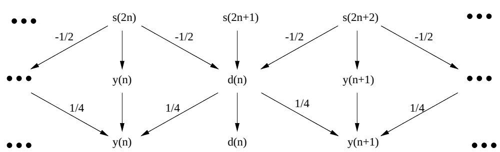

Fig. 4. Lifting Architecture for 5/3 wavelet

3. THE COST ESTIMATION

In order to fairly compare a non-distributed approach (direct transmission) and the proposed algorithm, a cost function that takes into account both local processing and transmission costs has to be defined. For the non-distributed case, no energy is spent with local processing or inter-sensor transmissions. The quantized measurements are simply transmitted directly to the central node. In the distributed approach, we introduce local processing and additional data transmission (over very short distances) in the hope that the obtained decorrelated data will require significantly less energy to be transmitted to the central node, compensating the extra energy spent with the decorrelation process.

Since energy dissipation for both transmission and processing is highly dependent on the processor being used, we take as an example the StrongARM SA-1100, described in detail in [4]. For this DSP, the energy dissipated with the transmission and reception of a k-bit packet over a distance D is

$$E_{Tx} = E_{elec} \cdot k + \varepsilon_{amp} \cdot k \cdot D^2$$
$$E_{Rx} = E_{elec} \cdot k$$

where $E_{elec} = 50nJ/b$ and $\varepsilon_{amp} = 100pJ/b/m^2$.

The energy dissipation due to computation is a function of the supply voltage $E_{lp} = NCV_{dd}^2$, where N is the number of clock cycles per task,C is the average capacitance switched per cycle, and V_{dd} is the supply voltage. For the StrongARM SA-1100, C is approximately $0.67nF$.

The total energy dissipated at each sensor will, therefore, be split into four main components:

$$E = E_{lp} + E_{lt} + E_{lr} + E_{rt}$$

where lp, lt, lr and rt stand for local processing, local transmission, local reception and remote transmission respectively. Obviously, for the non-distributed case, $E_{lp} = E_{lt} = E_{lr} = 0$.

4. SIMULATION AND RESULTS

In this section we present some preliminary results. We show that the proposed scheme gives us a flexible way of meeting requirements by choosing between different system configurations based on the trade-off point between processing and transmission.

To illustrate this point, the simulations considered a simple 5/3 wavelet, as described in section 2. The input process data was created using a second order AR model, with poles placed such that a reasonably smooth output would be generated from white noise (poles were at $0.99e^{\pm j\frac{\pi}{64}}$). The measurements at the sensors corresponded to a sampling of the output of the AR model, and consisted of 100 sensors. The sensors and the central node were assumed to be placed as in Fig 5. For the sake of simplicity, we considered $d_i \approx D$. Computation of actual distances in a practical application is straightforward. We considered, for this example, $d = 5m$ and $D = 1000m$. We used uniform quantization and no entropy coding at this point.

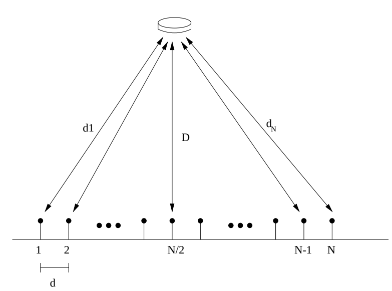

Fig. 5. Placement of sensors and central node

Also, based on the fact that a general processor can typically perform 150 instructions per bit communicated [4] energy-wise, and that the computation of a wavelet coefficient using the lifting scheme takes only 2 multiplications and 4 additions [10], we conclude that $E_{Tx} + E_{Rx} \gg E_{lp}$, and use, for this particular simulation, $E = E_{lt} + E_{lr} + E_{rt}$.

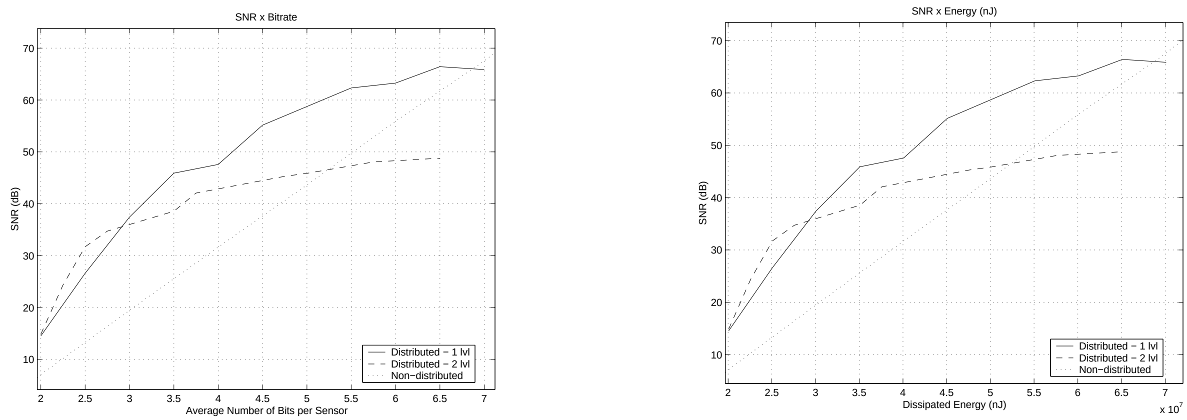

(a) SNR x Bitrate

(b) SNR x Energy consumption

Fig. 6. Distortion curves for simulation case.

The SNR values plotted on the following figures were averaged over a large number of different inputs generated from the same AR model. Fig. 6 shows the distortion curves against the bit rate and the dissipated energy for the cases of a 1 and 2-level wavelet decomposition. The similarity of the graphs comes from the fact that since $D \gg d$, the term E_{rt} dominates over $E_{lt} + E_{lr}$.

The figure also shows how different trade-off points can be used to choose between 1 and 2-level distributed or the non-distributed case. For high SNRs, the non-distributed method might be preferable over the non-distributed. However, SNRs of the order of 60dBs are usually far above typical system requirements, not to mention that they would also come at a considerably higher energy cost.

In the 2-level decomposition case, the inter-node communication cost was increased, but now more sensors have low-energy (detail) data, that can be coded using less bits. It can be seen that, in this particular example, for SNRs of above about 35dB, a 1-level decomposition performs better than a 2-level for the same energy dissipation. On the other hand, for a network with very restrictive energy consumption constraints, a 2-level decomposition would give lower distortion than one with only 1-level. In this example, energy savings can achieve values of up to 40%.

5. CONCLUSION AND FUTURE WORK

In this paper we have proposed a distributed wavelet compression algorithm for wireless, power-constrained, sensor networks. Preliminary results have shown that, by introducing inter-sensor communication over short distances, we can decorrelate data, and reduce the overall energy consumption of the network, while still achieving good distortion. We have used the lifting scheme to compute wavelet transforms as way of decorrelate data. The proposed algorithm allows a flexible way of exploiting trade-off points between processing and communication costs, so that energy savings can be achieved while maximizing the network performance for given specifications. The lifting scheme also enables the computation of partial coefficients [6], with the possibility of reorganizing the network into clusters, for even more efficient power consumption [4].

A number of topics are still to be addressed. Measurements at the sensors can consist of vectorial data, and the local correlation at each sensor could also be exploited. Clusterization of the network, as proposed in [4] along with efficient parallelization algorithms [6] seems the natural next step.

6. REFERENCES

[1] C. Chong and S. P. Kumar, "Sensor networks: Evolution, opportunities, and challenges," *Proceedings of the IEEE*, vol. 91, no. 8, pp. 1247–1256, August 2003.

[2] Sergio D. Servetto, "Distributed signal processing algorithms for the sensor broadcast problem," *Conference on Information Sciences and Systems, The Johns Hopkins University*, March 2003.

[3] M. Gastpar, P. Dragotti, and M. Vetterli, "The distributed karhunen-loève transform," in *Proceedings of the 2002 International Workshop on Multimedia Signal Processing*, St. Thomas, US Virgin Islands, December 2002.

[4] A. Wang and A. Chandraksan, "Energy-efficient dsps for wireless sensor networks," *IEEE Sigal Processing Magazine*, pp. 68–78, July 2002.

[5] S. S. Pradhan, J. Kusuma, and K. Ramchandran, "Distributed compression in a dense microsensor network," *IEEE Sigal Processing Magazine*, pp. 51–60, March 2002.

[6] W. Jiang and A. Ortega, "Lifting factorization-based discrete wavelet transform architecture design," *IEEE Transactions on Circuits ans Systems for Video Technology*, vol. 11, no. 5, pp. 651–657, May 2001.

[7] Sergio D. Servetto, "Sensing lena - massively distributed compression of sensor images," *ICIP - International Conference on Image Compression*, September 2003.

[8] J. Kusuma, L. Doherty, and K. Ramchandran, "Distributed compression for sensor networks," *ICIP-International Conference on Image Processing*, October 2001.

[9] W. Sweldens, "The lifting scheme: A construction of second generation wavelets," Technical report 1995:6, Industrial Mathematics Initiative, Department of Mathematics, University of South Carolina, (ftp://ftp.math.sc.edu/pub/imi_95/imi95_6.ps), 1995.

[10] C. Lian, K. Chen, H. Chen, and L. Chen, "Lifting based discrete wavelet transform architecture for jpeg2000," *ISCAS - International Symposium on Circuits and Systems*, May 2001.