A MULTISCALE DATA REPRESENTATION FOR DISTRIBUTED SENSOR NETWORKS

Raymond Wagner, Shriram Sarvotham, Richard Baraniuk

Department of Electrical and Computer Engineering, Rice University, Houston TX, USA

ABSTRACT

Though several wavelet-based compression solutions for wireless sensor network measurements have been proposed, no such technique has yet appreciated the need to couple a wavelet transform tolerant of irregularly sampled data with the data transport protocol governing communications in the network. As power is at a premium in sensor nodes, such a technique is necessary to reduce costly communication overhead. To this end, we present an irregular wavelet transform capable of adapting to an arbitrary, multiscale network routing hierarchy. Inspired by the Haar wavelet in the regular setting, our wavelet basis forms a tight frame adapted to the structure of the network. We demonstrate results highlighting the approximation capabilities of such a transform and the clear reduction in communication cost when transmitting a compressed snapshot of the network to an outside user.

1. INTRODUCTION

Wireless sensor networks have arisen in recent years as an application area rooted in such disparate fields as micro-electro-mechanical systems, digital signal processing, and wireless networking. With a variety of prototype systems already developed and in fledgling deployment, sensor networks represent a promising new tool for scientific study, given their ability to immerse themselves in an environment and densely sample phenomena of interest.

As such nodes must operate remotely without operator intervention, they must carefully budget power drains from their on-board batteries. Consisting of sensors to capture data, microprocessors to process data, and radios to establish network links, sensor nodes allocate by far the majority of their power to wireless communication. In fact, the disparity between the power expenditure for sensing/processing and communicating is typically orders of magnitude [1], with the clear implication that processing to reduce transmission loads will provide a substantial savings in power output.

Recently, sensor network literature has begun to investigate the application of multiscale structures to such tasks as querying a network for features of interest and compressing a measurement field to facilitate both in-network storage and out-network transmission of the field. These techniques typically impose a hierarchy on the network, whereby nodes are clustered and the clusters themselves are iteratively clustered until a root in the hierarchy is reached. To address the problem of routing communications in a sensor network, both the wavRoute [2] and COMPASS [3] protocols suggest such hierarchies for data flow. Similarly, multiscale data representations have been posed to address some of the typical data-processing applications in sensor networks. DIMENSIONS [4] proposes to use an in-network wavelet transform to facilitate both network querying and measurement compression for in-network data storage. The latter goal is also pursued by the related Wisden [5] system. In a similar spirit, the Fractional Cascading method of [6] employs a multiscale representation of the network to facilitate efficient processing of queries.

To date, wavelet data representations in sensor networks have assumed that sensors lie along a regular grid in the network and/or that the cost of computing each scale of the transform is negligible. Though such assumptions allow for direct application of traditional wavelet theory in two dimensions, they are not reasonable for realistic deployments, as detailed in [7]. While regular grid wavelet transforms do give excellent compression results for locally smooth fields, they are typically implemented as finite-length filters and by no means incur negligible communication overhead. Either all the samples must be known at a central location, or each sensor must collect neighboring measurements within the filter support prior to computing its wavelet coefficient. Moreover, regular grid wavelet transforms do not apply to data sampled on irregular grids.

Clearly, any practical multiscale transform must accommodate networks with arbitrary, irregular placement of sensors. Such a transform must also incur minimal communication overhead and be well matched to the routing protocol employed in the network. To this end, we propose a new distributed irregular wavelet transform related to the well-known Haar wavelet basis.

Given an arbitrary multiscale routing hierarchy, such as that determined by COMPASS [3], we develop a transform that restricts inter-sensor communication to within clusters formed by the hierarchy – areas where local communication is assumed to be least expensive. In Section 2, we overview prior work on computing centralized irregular wavelet transforms. Section 3 introduces the structure of hierarchically routed, irregularly sampled networks and discusses the constraints imposed by the routing topology. In Section 4 we develop the transform and provide bounds on its reconstruction error when the transform coefficients are thresholded and used to reconstruct an approximation of the field. Section 5 provides an example of our transform applied to the problem of compressing the measurement field and routing a compressed description outside the network. An analysis of the communication benefit provided by such a compression follows. Finally, Section 6 concludes with a discussion of our results and directions for the evolution of our transform.

2. CHALLENGES OF IRREGULAR DATA SAMPLING

Were sensor data sampled using a regular grid, they would be ideally suited to transformations using so called "first-generation" wavelets [8]. Such grids lend themselves to a regular multiscale

This work was supported by NSF, AFOSR, ONR, and the Texas Instruments Leadership University Program. Email: {rwagner,shri,richb}@rice.edu. Web: dsp.rice.edu.

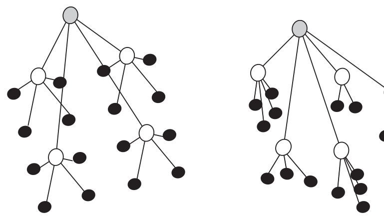

Fig. 1. Comparing the structure of a quadtree on regular-grid data (a) to that of an arbitrary hierarchy on irregular-grid data (b).

hierarchy like the quadtree illustrated in Figure 1(a), and coefficients of the wavelet transform are easily calculated across scales using integer translates and power-of-two dilations of a single mother wavelet function. When grid regularity is not the case, first-generation wavelets will no longer suffice, and we must appeal to the second generation of wavelet theory and develop irregular wavelet transforms based on the lifting scheme [8]. For example, sophisticated irregular wavelet transforms are proposed in both [9] and [10]. Unfortunately, such techniques are intended to operate in a centralized fashion on an entire dataset and are not easily distributable within a sensor network. Additionally, approaches of this type assume the freedom to impose the most convenient multiscale hierarchy on the data. In a sensor network with a specific, efficient hierarchical routing structure already in place, such a liberty can prove far too costly. Thus, any practical irregular wavelet transform for sensor networks must reconcile itself to the routing-optimal multiscale structure of the network.

3. MULTISCALE ROUTING HIERARCHIES

Multiscale routing protocols, such as wavRoute [2] and COMPASS [3], typically decompose the network of sensors into clusters, electing a cluster head in each. This set of cluster heads is then itself gathered into clusters with cluster heads, and the process repeats, forming a tree, until a single root node is reached. When nodes conform to a regular grid, such a hierarchy could take the form of a quadtree (Figure 1(a)), as suggested in [2]. A routing hierarchy of this form matches exactly that used for a regular wavelet transform, so that data flow during the transform requires the kinds of communications for which the routing hierarchy has been optimized.

Realistically, though, the network will comprise an irregularly sampled dataset, and the routing hierarchy will resemble that shown in Figure 1(b). Again, cluster formation imposes a specific communication economics whereby communications within clusters and up or down the hierarchy are considered relatively inexpensive, while communication across clusters is taken to be more costly. To implement an efficient multiscale transform, we require that data shared between sensors during the analysis phase of the transform correspond to communications favored by the hierarchy. Thus, in any given cluster, we must restrict the scope of our transform to the measurements within that cluster. This imposition, along with the need to accommodate a specific hierarchy, prevents us from directly applying the irregular wavelet theory discussed in Section 2. Instead, we look for inspiration to the most rudimentary wavelet basis from regular wavelet theory: the Haar wavelet.

4. TRANSFORM DETAILS

Conceptually, wavelets reduce a set of data into some sort of average and differences from that average, with the goal that these differences should be small and therefore both easier to encode and potentially more negligible than points in the original data set. This can be seen most directly by considering the Haar wavelet basis in the 1-D regular setting, as shown in Figure 2(a). Thinking in terms of averages and differences provides an intuitive starting point for developing a transform with the restrictions discussed in Section 3 – namely, that within a cluster, the transform has access only to cluster points, whose count is arbitrary and determined by the routing hierarchy. Although Haar wavelets cannot be applied directly to irregular samples with non-dyadic hierarchies, they do provide inspiration for a new irregular wavelet transform for sensor networks.

4.1. Extending Haar Past A Pair of Measurements

Figure 2(a) shows the standard Haar basis functions on a regular 1-D grid. Function s captures the sample average while w encodes a combination of differences of the two samples from the average. To adapt this methodology to our sensor network application, we must make two modifications. First, to account for the irregularity of the sampling, we must assign support sizes to each sample, essentially assuming that the sampled signal is piecewise constant over the supports of the samples. This is illustrated at the top of Figure 2(b). The discrete inner product sums of the regular Haar analysis become continuous inner product integrals of piecewise constant measurement and basis functions. In fact, regular Haar can be expressed the same way, with the uniform support size from the continuous integral factoring out at the end of the analysis to effectively give the discrete inner product. To see the need for this modification, we must realize that no longer are all samples equal – the average value must not be biased toward samples in regions of high sampling density.

Thus, we can form a set of piecewise-constant basis functions in the 1-D example, as given in Figure 2(b). The function s, everywhere constant, serves as the scaling analysis function. The remaining functions w_1, w_2, and w_3, each constant save for a single support area, act as wavelet functions. Magnitudes of the piecewise-constant regions in each function are given in Section 4.2, where we state the explicit form of analysis and synthesis matrices. Assigning support sizes in these 1-D examples is trivial — for any given sensor, its support extends halfway to either of its neighboring sensors. This effectively makes each sensor responsible for the space closer to it than to any other sensor according to the Euclidian norm. Extending this idea to 2-D, we can assign to each sensor the polygon surrounding it in a Voronoi tessellation [11] of the sensor field. Again, the polygon enclosing a sensor contains all points closer to that sensor than any other. Voronoi areas can be computed for the sensor network during a startup phase, and as the polygon for a sensor depends only on locations of sensors whose own polygons abut it, the Voronoi calculation should be distributable.

Finally, we note that Haar wavelets form a non-redundant, complete orthonormal basis. As a result, the sum of the squared wavelet coefficients for any signal expansion equals the energy of the original signal. Maintaining such a Parseval relation is key for any useful sensor network wavelet transform, as it guarantees optimal approximation results via thresholding, to be discussed in

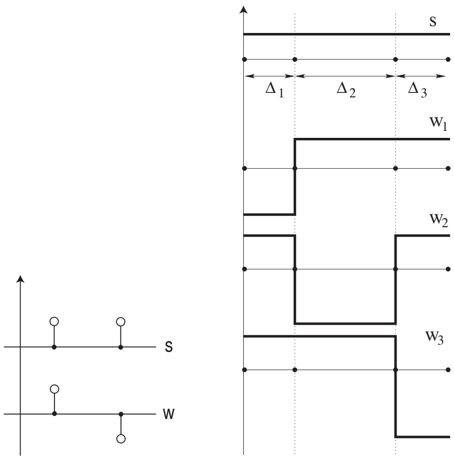

(a) (b)

Fig. 2. Relating the discrete 1-D regular Haar basis $\{s, w\}$ (a) to the continuous irregular tight frame basis $\{s, w_1, w_2, w_3\}$ (b), which is piecewise constant over the intervals $\{\Delta_1, \Delta_2, \Delta_3\}$.

Section 4.3. And while our basis does obey Parseval, we introduce a bit of redundancy, forming a total of $N+1$ coefficients for N measurements in a cluster – that is, we produce 1 average value and N differences for a set of N sensors. This tiny amount of redundancy, tolerated to yield a clean expansion set, means that we have an overcomplete basis — in this case, a Parseval tight frame [12]. We now turn to the details of computing this tight frame expansion.

4.2. Tight Frame Expansion

For a cluster of N sensors, we wish to take the vector of sensor measurements $\underline{m} = [m_1, m_2, ..., m_N]^T$ and transform it into the coefficient vector $\underline{c} = [s, w_1, w_2, ..., w_N]^T$, where s is the scaling coefficient for the cluster and $\{w_j\}_{j=1}^N$ are the wavelet coefficients. We begin by defining the $(N+1)$-by-N matrix $\mathbf{K}$:

$$\mathbf{K} = \begin{pmatrix} k_0 & k_0 & k_0 & \cdots & k_0 \\ k_1^{'} & k_1 & k_1 & \cdots & k_1 \\ k_2 & k_2^{'} & k_2 & \cdots & k_2 \\ k_3 & k_3 & k_3^{'} & \cdots & k_3 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ k_N & k_N & k_N & \cdots & k_N^{'} \end{pmatrix}.$$

The first row of $\mathbf{K}$ represents the constant-valued scaling function, while each $(i,j)^{th}$ element in the remaining rows gives the value of wavelet function $i-1$ over the Voronoi cell surrounding sensor j (see Figure 2(b)). Utilizing the set of N Voronoi tessellation areas $\{\Delta_j\}_{j=1}^N$ assigned to the sensors as described above, the majority of row coefficients in $\mathbf{K}$ are given by

$$k_i = \begin{cases} \frac{1}{\sqrt{\Delta_{\mathrm{tot}}}}, & i = 0 \\ \frac{\sqrt{\Delta_i}}{\Delta_{\mathrm{tot}}}, & i > 0 \end{cases},$$

with $\Delta_{\mathrm{tot}} = \sum_{j=1}^N \Delta_j$. Entries on the first sub-diagonal are described as

$$k_i^{'} = -\frac{k_i(\Delta_{\mathrm{tot}} - \Delta_i)}{\Delta_i}, \quad i > 0.$$

The analysis matrix is formed as $\mathbf{T_A} = \mathbf{K}\boldsymbol{\Delta}$, where $\boldsymbol{\Delta}$ is an N-by-N diagonal matrix of the areas $\{\Delta_j\}_{j=1}^N$. Thus, wavelet and scaling coefficients can be calculated as $\underline{c} = \mathbf{T_A}\underline{m}$. The synthesis matrix takes the form $\mathbf{T_S} = \boldsymbol{\Delta}^{-1}\mathbf{T_A}^T$, so that measurements can be recovered as $\underline{m} = \mathbf{T_S}\underline{c}$. For a detailed proof showing that this transform does indeed form a Parseval tight frame, we refer the reader to [13].

4.3. Approximation via Thresholding

Approximation of wavelet transforms typically amounts to discarding the wavelet coefficients which are the "least important," a distinction the tight frame allows us to quantify. When the wavelet-domain signal formed in Section 4.2 is approximated by zeroing out a set of M wavelet coefficients $\{w_{z,i}\}_{i=1}^M$, the energy difference $\|S\|^2 - \|\widehat{S}\|^2$ between the reconstructed and original signals is bounded as

$$\|S\|^2 - \|\widehat{S}\|^2 \leq 2\sum_{i=1}^M w_{z,i}^2 \quad ,$$

where S describes the original field of measurements, and $\widehat{S}$ gives the field reconstructed following approximation. This result ensures that thresholding may be used as an optimal approximation technique: given that we are able to keep a only a certain number of coefficients, those best discarded are the M smallest, since the bound depends on the magnitude of zeroed coefficients. Again, for a proof of the energy difference bound, the reader should consult [13].

4.4. Iteration

To build up the complete multiscale representation, the transform outlined in Section 4.2 is first applied at each cluster of sensors. Sensors in a cluster transmit their measurements and areas to a clusterhead, which generates wavelet and scaling coefficients. Each clusterhead then retains wavelet coefficient values and passes its scaling coefficient s and aggregate area Δ_{tot} to the clusterhead for the next higher level in the hierarchy. Coefficients are again calculated, and the process iterates on the scaling coefficients and summed areas for the set of higher level clusters, terminating when the root of the hierarchy is reached, at which point the single remaining scaling coefficient is recorded, and the transform is complete.

For the compression application, a user outside the network broadcasts to the network a threshold below which wavelet coefficients may be discarded. Clusterheads at all levels then compare their stored wavelet coefficients against the threshold and begin to send up the hierarchy those values which are significant. Values are tagged as to their position in the transform, and the user may reconstruct an approximation of the field, assuming knowledge of sensor locations, which may be transmitted during network initialization. In the next section, we present an example of such network compression.

5. EXAMPLE

To evaluate the effectiveness of our proposed transform, we randomly assign coordinates on the $(0,1) \times (0,1)$ square to 2500

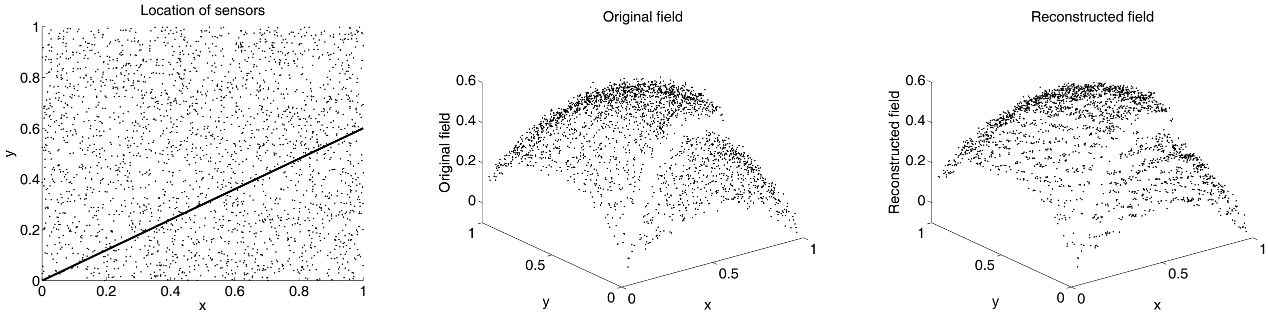

(a) (b) (c)

Fig. 3. Irregular sensor locations (a), samples from a piecewise smooth quadratic signal with additive noise (b), and reconstruction using largest 15% of the wavelet coefficients (c).

sensors. This field is depicted in Figure 3(a). Sensors are assigned to a 6-level hierarchy, with 2500, 500, 100, 20, and 4 clusters in each of the fine-to-coarse levels and 1 root cluster. The hierarchy is allowed to form randomly, clustering spatially proximate sensors. Sensor locations are then used to samples a smooth, noisy quadratic with a discontinuity along the line $y = 0.6x$, yielding the field of sensor measurements shown in Figure 3(b).

To evaluate approximation accuracy, we set approximately 85% of the wavelet coefficients to zero via thresholding of the smallest coefficients. The reconstructed field, depicted in Figure 3(c), closely resembles the original and has 99.74% of the original signal energy. The energy difference between the original and approximated signals is only 1.08 times that of the discarded wavelet coefficients – well within the bounds calculated in Section 4.3.

To understand the communication benefits of such a compression, we compare the cost to route out all the original measurements to that of routing out significant wavelet coefficients. Assuming an exponential increase in the cost of links as we traverse up the hierarchy, the cost of transmitting only the non-zero wavelet coefficients is 14.9% of the cost to transmit all measurements – an excellent reduction in transmission cost for such a small approximation error.

6. CONCLUSIONS AND FUTURE WORK

To summarize, we have proposed what is, to our knowledge, the first irregular wavelet transform well-suited to distributed operation within a sensor network. Rather than requiring any specific multiscale hierarchy, our transform adapts to the hierarchy best suited to data transport within the network, tightly coupling transform with routing and ensuring that transform overhead requires minimal communication energy. We have presented experimental results indicating that the approximation power of our transform is substantial and that communication cost is substantially reduced by employing such a technique when transmitting a snapshot of the measured field off-network. We intend to extend the results seen here to accommodate higher orders of approximation than piecewise constant.

7. REFERENCES

[1] G. J. Pottie and W. J. Kaiser, "Wireless integrated network sensors," *Comm. of the Assoc. for Computing Machinery*, vol. 43, no. 5, pp. 51–58, May 2000.

[2] D. Ganesan, D. Estrin, and J. Heidemann, "DIMENSIONS: Why do we need a new data handling architecture for sensor networks?," *ACM SIGCOMM Comput. Commun. Rev.*, vol. 33, no. 1, pp. 143–148, 2003.

[3] "COMPASS project website," `http://compass.cs.rice.edu`, Rice Univeristy (Houston, TX).

[4] D. Ganesan, B. Greenstein, D. Perelyubskiy, D. Estrin, and J. Heidemann, "An evaluation of multi-resolution storage for sensor networks," in *Proc. of the ACM SenSys Conf.*, Los Angeles, California, USA, November 2003, pp. 89–102.

[5] N. Xu, S. Rangwala, K. Chintalapudi, D. Ganesan, A. Broad, R. Govindan, and D. Estrin, "A wireless sensor network for structural monitoring," in *Proc. of the ACM Conf. on Embedded Networked Sensor Sys.*, November 2004.

[6] J. Gao, L. J. Guibas, J. Hershberger, and L. Zhang, "Fractionally cascaded information in a sensor network," in *Proc. of the Third Int. Symp. on Information Proc. in Sensor Networks (IPSN)*, Berkeley, California, USA, 2004, pp. 311–319.

[7] D. Ganesan, S. Ratnasamy, H. Wang, and D. Estrin, "Coping with irregular spatio-temporal sampling in sensor networks," *SIGCOMM Comput. Commun. Rev.*, vol. 34, no. 1, pp. 125–130, 2004.

[8] W. Sweldens, "The lifting scheme: A construction of second generation wavelets," *SIAM J. on Math. Analy.*, vol. 29, no. 2, pp. 511–546, March 1998.

[9] I. Daubechies, I. Guskov, P. Schröder, and W. Sweldens, "Wavelets on irregular point sets," *Phil. Trans. R. Soc. Lond. A*, vol. 357, no. 1760, pp. 2397–2413, 1999.

[10] V. Delouille, M. Jansen, and R. von Sachs, "Second generation wavelet methods for denoising of irregularly spaced data in two dimensions," Tech. Rep., Institut de Statistique, Université Catholique de Louvain, Belgium, January 2003.

[11] A. Okabe, B. Boots, K. Sugihara, and S.N. Chiu, *Spatial Tessellations: Concepts and Applications of Voronoi Diagrams*, Wiley, second edition, 2000.

[12] P. G. Casazza and J. Kovacevic, "Equal-norm tight frames with erasures," *Advances in Computational Mathematics*, vol. 18, no. 2-4, pp. 387–430, February 2003.

[13] S. Sarvotham, R. Wagner, and R. Baraniuk, "A multiscale data representation for distributed sensor networks: Proofs of basis characteristics and error bounds," Tech. Rep., Rice University (Houston, TX), `http://cmc.rice.edu/docs/docinfo.aspx?doc=Sar2004Sep9AMultiscal`, September 2004.