

Energy-Scalable Protocols for Battery-Operated MicroSensor Networks

**Alice Wang, Wendi Rabiner Heinzelman, and
Anantha P. Chandrakasan**

Department of Electrical Engineering
Massachusetts Institute of Technology
Cambridge, MA

Abstract - To maximize battery lifetimes of distributed wireless sensors, network protocols and data fusion algorithms should be designed with low power techniques. Network protocols minimize energy by using localized communication and control and by exploiting computation/communication trade-offs. In addition, data fusion algorithms such as beamforming aggregate data from multiple sources to reduce data redundancy and enhance signal-to-noise ratios, thus further reducing the required communications. We have developed a sensor network system that uses a localized clustering protocol and beamforming data fusion to enable energy-efficient collaboration. We have implemented two beamforming algorithms, the Maximum Power and the Least Mean Squares (LMS) beamforming algorithms, on the StrongARM (SA-1100) processor. Results from our experiments show that the LMS algorithm requires less than one-fifth the energy required by the Maximum Power beamforming algorithm with only a 3 dB loss in performance. The energy requirements of the LMS algorithm was further reduced through the use of variable-length filters, a variable voltage supply, and variable adaptation time.

1. INTRODUCTION

Networks of microsensors can greatly improve environment monitoring for many civil and military applications [1]. For example, a wireless sensor system can be used for boundary surveillance, for target detection and classification, or for patient monitoring. Multiple sensors provide fault tolerance and can provide valuable inferences about the physical world to the end-user.

In order to prolong the lifetimes of the wireless sensors, all aspects of the sensor system should be energy-efficient. This includes the sensor, data conversion, signal processors, network protocols, and RF communication. Energy scalability further allows the sensor network and individual sensors to adapt as energy resources of the system diminish. This allows for longer battery lifetimes and more efficient sensor systems.

Prepared through collaborative participation in the Advanced Sensors Consortium sponsored by the U.S. Army Research Laboratory under Cooperative Agreement DAAL01-96-2-001.

A network protocol layer allows for sensor collaboration. If the distance between neighboring sensors is less than the distance between the sensors and the end-user, then transmission power can be saved if the sensors collaborate locally. We have developed a clustering communication protocol whereby sensors communicate with a local control center (called a “cluster-head”). Since it is likely that the sensors in the local cluster share highly correlated data, the cluster-head aggregates the data and then transmits the aggregate data to the end-user. In addition to reducing transmission power, effective data aggregation can improve signal enhancement, detection and classification.

Beamforming is one method of combining data from multiple sensors in order to satisfy a given performance criteria. The advantage of beamforming is that the desired signal is enhanced while the uncorrelated noise is reduced, which in turn improves detection and classification of the source. An extension of beamforming also allows for source localization and tracking [2]. However, beamforming algorithms are computationally complex, often involving matrix operations, and this large amount of computation results in large power dissipation. Thus, there are tradeoffs between performance and power dissipation which should be considered when implementing beamforming algorithms for sensor networks.

2. LOW POWER NETWORK PROTOCOLS

Often, sensor networks are used to monitor remote areas or disaster situations. In both these scenarios, the end-user cannot be located near the sensors. Thus, direct communication between the sensors and the end-user, as shown in Figure 1a, is extremely energy-intensive, since transmission energy goes as R^n (n typically 2-4). In addition, direct communication may not be feasible for large-scale sensor networks. If, for example, frequency-division is used (e.g., each sensor obtains a certain bandwidth in which to transmit data), the amount of information that can be sent from each sensor to the end-user becomes negligible as the number of sensors increases, because each sensor’s bandwidth is reduced down to zero. Thus new methods of communication need to be developed.

A common method of communication in wireless networks is multi-hop routing, where sensors act as routers for other sensors’ data in addition to sensing the environment, as shown in Figure 1b. Multi-hop routing minimizes the distance an




Figure 1. (a) Direct communication with basestation. (b) Multi-hop communication with basestation. (c) Clustering algorithm. The grey nodes represent “cluster-heads”, and the function $f(A,B,C)$ represents the data fusion algorithm.




Figure 2. Data aggregation done locally can reduce energy dissipation.

individual sensor must transmit its data, and hence minimizes the dissipated energy for that sensor. However, multi-hop requires that several sensors transmit and receive a particular signal; hence this does not achieve global energy-efficiency. For example, the sensors near the end-user will be used as routers for a large number of the other sensors, and their lifetimes will be dramatically reduced using such a multi-hop protocol.

Since data from neighboring sensors will often be highly correlated, it is possible to aggregate the data locally using an algorithm such as beamforming and then send the aggregate signal to the end-user to save energy. Figure 2 shows the amount of energy required to aggregate data from 2, 3, and 4 sensors and to transmit the result to the end-user, as compared to all of the individual sensors transmitting data to the end-user. As shown in this plot, there is a large advantage to using local data aggregation (beamforming), rather than direct communication. In this scenario, we assume that the transmission energy dissipated is 10pJ/bit/m^4 and the reception energy dissipated is 10pJ/bit .

We have developed a clustering algorithm that utilizes the energy savings from data aggregation to greatly reduce the energy dissipation in a sensor system. In our algorithm, the sensors self-organize into local clusters, as shown in Figure 1c. Each cluster has a “cluster-head”, a sensor that receives data from all other sensors in the cluster, performs data fusion (e.g., beamforming), and transmits the aggregate data to the end-user. This greatly reduces the amount of data that is sent to the end-user and thus achieves a global energy minimization. Furthermore, the clusters can be organized hierarchically such that the cluster-heads transmit the aggregate data to “super-cluster-head” nodes, rather than directly to the end-user so as to further reduce energy dissipation.

Figure 3a shows the total energy dissipated in the sensor network as the diameter of the network is increased using a direct transmission protocol, a multi-hop routing protocol, and our clustering algorithm. This plot shows that our clustering algorithm achieves greater than a factor of 6 reduction in energy compared with a direct communication approach and a multi-hop routing protocol. In addition to reducing energy dissipation, Figure 3b shows that our clustering algorithm is able to double the system lifetime compared with the other protocols.

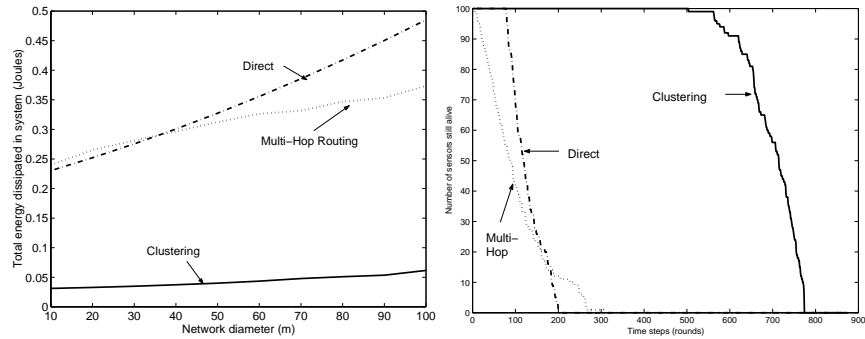


Figure 3. A comparison of (a) total energy dissipated as the diameter of the sensor network is increased and (b) system lifetime for a direct communication protocol, a multi-hop routing protocol, and our clustering algorithm.

In addition to minimizing energy dissipation, our clustering algorithm has several other advantages over tradition routing protocols. The clusters are self-organizing and use localized coordination and control, which not only enables scalability of the network (as no reorganization of the network is required when nodes are added to the system), it also enhances the fault tolerance of the system. This protocol can easily handle trade-offs in computation and communication. If computation is expensive compared to communication costs, the network can have the cluster-head transmit all data directly to the basestation. On the other hand, if computation is cheap compared to communication costs, the cluster-head can perform signal processing functions to compress the data from all the sensors in the cluster and transmit the compressed (aggregated) data to the end-user. For example, any of the beamforming algorithms discussed in the next section can be used by the cluster-head to aggregate the data from the sensors in the cluster.

3. BEAMFORMING ALGORITHMS

Beamforming algorithms combine signals from multiple sensors in order to satisfy some optimization criteria. Example criteria are minimizing mean squared error (MSE), maximizing signal-to-noise ratio (SNR), and minimizing variance. Figure 4 shows a block diagram that describes how beamforming algorithms can be applied to a wireless network of sensors. Assume there are M acoustic sensors




Figure 4. Beamforming at the cluster-head.

which have detected a target, and each sensor transmits its data using the wireless channel to the cluster-head, as described in Section 2. At the cluster-head, the beamforming algorithm chooses the L -tap FIR filters, $w_i(n)$, to optimize a selected criteria. Each $w_i(n)$ is applied to the i th sensor data, $s_i(n)$, and the resulting signals are summed for all M sensors, to get the beamformed signal, $y(n)$:

$$y(n) = \sum_{i=1}^M \sum_{l=0}^{L-1} w_i(l) s_i(n-l) \quad (1)$$

We have benchmarked the energy requirements and performance of two beamforming algorithms that are suitable for the application of distributed sensor nodes.

3.1 Maximum Power Beamforming Algorithm

In [3], Yao et al. propose an eigenvector-based method to perform Maximum Power beamforming for a randomly spaced sensor network. The algorithm uses the correlation matrix of the sensor data to find the weighting filters that pick out the signal with the highest peak power spectral density. The weighting filters are chosen to solve the following maximization problem,

$$\text{maximize } \left\{ \mathbf{w}_{ML}^T \mathbf{R}_{ML} \mathbf{w}_{ML} \right\}, \text{ subject to } \|\mathbf{w}_{ML}\| = 1. \quad (2)$$

where \mathbf{R}_{ML} is the space-time correlation matrix of the sensor data, given by

$$\mathbf{R}_{ML} = E \left[s_M(n) s_M(n)^T \right] \quad (3)$$

and $s_M(n)$ is the sensor data.

The desired weighting vector \mathbf{w}_{ML} is given by the eigenvector corresponding to the largest eigenvalue of \mathbf{R}_{ML} . A detailed proof of this solution can be found in [3].

The bulk of the computation involved in Maximum Power beamforming is involved in the following steps: (1) computing \mathbf{R}_{ML} , the correlation matrix, from the sensor data and (2) performing the eigenvector decomposition of \mathbf{R}_{ML} . The power method of eigenvector decomposition provides a low-computation, iterative method to find the eigenvector with the largest eigenvalue [3].

3.2 Least Mean Squares (LMS) Algorithm

Another algorithm used in antenna array processing is the time-domain LMS adaptive algorithm [4]. The LMS algorithm uses a minimum mean squared error criterion to determine the appropriate array weighting filters. This algorithm is considered an optimum algorithm because the solution minimizes the error between the array output and the desired signal. Therefore, it is assumed that the desired signal is known, or a signal containing the desired signal characteristics is available.

The LMS iterative equations are :

$$\mathbf{w}_{ML}(n+1) = \mathbf{w}_{ML}(n) + 2\mu s_M(n) \boldsymbol{\varepsilon}(n) \quad (4)$$




Figure 5. Energy requirements for Beamforming algorithms on SA-1100.

where μ is the stepsize, which governs the rate of convergence of this iterative process, and $\epsilon(n)$ is the error function between the output and the desired signal.

3.3 Energy Requirements

We ran the algorithms on the StrongArm-1100 (SA-1100) processor. Figure 5 shows the energy dissipated (in $\mu\text{J}/\text{sample}$) for 32 tap sensors as the number of sensors, M , is varied from 2 to 10. This figure shows that the LMS algorithm requires one-tenth the energy of the Maximum Power beamforming algorithm. In addition, the energy requirement for the LMS algorithm is linear with the number of sensors, while the Maximum Power algorithm has a quadratic dependence on the number of sensors.

3.4 Energy - Quality Tradeoff

Through the use of beamforming, the source signature is enhanced, leading to improved detection and classification. We benchmarked the performance of the two beamforming algorithms using the mean squared error (MSE) quality measure and using acoustic data collected of tracked vehicles. Figure 6 shows the performance of the two algorithms for different SNR's, as we increase the number of sensors involved in beamforming from no beamforming and 2 sensor beamforming up to 5 sensor beamforming. We assume that there is also communication energy dissi-

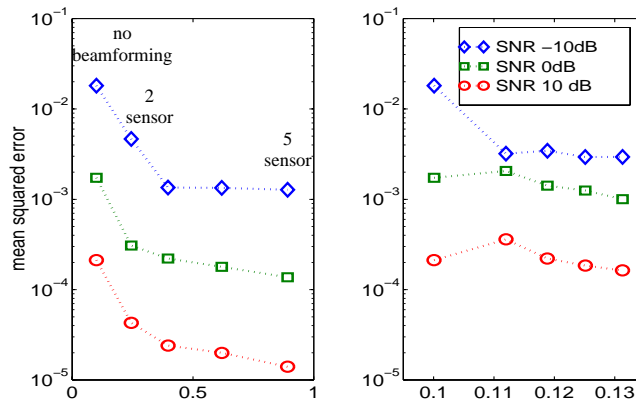


Figure 6. Energy vs. MSE for Maximum Power and LMS algorithm.

pated, and the assumptions made are similar to those found in Section 2.

Figure 6 demonstrates two key findings. First, there is a large improvement in MSE between no-beamforming and 2 sensor beamforming, which shows that beamforming algorithms can be used enhance the signal and separate the desired signal from the uncorrelated noise. Second, the Maximum Power beamforming algorithm achieves better than 3 dB improvement in performance compared to the LMS algorithm, but at the cost of dissipating 5 times more energy.

4. LOW POWER TECHNIQUES FOR LMS BEAMFORMING

Energy scalability can be achieved by monitoring energy resources, latency and performance requirements to dynamically reconfigure an algorithm. Due to its low complexity, the LMS algorithm is better suited for low power applications. In addition, the LMS algorithm is flexible because it dynamically changes the value of the filter coefficients to adapt to a changing environment. We have developed a variable-length filter architecture that can dynamically adjust the filter order and a variable adaptation time approach to power down the update computation if the error is below a certain threshold.

An energy-aware controller monitors the LMS error function, $\epsilon(n)$, as well as the sensor's energy resources/performance requirements and dynamically changes the datapath parameters involved in the LMS algorithm. The LMS beamforming algorithm block can be implemented using a tapped delay line approach as shown in Figure 7. This approach of approximate signal processing architectures has been found in a variety of related work in adaptive filtering [5], adaptive equalizers in VDSL [6] and broadband modems [7]. In our implementation, the LMS algorithm is implemented in software on a low power embedded processor.

4.1 Variable-Length Filtering

The length of the adaptive filter can affect the performance and energy requirements of the LMS algorithm. Increasing the length of the adaptive filter improves the frequency resolution of the signal processing done, thus reducing MSE and improving performance. However, this comes at the cost of an increase in energy dissipation. In a software implementation, the number of cycles increases linearly

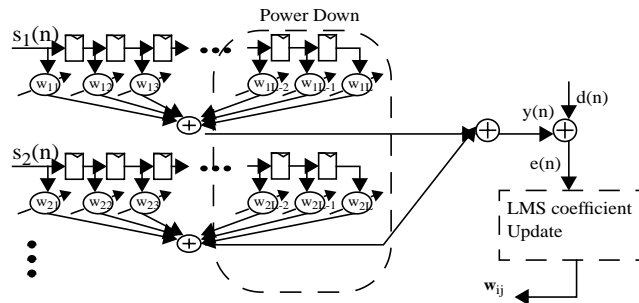


Figure 7. Tapped delay line structure of the LMS beamforming algorithm.




Figure 8. MSE versus filter length for different number of sensors.

as the filter length is increased. Thus given a specified performance requirement, the latter parts of the tapped delay line can be disabled to reduce the number of processor cycles. This, in turn, reduces the energy dissipated.

Figure 8 shows the relationship between filter length and MSE. This plot shows that there is an optimal filter length which minimizes the MSE and the filter length, providing both low power and the required performance. The optimal filter length is highly data dependent, but in general, a filter that is too short may not provide enough frequency resolution, but a filter that is too long takes longer to converge to the optimal solution.

A simple variable-length filter controller computes the MSE :

$$MSE = \frac{1}{L} \sum_{n=1}^L \varepsilon^2(n) \quad (5)$$

where the error function, $\varepsilon(n)$, is given in Section 3.2. A programmable threshold, α , is set and the filter length is set initially to the maximum length, L_{max} . On a frame to frame basis, the filter length is decreased until the MSE is greater than α .

4.2 Variable Voltage Supply

Since latency is linearly related to filter length and the number of sensors, we can use a variable voltage supply and variable clock rate to further reduce energy dissipation [8]. The total energy dissipated by a digital circuit is given by

$$E_{TOT} = C_{TOT} V_{DD}^2 + V_{DD} I_{leak} \Delta t, \quad (6)$$

where E_{TOT} , the total energy dissipated, is the sum of the energy lost to switched capacitance (C_{TOT}) and the energy lost to sub-threshold current leakage (I_{leak}). V_{DD} is the voltage supply and Δt is the latency.

Assume that the throughput is fixed for the worst case scenario, where $L = L_{max}$, the worst case filter length, and for $M = M_{max}$, the maximum number of sensors. If we have a variable-length filter architecture or receive data from fewer than M_{max} sensors, then there is less computation required than in the worst case scenario, and it will be completed in fewer processor cycles. Ideally, if we reduce the clock rate, then the energy dissipated should be the same, but due to leakage and increased latency, the energy dissipated will increase. Thus, when we reduce the clock rate, it




Figure 9. Latency vs. energy for a variable voltage supply on the StrongARM SA-1100.

is also necessary to reduce the voltage supply level.

We have modified the StrongARM processor to run at different voltage levels. Figure 9 shows the energy versus latency for both a fixed voltage supply and a variable voltage supply, for the LMS algorithm run on the StrongARM processor. For the fixed voltage supply, the StrongARM processor frequency was reduced while the voltage supply was held at a constant 1.42 V level. The increase in energy reflects the leakage currents and latency effects on energy dissipation. For a variable voltage supply, we can see an inverse squared relationship between the latency and the energy.

4.3 Variable Adaptation Time

Another way to save energy is to power down the LMS coefficient update computation (see Figure 7). The iterative equations in the LMS algorithm adaptively approach the optimum weighting filters by using the steepest gradient descent. As the weighting filters approach the optimum solution, the error function, $\epsilon(n)$, approaches zero. Thus, $\epsilon(n)$ can be monitored and a programmable threshold, β , can be set, such that when the error falls below β , the LMS coefficients update computation is powered down. If the error rises above β , the LMS coefficients update computation is restarted.

Figure 10 shows the tradeoff between performance and energy for the constant adaptation versus variable adaptation for $\beta = 10^{-3}$, and 5×10^{-4} . This plot shows




Figure 10. Energy vs. performance for constant adaptation and variable adaptation.

that as we increase the threshold, the performance worsens and there is less power dissipated. Thus, if the performance requirement can be relaxed, then a variable adaptation architecture can help to reduce power dissipation.

5. CONCLUSIONS

Simulations on the StrongARM SA-1100 processor have shown that the LMS beamforming algorithm is a suitable data aggregation algorithm for the application of multiple distributed acoustic sensors. The LMS algorithm provides signal enhancement and has low complexity, when compared to the Maximum Power beamforming algorithm. We have also shown how the flexibility of the LMS algorithm can be exploited to further reduce the power dissipated. Simple controllers have been suggested to implement variable-length filters, variable voltage and variable adaptation time architectures.

ACKNOWLEDGMENTS

We would like to thank Professor Kung Yao at UCLA for valuable help and discussions. We also like to thank the Acoustic Signal Processing Branch at ARL for providing us with acoustic data from the ACIDS trials. A. Wang is supported by the Lucent Fellowship and W. Heinzelman is supported by the Kodak Fellowship.

References

- [1] K. Bult, et. al., "Low Power Systems for Wireless Microsensors," *IEEE/ACM International Symposium on Low Power Electronics and Design*, pp. 17-21, Aug 1996.
- [2] K. Yao et. al., "Array Signal Processing for a Wireless MEM Sensor Network," *1998 IEEE Workshop on Signal Proc. Systems*, invited paper, pp. 11-20.
- [3] K. Yao et. al., "Blind Beamforming on a Randomly Distributed Sensor Array System", *IEEE Journal on selected Topics in Communications*, October 1998.
- [4] S. Haykin, J. Litva, T.J. Shepherd, *Radar Array Processing*, Springer-Verlag, 1993
- [5] J.T. Ludwig, S.H.Nawab, and A.P.Chandrakasan, "Low-Power Digital Filtering Using Approximate Processing", *IEEE J. Solid-State Circuits*, vol. 31, no.3, pp.395-399, Mar 1996.
- [6] M. Goel and N.R. Shanbhag, "Low-power Equalizers for 51.84 Mb/s Very High-Speed Digital Subscriber Loop (VDSL) Modems", *1998 IEEE Workshop on Signal Proc. Systems*, pp. 317-326.
- [7] C.J.Nicol, P.Larsson, K.Azadet, J.H.O'Neill, "A Low Power 128-Tap Digital Adaptive Equalizer for Broadband Modems", *1997 IEEE International Solid-State Circuits*, pp. 94-95.
- [8] V. Gutnik, A.P. Chandrakasan, "Embedded power supply for low-power DSP," *IEEE Trans. on VLSI Systems*, vol. 12, pp. 425-435, Dec. 1997.

The views and conclusions contained in this document are those of the authors and should not be interpreted as presenting the official policies, either expressed or implied, of the Army Research Laboratory or the US Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation thereon.