DECENTRALIZED SOURCE LOCALIZATION AND TRACKING

Michael G. Rabbat & Robert D. Nowak

Electrical and Computer Engineering
University of Wisconsin
Madison, WI 53706

ABSTRACT

This paper describes a new approach to the source localization and tracking problem in wireless sensor networks. A fast, easy-to-implement algorithm for localizing a source using received signal strength measurements is presented. The algorithm is based on incremental subgradient optimization methods. Using theory on the convergence rates of these methods we characterize the amount of in-network communication required to achieve an accurate estimate of the source's location. In comparison to other localization and tracking algorithms described in the literature, the amount of communication (and thus energy and bandwidth) used by our algorithm is much lower than that used by other schemes, especially as network size grows.

1. INTRODUCTION

The problem of localizing and tracking a target has been widely sighted as a canonical application of wireless sensor networks, as sensor networking provides an attractive approach to spatial monitoring. Wireless technology makes these systems relatively easy to deploy, but also places heavy demands on energy consumption for communication. In this paper we present a new approach to source localization and tracking using received signal strength measurements. Based on incremental gradient descent-like optimization methods, our algorithm only requires that small amounts of data be communicated over short distances.

In order to solve the tracking/localization problem, individual sensors must first detect the presence of a source from their local data. Then collaborative signal processing at the network level involves routing information through the network and fusing the data from different sensors to generate an estimate of the source location or bearings. Previously proposed schemes for collaborative data processing and communication have been either hierarchical or based on collecting and fusing the data from all sensors at a central location. The work presented in this paper adopts a decentralized approach to collaborative signal processing.

We motivate taking a decentralized approach for the following two main reasons. First, in-network processing of data is more efficient in terms of energy and bandwidth usage, especially as the number of sensors grows. Also, distributed algorithms are arguably more robust than purely centralized or hierarchical processing schemes since all nodes in the network bear equal importance in a distributed approach. Thus, in a decentralized scheme the network has a higher tolerance to individual node failures.

Supported by the National Science Foundation, grants CCR-0310889 and ANI-0099148, DOE SciDAC, and the Office of Naval Research, grant N00014-00-1-0390

In, [1], Sheng and Hu consider the problem of acoustic source localization in sensor networks. They formulate a Maximum Likelihood approach which requires sensors to transmit all of their data to a fusion center for processing. In large networks, the massive amount of energy and bandwidth required make such an approach impractical. See [2] for many other articles related to source localization and tracking. An information-driven approach to routing and collaboration is described by Zhao et al. in [3] which trades off the utility of publishing data at individual sensors with the cost of communicating it. They describe a particle-filtering method which requires that many particle weights (perhaps hundreds) be communicated through the network to track the source location. While this approach is viable from the perspective that it is distributed and thus robust, it is unclear whether the amount of energy and bandwidth expended in transmitting numerous particle weights through the network is practical for large networks.

The driving philosophy behind our approach is to balance estimator accuracy with the amount of communication required. This is accomplished via in-network processing of data. An estimate of the source position is circulated through the network. Each node makes a small adjustment to the estimate based on its local data, and then passes the modified estimate to its neighbors. Similar to Sheng and Hu[1], assume that each sensor knows its coordinates (either by GPS or some other mechanism), and that the received signal strength measurements behave according to a far-field model. That is, the j-th measurement at sensor i takes the form $y_{i,j} = s_i + w_{i,j}$, where $w_{i,j}$ are i.i.d samples of a white Gaussian noise process, and

$$s_i = \frac{A}{||\theta - r_i||^\beta}, \tag{1}$$

where A is the amplitude (energy) of the signal emitted by the source, θ is the source location, r_i is the location of the i-th sensor, and β is related to the attenuation characteristics of the medium through which the signal is begin transmitted. Under these assumptions the maximum likelihood estimate for the location of a stationary source is

$$\hat{\theta} = \arg\min_\theta \frac{1}{nm} \sum_{i=1}^{n} \sum_{j=1}^{m} \left(y_{i,j} - \frac{A}{||\theta - r_i||^\beta} \right)^2, \tag{2}$$

where n is the number of sensors and m is the number of measurements at each sensor. We use a distributed incremental gradient algorithm to solve this non-linear least squares optimization iteratively.

2. DECENTRALIZED INCREMENTAL OPTIMIZATION

In our algorithm a parameter estimate is cycled through the network. When each sensor receives the current estimate it makes a small adjustment based on its local data and then passes the updated estimate on to its neighbors. Without loss of generality, assume that sensors have been numbered $i = 1, 2, \ldots, n$ according to their order in the cycle. On the k-th cycle, sensor i receives an estimate $\psi_{i-1}^{(k)}$ from its neighbor and computes an update according to

$$\psi_i^{(k)} = \psi_{i-1}^{(k)} - \alpha \nabla f_i(\psi_{i-1}^{(k)}), \tag{3}$$

where $\alpha > 0$ is the step size, $\nabla f_i(\psi_{i-1}^{(k)})$ denotes the gradient of f_i evaluated at $\psi_{i-1}^{(k)}$, and

$$f_i(\psi) = \sum_{j=1}^{m} \left(y_{i,j} - \frac{A}{||\psi - r_i||^{\beta}} \right)^2. \tag{4}$$

The algorithm begins with arbitrary initial condition $\psi_0^{(0)} = \hat{\theta}^{(0)}$, and after a complete cycle through the network we get $\hat{\theta}^{(k)} = \psi_n^{(k)} = \psi_1^{(k+1)}$. Note the similarities between this algorithm and standard gradient descent. Essentially, at each subiteration an update is made by optimizing the function f_i (which only depends on local data at sensor i) for the current estimate. Now, the amount of in-network communication required for this algorithm is a function the number of cycles used therefore it is useful to understand the convergence behavior of this algorithm.

As presented above, the decentralized algorithm is a special case of incremental subgradient optimization methods. The general form of these optimization problems are concerned with minimizing functions of the form $f(\theta) = \sum_{i=1}^{n} f_i(\theta)/n$. We use the following result presented by Nedić and Bertsekas in [4] to characterize the tradeoff between estimator accuracy and resource usage. Assume that $||\nabla f_i||/n \leq c$. Then for a step size α, after K cycles with

$$K = \left\lfloor \frac{||\theta^* - \theta^{(0)}||^2}{\alpha^2 c^2} \right\rfloor \tag{5}$$

we have

$$f(\theta^{(K)}) \leq f(\theta^*) + \alpha c^2. \tag{6}$$

Observe that as αc^2 goes to zero, the estimate $\hat{\theta}^{(K)}$ gets arbitrarily close to minimizing our objective function, f. Thus, using (6) we can set a step size α in order to be within $\epsilon = \alpha c^2$ of the optimal cost, and (5) bounds the number of cycles required to achieve this accuracy. See [4] for a proof.

It should be noted that the description we give here is in terms of the gradient of f since in this paper we are optimizing a differentiable function, but in general these methods can be applied to nondifferentiable functions using subgradients rather than the gradient. Thus, the algorithm could easily extended to optimize a more complicated cost function – for instance, one depending on multi-modal measurements. The framework can also be modified to compute the MAP estimate by considering a prior density at each sensor.

Using these convergence results we can precisely quantify the performance of our algorithm for estimating the location of a stationary source. In general it is much more difficult to theoretically analyze the tracking performance of an algorithm without making assumptions about the source's dynamic behavior.

3. DISTRIBUTED VS. CENTRALIZED PROCESSING: ENERGY-ACCURACY ANALYSIS

We assume that all communication of data through the network is packet-based and multihop as opposed to direct communication. Based on this model, the expected total energy used for in-network communication by any sensor network algorithm can be written as

$$\mathcal{E}(n) = b(n) \times h(n) \times e(n), \tag{7}$$

where $b(n)$ is the total number of packet transmissions, $h(n)$ is the average number of hops over which each packet is transmitted, and $e(n)$ is the average amount of energy required to transmit one packet over one hop. In general, $e(n)$ depends on the density of nodes and the actual communication system being used, thus we express energy consumption in terms of $e(n)$.

In one cycle of our incremental algorithm each node makes a single communication – the current location estimate – to its nearest neighbor. Thus, the average number of packets transmitted is $b_{incr}(n) = O(nK)$, where K is the number of cycles required for the desired level of accuracy determined according to (5) and (6). All communications in this scheme are between neighboring nodes so $h_{incr}(n) = 1$. The total expected energy usage with the incremental algorithm for accuracy ϵ is

$$\mathcal{E}_{incr}(n) = O\left(\frac{n}{\epsilon^2}\, e(n)\right). \tag{8}$$

That is, the total energy used in our decentralized approach grows linearly with the number of sensors, and is inversely related to the desired level of accuracy. Note that there is no dependence on the amount of data collected at each sensor since only the source location estimate is communicated through the network and not raw data.

Alternatively, consider the the approach where each sensor transmits all of its data to a fusion center. The number of packets transmitted is $b_{cen}(n) = O(mn)$. If the n nodes are uniformly distributed over a two-dimensional field then the average number of hops from any sensor to the network perimeter is $h_{cen} = O(\sqrt{n})$, and the expected total communication energy for the centralized approach is $\mathcal{E}_{cen}(n) = O(mn^{3/2} e(n))$. In comparison to (8), we see that when $\epsilon^2 \geq 1/(m\sqrt{n})$ the incremental approach is more efficient in terms of resource consumption. Thus, as both the size of the network and the amount of data collected at each node grow, decentralized in-network processing is much more advantageous than centralized processing from the perspective of energy and bandwidth consumption.

4. SIMULATION RESULTS

4.1. Locating a Stationary Source

We have simulated a scenario where 100 sensors are uniformly distributed over a 100×100 field, with a stationary acoustic source positioned at a random location. From this point forward we also assume that $\beta = 2$. Of the 100 sensors in the field, those located with a radius of 15 from the source make 2 measurements of received signal strength and then implement the incremental subgradient method described above to find a maximum likelihood estimate for the source location (detecting the presence of a source from local data is a separate problem not addressed in this paper). The first sensor to detect the presence of an acoustic source and generates an initial location estimate by choosing a point in a random direction at a distance from itself based on its local data. Figure 1 shows contours of the global cost function being optimized for one simulation. Solid black dots indicate active sensor locations and the black square in the center of the figure is the source location. In this case, 11 of the 100 sensors were in active range of the source. Parameters A and σ_w^2 were such that the average SNR at active sensors was -11dB. The dashed line shows the solution trajectory with step size $\alpha = 1$. Circles along the trajectory denote the location estimate after each cycle. The estimate converges to the optimal solution after only a few cycles.

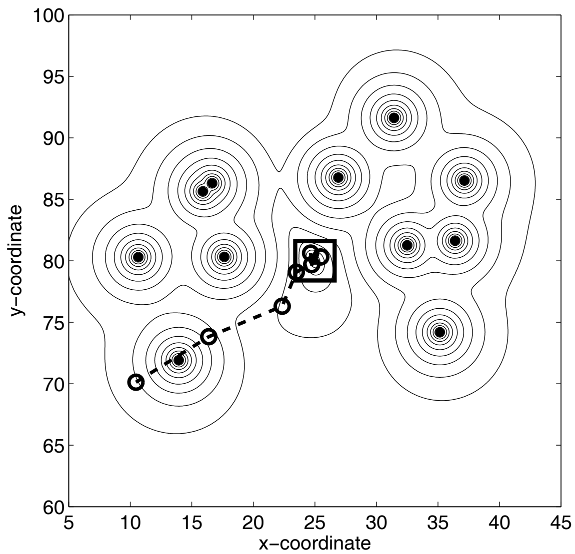

Fig. 1. Cost function contours and the solution trajectory for one simulation. In this case 15 of the 100 sensors in the 100×100 field detect the source and collaborate to identify its location. The average SNR at an active sensor was -11dB.

Figure 2 illustrates the performance of our algorithm at various step sizes. For each step sizes α between 0.1 and 5 (in increments of 0.1) the simulation as described above was repeated 500 times. The solid line shows the average root mean squared error (RMSE) and the dashed line represents the average number of optimization cycles before the decentralized algorithm reached a solution with the desired level of accuracy. As suggested by the theory, a more accurate solution can be achieved with smaller step sizes, however more cycles (and thus more energy and bandwidth) are required.

4.2. Tracking a Moving Source

Next we analyze the performance of incremental subgradient methods for tracking a moving source. In the case of a stationary source we saw the trade-off between the accuracy of our position estimate and the number algorithmic cycles required to reach this accuracy. The number of cycles is directly proportional to the amount of communication required. In general it is difficult to theoretically analyze the performance of tracking algorithms without making some assumptions about source dynamics (e.g. assumptions on maximum source velocity or change of direction speeds). Through simulations we find that accuracy is still directly proportional to the step size, α. However, as a consequence of the time con-

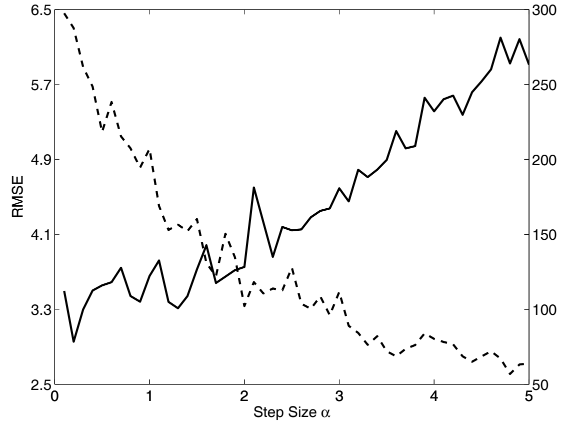

Fig. 2. The energy-accuracy tradeoff for locating a static source using incremental subgradient methods. Both RMSE (solid line) and the average number of cycles (dashed line) before reaching the desired level of accuracy are shown as functions of the step size, α. Each data point represents the average over 500 simulations with different sensor configurations and source locations.

strained nature of the tracking application, the acceptable number of cycles before reaching an accurate solution is limited. Thus, for on-line tracking there is a restriction on processing time in addition to energy and bandwidth constraints.

To illustrate these tradeoffs we have simulated the situation where a source moves across a 100×100 field, taking the dashed line trajectory displayed in Figure 3. The source moves from left to right across the field with a constant horizontal velocity, so that the true source velocity is proportional to the slope of it's trajectory. At each point in time, sensors within a radius of 15 of the source detect its presence, are activated, and make two signal strength measurements. Parameters A and σ_w^2 are again set so that the average SNR at active sensors is -11dB. To account for the time constrained nature of the tracking problem we only allow one optimization cycle at each step in time. The active nodes circulate an estimate of the source location once, and this position estimate is then used as the initial estimate at the next point in time. Thus, the variable k in the incremental update equations acts as a discrete time variable and $\widehat{\theta}^{(k)}$ is the estimated source location at time k.

We repeated this experiment 100 times, with new sensor locations drawn randomly on each trial. Figure 4 shows the average RMSE at each point in time, for different step sizes. For smaller step size values the algorithm initially takes longer to "lock on" to the source location, but after locking on the average estimation error is lower than the error for larger step sizes. The increase in RMSE at later times (as the source reaches the right side of the sensor field) is due to the fact that the source velocity increases as it moves along this portion of the path (roughly from $x = 70$ to 100). When a smaller step size is used the tracking algorithm can no longer keep up with the source, so the RMSE increases.

Figure 5 depicts the average RMSE over all time steps at different step sizes. The energy-accuracy analysis for the static source version of the decentralized algorithm revealed that estimator accuracy is inversely proportional to step size, α, and that the desired level of accuracy increases like $1/\sqrt{\alpha}$. However, we see effects due to the time-constraints in source tracking at very small step sizes. Specifically, in the static case small step sizes allowed for a

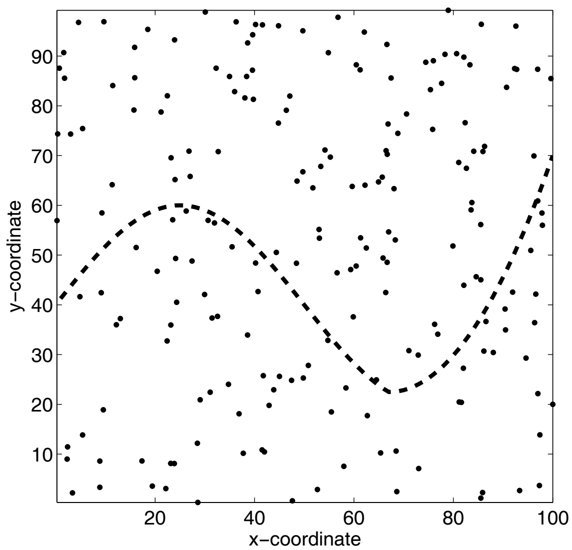

Fig. 3. In this experiment a source moves from left to right across a field taking the trajectory shown by the dashed line. Dots represent the locations of 200 sensors for one realization of the experiment.

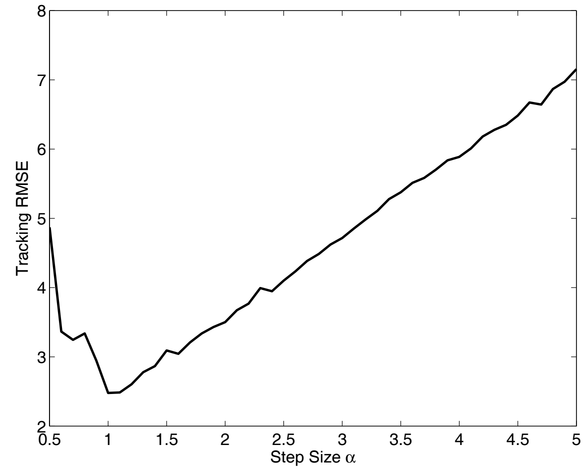

Fig. 5. RMSE averaged over the entire trajectory, for different step sizes. From the section on locating a stationary source we expect accuracy to increase linearly with step size. However in the tracking application there are also time constraints which limit the number of incremental subgradient cycles which can be performed at each step in time. Consequently, average accuracy is hindered for very small step sizes which require more time to reach an accurate estimate.

more accurate solution but required more optimization cycles. Because the number of cycles is constrained by time in the tracking application, the RMSE increases for very small step sizes ($\alpha < 1$).

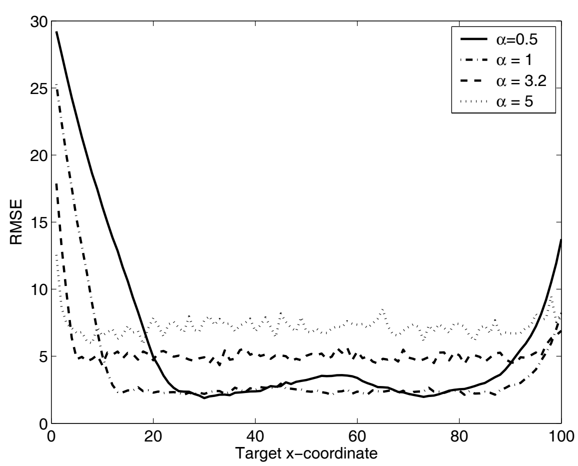

Fig. 4. The RMSE at each source position, plotted for different step sizes. With larger step sizes the algorithm locks on to the source faster but with a greater average error.

5. CONCLUSION AND FUTURE WORK

In this paper we have described a new decentralized algorithm for localizing and tracking a source. Sensors circulate a location estimate through the network and perform an update which resembles one step of a gradient descent algorithm using only their local data. The tradeoff between accuracy and the amount of communication is quantified and can be controlled through the choice of a step size. We show that the expected amount of communication for our decentralized algorithm is independent of the amount of data collected at each sensor and grows linearly with the number of sensors in the network. In comparison to a centralized processing scheme, the decentralized algorithm makes much better use of energy and bandwidth resources as either the size of the network and the amount of data collected by each sensor grow.

As described here, sensor communication occurs in a cycle through the network. Future work will be focused on developing an asynchronous version of the algorithm, as well as studying different scheduling schemes which balance the utility of information at each sensor with the cost of transmitting it. In this work we have also assumed that individual sensors can detect the presence of a source in their data, and this task will also be addressed in future work. Finally, our approach to the tracking problem thus far has been based on the assumption that the source is stationary over a short period of time. A future approach to tracking will include target dynamics in the model.

6. REFERENCES

[1] X. Sheng and Y.H. Hu, "Energy based acoust source localization," in *Proc. IPSN*, Palo Alto, CA, April 2003.

[2] F. Zhao and L. Guibas, Eds., *Information Processing in Sensor Networks*. Springer LNCS, 2003.

[3] F. Zhao, J. Shin, and J. Reich, "Information-driven dynamic sensor collaboration for tracking applications," *IEEE Sig. Proc. Mag.*, March 2002.

[4] A. Nedić and D. Bertsekas, *Stochastic Optimization: Algorithms and Applications*, chapter Convergence Rate of Incremental Subgradient Algorithms, pp. 263–304, Kluwer Academic Publishers, 2000.