Distributed Detection with a Minimalistic Signal Model: A Framework for Exploiting Correlated Sensing

Sriram Venkateswaran and Upamanyu Madhow
Department of Electrical and Computer Engineering
University of California Santa Barbara,
Santa Barbara, CA 93117, USA
Email: {sriram, madhow}@ece.ucsb.edu

Abstract— **We propose and investigate a minimalistic model for distributed detection using a sensor network. The signal of interest is a priori unknown. When a signal is present, sensors receive scaled, delayed and noisy versions of it, and signal presence is decided solely based on the correlation between sensor observations. We obtain encouraging performance results for both centralized and distributed detection, and subsequent centralized signal estimation. We observe that temporal alignment of sensor observations prior to combining is the bottleneck that determines the SNR threshold after which these methods work well. For ideal temporal alignment, detection performance improves exponentially with the number of sensors.**

I. Introduction

We consider a problem in distributed detection and estimation motivated by the goal of designing sensor networks which can "discover" new phenomena for which explicit signal models are not yet available. In order to undertake such a design, we must first answer the following fundamental question: how do we draw a distinction between "signal" (i.e., a manifestation of an interesting phenomenon) and "noise" (i.e., contributions to sensor readings not associated with interesting phenomena)? The need for making a careful distinction is not as important if the signal-to-noise ratio (SNR) is high enough (e.g., if we can define signal presence simply by a sensor reading exceeding a threshold), but becomes critical at low SNR, which might be typical for a sparse sensor deployment. In this paper, we consider one possible approach for making such a distinction, based purely on the correlation between neighboring sensors: an interesting signal is present if neighboring sensors show correlated readings, whereas uncorrelated sensor readings correspond to noise. We validate our approach in a simple setting in which all sensors received a scaled and delayed version of the signal, corrupted by additive white Gaussian noise (AWGN) which is uncorrelated across sensors. In each time window, or epoch, the sensor network must decide whether or not a signal is present. If it decides that a signal is present, then it must construct a good estimate of the signal from the sensor readings. One of the key requirements is temporal alignment of the sensor data prior to combining for detection or estimation.

In order to obtain fundamental performance benchmarks, we first consider a centralized system with temporally aligned sensor data. Treating the unknown signal as a nuisance parameter, we consider a Generalized Likelihood Ratio Test (GLRT) which reduces to comparing the largest eigenvalue of the sample covariance matrix to a threshold. Using asymptotic results from random matrix theory, we show that this decision rule guarantees that the probability of miss decays exponentially (with the number of sensors) for a fixed probability of false alarm. When the received signals are not temporally aligned, we continue with the GLRT approach, treating the unknown delays as additional nuisance parameters. The delay estimation required for this approach turns out to be the technical bottleneck that determines the SNR threshold beyond which the system exhibits good detection performance. We then investigate a distributed system in which the communication costs for monitoring for signal presence are reduced by restricting only a subset of the sensors to broadcast their observations in each epoch, in order to initiate a procedure for arriving at a consensus as to whether a signal is present. Finally, we show that, once the data is synchronized, the sensor observations can be combined to produce a reasonable estimate of the signal. Performance results for synthetic signals, as well as for an acoustic bird call, show that the performance of both signal detection and estimation is significantly improved relative to that obtained using a single sensor, despite the minimalistic signal model.

Related work: Most prior research on distributed detection (e.g., see [1],[2], and the references therein) is based on explicit signal models, unlike the implicit modeling here. Reference [3] investigates detection of an unknown signal in a multi-channel system, but assumes that the data on the various channels are synchronized (as we point out, temporal alignment is a major bottleneck in our system). Source localization and estimation with multiple sensors, where there is no explicit signal model, is investigated in [4], [5], [6].

This work was supported in part by the Institute for Collaborative Biotechnologies under grant DAAD19-03-D-0004 from the US Army Research Office, the Office for Naval Research under grant N00014-06-1-0066 and the National Science Foundation under grant CCF-0431205.

However, the limits of signal detection are not considered and the issue of temporal alignment is treated with different assumptions. Furthermore, the impact of communication costs is not considered in this prior literature, which focuses on sensor arrays rather than sensor networks.

II. System Model

There are N sensors, labeled $\mathcal{S}_1, \mathcal{S}_2, \ldots, \mathcal{S}_N$. We consider a discrete time system with real-valued observations, in which the observation window over which a detection decision is to be made spans T time instants. Let $\mathbf{y}_i$ denote the received signal at the i^{th} sensor over the observation window i.e. $\mathbf{y}_i = \big(y_i[0], y_i[1], \ldots, y_i[T-1]\big)^T$, modeled as follows:

Null Hypothesis H_0: When no event of interest occurs in the time window, the signals at the sensors are independent, discrete-time, unit variance, White Gaussian Noise (WGN) i.e. $\mathbf{y}_i = \mathbf{n}_i \ (1 \leq i \leq N)$ where $\mathbf{n}_i \sim N(\mathbf{0}, \mathbb{I})$ and $E(\mathbf{n}_i \mathbf{n}_j^T) = 0$ when $i \neq j$.

Alternate Hypothesis H_1: When an event of interest occurs within the window, the received signals are scaled and delayed versions of a common signal corrupted by noise. To avoid edge effects, we assume that the non-zero portions of the signal are received entirely in the same time window by all sensors. With this assumption, we can simplify notation, denoting the common signal $\mathbf{s} = \big(s[0], s[1], \ldots, s[T-1]\big)^T$ as spanning the entire observation window of length T (even though, in practice, the observation window is chosen to be longer than the expected signal length), and interpreting delays as circular shifts. Let h_i denote the channel gain and d_i denote the delay seen by the i^{th} sensor. Let $\mathcal{D}^k$ denote the delay operator, that delays a signal (circularly) by k discrete-time instants. We then have the following model under H_1:

$$\mathbf{y}_i = h_i \mathcal{D}^{d_i} \mathbf{s} + \mathbf{n}_i, \quad 1 \leq i \leq N \tag{1}$$

We make a couple of general observations: (1) We can only measure relative delays between sensors and not the "absolute" delays themselves i.e we can only estimate $\{d_2 - d_1, d_3 - d_1, \ldots, d_N - d_1\}$. Therefore, without loss of generality, we can set $d_1 = 0$ and estimate only the other variables. (2) Once the signal has been detected, it can only be estimated up to a scale factor, since the channel gains and the signal waveform are unknown. Therefore, when performing signal reconstruction, we constrain it to have unit energy, and use the correlation coefficient between the true signal and the estimated waveform to judge the fidelity of our estimate.

III. Centralized Detection & Estimation

In this section, we develop detection and estimation algorithms assuming that the sensor observations $\{\mathbf{y}_i, i = 1, ..., N\}$ are available at a centralized location. Before addressing the relative delays between the sensor observations, we first consider temporally aligned signals to develop performance benchmarks, and insight into the structure of the decision statistics.

A. Temporally Aligned Signals

As a first step, suppose that, when there is a common signal of interest, the sensed signals are temporally aligned, so that we can set $d_1 = d_2 = \ldots = d_N = 0$ without loss of generality. Even in this simplified setting, it is difficult to come up with detection rules that satisfy precise definitions of optimality in the absence of an explicit signal model [7], [1]. Therefore, we first discuss the optimal estimate of the signal waveform in this scenario and use this to obtain a GLRT detection rule.

1) Estimation: When the received signals are temporally aligned, the model simplifies to: $\mathbf{y}_i = h_i \mathbf{s} + \mathbf{n}_i \ (1 \leq i \leq N)$. The optimal solution for this problem is discussed in [5], [8], but is briefly reviewed here. Let us denote the matrix $[\mathbf{y}_1 | \mathbf{y}_2 | \ldots | \mathbf{y}_N]$ by $\mathbb{Y}$. The optimal estimate of the signal amounts to choosing the dominant left singular vector of $\mathbb{Y}$, which is the same as the dominant eigenvector of the sample covariance matrix(upto a scale factor) $\mathbb{R}_y = \mathbb{Y}\mathbb{Y}^{\mathbb{T}} = \sum_{i=1}^{N} \mathbf{y}_i \mathbf{y}_i^T$.

2) Detection: Given that the best estimate of the signal waveform is the dominant eigenvector of $\mathbb{R}_y$, we propose a GLRT rule that compares λ_1, the largest eigenvalue of $\mathbb{R}_y$, to a threshold to decide between H_0 and H_1. Owing to space constraints, we omit the proof of this fact. The rule is designed as follows:

3) Setting the threshold: We use the asymptotic distribution (as $N, T \rightarrow \infty$) of the largest eigenvalue under the null hypothesis to set the threshold for a given probability of false alarm. The rationale behind this is that, even for small values of N and T (for example, $N = T = 10$), the asymptotic distribution serves as a very good approximation to the true distribution ([9], [10]). We present the results regarding the asymptotic distribution that have been put together from [11], [9], [10] in a form suitable for our purpose. We consider $T > N$. The scenario $N > T$ is exactly analogous, and can be handled by interchanging the time and space (i.e., sensor label) variables.

Result: Let λ_1 denote the largest eigenvalue of $\mathbb{R}_y = \sum_{i=1}^{N} \mathbf{y}_i \mathbf{y}_i^T$ where $\mathbf{y}_i \in \mathbb{R}^T$, $\mathbf{y}_i \sim N(\mathbf{0}, \mathbb{I})$ and $E(\mathbf{y}_i \mathbf{y}_j^T) = 0$ $\forall\ i \neq j$ Let $\gamma = \frac{T-1}{N}$. We define center and scaling constants, μ_{NT} and σ_{NT}, as:

$$\mu_{NT} = N \times (1 + \sqrt{\gamma})^2, \sigma_{NT} = N^{\frac{1}{3}} \times (1 + \sqrt{\gamma})\big(1 + \frac{1}{\sqrt{\gamma}}\big)^{\frac{1}{3}}$$

As $N, T \rightarrow \infty$ with $\frac{T}{N} \rightarrow \gamma$, the cumulative density function of the random variable $Z = \frac{\lambda_1 - \mu_{NT}}{\sigma_{NT}}$ converges to the Tracy-Widom distribution of order 1 which we denote by $F_1(z)$ [11]. We propose a detection test of the form $Z \gtrless_{H_0}^{H_1} \eta$, where η is a threshold to be decided. Given a tolerable false alarm probability ϵ_{fa}, we set the threshold (η) to be $F_1^{-1}(1 - \epsilon_{fa})$. A key point to note is that, for a given value of ϵ_{fa}, η doesn't change with the parameters of the problem (N, T) as long as they are sufficiently large. We now bound the probability of miss of this decision rule and examine its asymptotics with the number of sensors.

Let $\mathbf{e}_s = \frac{\mathbf{s}}{||\mathbf{s}||}$ denote a unit vector in the direction of $\mathbf{s}$

and $E_s = ||\mathbf{s}||^2$ denote the energy in the signal. Let $E_r = E_s \times \left(\frac{1}{N}\sum_{i=1}^N h_i^2\right)$ denote the average received energy (of the signal component) at the sensors. Under H_1, we lower bound λ_1 (which is a random variable) for any realization of the noise. By the definition of the largest eigenvalue of $\mathbb{R}_y$, we have

$$\lambda_1 \geq \mathbf{e}_s^T \mathbb{R}_y \mathbf{e}_s = \mathbf{e}_s^T \left(\sum_{i=1}^N (h_i\mathbf{s} + \mathbf{n}_i)(h_i\mathbf{s} + \mathbf{n}_i)^T \right) \mathbf{e}_s \quad (2)$$
$$= \sum_{i=1}^N (h_i\sqrt{E_s} + w_i)^2 \triangleq V$$

where $w_i = \mathbf{e}_s^T \mathbf{n}_i \sim N(0,1)$ and $E(w_i w_j) = \delta[i-j]$. Thus, V is a lower bound to λ_1 for any realization of the noise. The random variable V has a non-central chi squared distribution with N degrees of freedom and a non-centrality parameter $\theta = \sum_{i=1}^N h_i^2 E_s = NE_r$. Let us define $\beta = \frac{\mu_{NT} + \eta\sigma_{NT}}{N}$. Using $\lambda_1 \geq V$ and a Chernoff bound to upper bound $P(V < N\beta)$, we can show that for $E_r > \beta - 1$, the probability of miss is upper bounded as $P_{miss} \leq e^{-N\psi_{unknown}}$ with $\psi_{unknown} > 0$ given by

$$\psi_{unknown} = \beta s_{min} + \frac{1}{2}\ln(1 - 2s_{min}) - \frac{E_r s_{min}}{1 - 2s_{min}}$$
$$\text{where } 1 - 2s_{min} = \frac{1 + \sqrt{1 + 4\beta E_r}}{2\beta} \quad (3)$$

For large N, $\beta \approx (1 + \sqrt{\gamma})^2$ and the constant $\psi_{unknown}$ becomes independent of N.

While we have fixed the probability of false alarm in the preceding design, it is also possible to choose the threshold to achieve an asymptotic decay in both the false alarm and miss probabilities in the number of sensors.

The price of minimalism: We compare the Chernoff bound above with the rate of decay of the miss probability when the signal model is known. To focus on the key issue, we assume that the channel gains are all unity. The rate of decay of the miss probability when the signal model is known, which we denote by ψ_{known}, is given by,

$$\psi_{known} = \frac{1}{2} \times \left[\sqrt{E_r} - \frac{Q^{-1}(\epsilon_{fa})}{\sqrt{N}}\right]^2 \approx \frac{E_r}{2} \quad \text{(for large } N\text{)}.$$

As Figure 1 shows, while an exponential improvement in performance can be achieved in our system by increasing the number of sensors (as long as the SNR is above a minimum value), there is indeed a steep price paid for not knowing the signal model.

B. Temporal Alignment Algorithm

When the sensed signals are not temporally aligned (under H_1), we need to correct for the timing offsets before we can use the largest eigenvalue of the sample covariance matrix as a decision statistic. The Maximum Likelihood (ML) estimate of the relative delays (jointly estimated with the signal and the channel gains) would involve searching over all possible delays for each sensor, and the complexity grows exponentially with the number of sensors. We therefore propose a GLRT-like detector, treating the delays as nuisance parameters, based

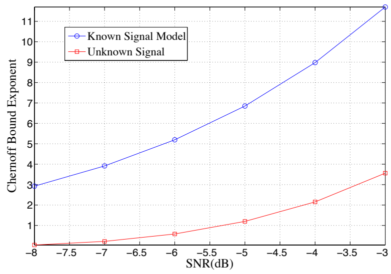

Fig. 1. A comparison of the Chernoff bound exponents with $N = 20$ sensors and $\gamma = 3$, as a function of SNR. Since the channel gains are taken to be unity, the SNR is same at all the sensors.

on a simpler two-stage delay estimation algorithm, as follows. In the first stage, we make a coarse estimate of the relative delay between sensor i and sensor j. After estimating the delay between every pair of sensors, we use consistency checks among the estimated delays to arrive at a set of delay estimates with respect to sensor $\mathcal{S}_1$.

1) Stage 1: Pairwise Delay Estimation: Let $R_{ij}[k]$ denote cross-correlation of the signals received at sensors i and j which is given by $R_{ij}[k] = \sum_n y_i[n]y_j[n-k]$. The coarse estimate of the relative delay $d_i - d_j$ is given by, $\hat{d}_{ij} \triangleq \arg\max_k |R_{ij}[k]|$

2) Stage 2: Pruning the estimates: We note that the delay estimates that are made are not all independent and need to be internally consistent. These constraints help us make a more reliable estimate of the delays with respect to $\mathcal{S}_1$. Let e_{ij} be the error incurred in the estimation of $d_i - d_j$ i.e., $d_i - d_j = \hat{d}_{ij} + e_{ij}$. The errors are not independent and in fact, their statistics (conditioned on H_1) would depend on the signal and its autocorrelation properties. Numerical simulation indicates that the probability density functions of the errors are well modeled by a Laplacian distribution. Ignoring the dependence between the errors, this motivates an "ML" estimator at the second stage that minimizes the cost function: $J = \sum_{i=1}^N \sum_{j=1}^N |e_{ij}| = \sum_{i=1}^N \sum_{j=1}^N |\hat{d}_i - \hat{d}_j - \hat{d}_{ij}|$. The preceding cost function depends on the l_1 norm of the residuals (e_{ij}) rather than the l_2 norm, which handles outliers due to large estimation errors in the first stage better [12], which is especially crucial at low SNR. Since the objective function J is convex in the delays, the delay estimation in the second stage is accomplished efficiently using a convex optimization solver [13].

The temporal alignment algorithm, which largely dictates the overall system performance, performs very well for moderate and large values of SNR, but fails for extremely low values of SNR. Thus, the detection performance with our GLRT approach exhibits a distinct thresholding effect, improving rapidly once the SNR is above a minimum value required for effective temporal alignment.

C. Summarizing Centralized Algorithms

In this part, we summarize the centralized algorithms for the case when the signals are not temporally aligned.

(1) Number the sensors arbitrarily from 1 to N. Run the two-stage time synchronization algorithm to estimate the delays at various sensors relative to sensor 1. Call the estimates $(\hat{d}_1 = 0, \hat{d}_2, \ldots, \hat{d}_N)$. (2) Form the sample-covariance matrix of the delay-corrected version of the sensed signals $\mathbb{R}_y = \sum_{i=1}^{N} \big(\mathcal{D}^{-\hat{d}_i}\mathbf{y}_i\big)\big(\mathcal{D}^{-\hat{d}_i}\mathbf{y}_i\big)^T$. (3) Let λ_1 denote the largest eigenvalue of $\mathbb{R}_y$ and $\mathbf{v}$ denote the corresponding eigenvector. (4) **Detection:** Fix a threshold η based on the tolerable false-alarm probability using the Tracy-Widom distribution. Compare $Z = \frac{\lambda_1 - \mu_{NT}}{\sigma_{NT}}$ with η and declare H_1 if $Z > \eta$ and H_0 otherwise. (5) **Estimation:** Declare $\frac{\mathbf{v}}{||\mathbf{v}||}$ to be estimate of the signal. We will use the normalized inner product maximized over all possible shifts, given by $\rho = \max_\delta \frac{|\mathbf{v}^T(\mathcal{D}^\delta \mathbf{s})|}{||\mathbf{v}|| \times ||\mathbf{s}||}$ to quantify the fidelity of our estimate.

IV. Distributed Detection

For a network of distributed sensors, centralized detection requires that each sensor send its observation over each window when monitoring for signal presence. The associated energy consumption may be excessive, especially when signals of interest appear rarely. We therefore investigate a distributed system in which the amount of communication associated with monitoring for signal presence is reduced: in a given window, only a subset (of size $M < N$) of the sensors broadcast their received signals to all other sensors. The broadcasting subset can change over different observation windows in order to increase the network lifetime. For our purpose, we consider a specific observation window, and describe a distributed detection algorithm, assuming, without loss of generality, that sensors $\mathcal{S}_1, \mathcal{S}_2, \ldots, \mathcal{S}_M$ have broadcast their observations.

1) Sensors $\mathcal{S}_1, \mathcal{S}_2, \ldots, \mathcal{S}_M$ broadcast their sensed signals $\mathbf{y}_1, \mathbf{y}_2, \ldots, \mathbf{y}_M$ respectively.
2) **Round 1:** Sensor $\mathcal{S}_i$, $i > M$, uses the broadcast signals along with its own received signal ($\mathbf{y}_i$) to make an estimate of delays of the signals received at $\mathcal{S}_2, \mathcal{S}_3, \ldots, \mathcal{S}_M$ and $\mathcal{S}_i$ with respect to $\mathcal{S}_1$ using the two-stage time synchronization algorithm. We shall denote these estimates by $\tau_1^{(i)} = 0, \tau_2^{(i)}, \ldots, \tau_M^{(i)}, \tau_i^{(i)}$ where $\tau_j^{(i)}$ denotes the delay of the signal received at $\mathcal{S}_j$ with respect to $\mathcal{S}_1$ as estimated by $\mathcal{S}_i$.
   Similarly, sensors $\mathcal{S}_1, \mathcal{S}_2, \ldots, \mathcal{S}_M$ use the broadcast messages to estimate the delays with respect to $\mathcal{S}_1$. Note that the information available at these sensors is the same and therefore, the delay estimates made at these sensors are the same. We denote these delay estimates by $\tau_1^{(0)}, \tau_2^{(0)}, \ldots, \tau_M^{(0)}$.
3) **Round 2:** At the end of the first stage, every sensor broadcasts the delay estimates that it made. This is done so that more reliable estimates of the delays are made by pooling in the information that is common to all sensors. Sensor $\mathcal{S}_i$ broadcasts $(\tau_1^{(i)} = 0, \tau_2^{(i)}, \ldots, \tau_M^{(i)})$. Each sensor estimates the delay of the signal received at $\mathcal{S}_k$ $(1 \le k \le M)$ with respect to $\mathcal{S}_1$, to be $\hat{d}_k \triangleq \text{median}(\tau_k^{(0)}, \tau_k^{(M+1)}, \ldots, \tau_k^{(N)})$
4) **Round 3:** Each sensor forms the matrix $\tilde{\mathbb{R}}_y$ given by $\tilde{\mathbb{R}}_y = \sum_{i=1}^{M} \big(\mathcal{D}^{-\hat{d}_i}\mathbf{y}_i\big)\big(\mathcal{D}^{-\hat{d}_i}\mathbf{y}_i\big)^T$. Let $\mathbf{v}_y$ denote the dominant eigenvector of $\mathbb{R}_y$ which has unit norm. Sensor $\mathcal{S}_k$ now transmits a single number quantifying its confidence in certifying the presence of a common signal. The "confidence indicator" (μ_k) is given by $\mu_k = \max_\delta |\mathbf{v}_y^T \mathcal{D}^{-\delta}\mathbf{y}_k|^2$.
   Now sensor $\mathcal{S}_1$ forms $\mu = \sum_{k=1}^{N} \mu_k$ and makes a decision on the existence of a common signal based on the value of μ.

V. Results

In this section, we present simulation results for the detection and the estimation problems. We consider both "wideband" and "narrowband" signals, where we use the terminology to refer to signals whose autocorrelation functions exhibit sharp and broad peaks, respectively. For known signal models, it is well known that wideband signals lead to better delay estimates [7]. However, in our setting, where the delays and the signal model are unknown, wideband signals actually lead to performance deterioration: missing a sharp autocorrelation peak leads to a potentially large error in the delay estimate, which in turn causes a deterioration in detection and estimation performance. Effectively, under our GLRT approach, the system has to work harder at implicitly estimating more degrees of freedom for a wideband signal, and hence performs worse, as illustrated by the results below.

In our simulation model, we set the number of sensors, $N = 20$. The wideband signal model consists of i.i.d. Gaussian random variables embedded in the "middle" of the observation window for different lengths of time to vary the SNR. The narrowband signals we have considered include synthetic signals such as triangular and rectangular waveforms and a portion of the pheasant's bird-call. For the bird call, the observation window length is taken to be $T = 400$. For all other signals, $T = 256$. The channel gains are modeled as independent random variables uniformly distributed over $[0.5, 1]$. For reference, the channel gain at sensor $\mathcal{S}_1$ is taken to be 1 and the SNR values reported correspond to those measured at sensor $\mathcal{S}_1$. The delays with respect to $\mathcal{S}_1$ are modeled as independent random variables that take values uniformly in the range [4,30]. The noise is modeled as independent and identically distributed (across time and sensors) standard Gaussian random variables. For the distributed detection algorithm, we set the number of broadcasting sensors $M = 6$.

We now briefly discuss the results.

1) The probability of a miss is quite small even for rather low values of tolerable false alarm probabilities.
2) The performance is worse for wideband signals due to the large losses of SNR pooling gain resulting from even small synchronization errors.
3) In both the centralized and distributed detection schemes, we observe the thresholding effect associated with temporal alignment. We observe that there is a significant improvement in the detection performance for an SNR increase of 0.5 dB (from -5.78 dB to -5.27 dB) in the wideband case. This occurs because of better

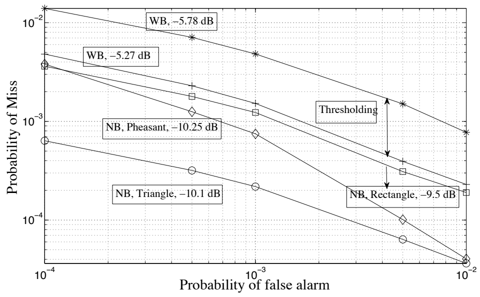

Fig. 2. Performance of Centralized Detection scheme with wideband and narrowband signals. The SNR values marked refer to those measured at the sensor where the signal strength is maximum.

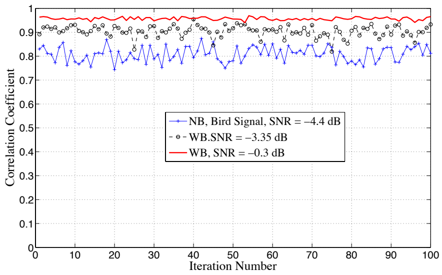

Fig. 3. Plot of correlation coefficient (ρ) between the reconstructed waveform and the original noiseless waveform

synchronization at a slightly higher SNR, which results in significantly better performance.

4) Figure 5 shows the noiseless section of the pheasant bird call and a noisy version of the same. We see that the reconstructed version of the bird-call closely resembles the original noiseless version even though the received signal at any individual sensor is extremely noisy.

5) From Figure 5, we see that correlation coefficient between the true signal and the estimated waveform is consistently high even when the SNR is low. However, it must be pointed out that estimation accuracy exhibits the same thresholding effect and excellent reconstruction requires a certain minimum SNR.

VI. Conclusion

We have shown that effective detection and estimation is possible using a minimalistic signal model based on correlation between neighboring sensors. The performance improves exponentially as a function of the number of sensors, assuming that temporal alignment of sensor readings can be achieved. The problem of temporal alignment appears to be the key bottleneck in our system, and determines the SNR threshold beyond which rapid performance improvements can be obtained by increasing the number of sensors. The example system considered here illustrates that sensor correlation potentially provides a powerful mechanism for discovering new phenomena of interest. Future research includes developing more detailed insight into the model considered here and its natural generalizations, such as multiple signals, colored noise, and multimodal sensing.

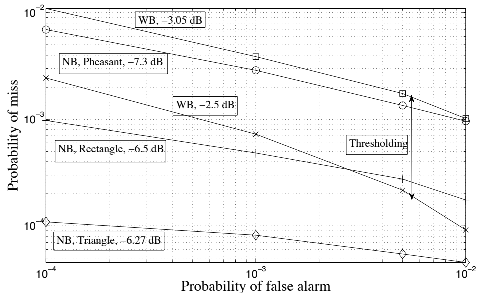

Fig. 4. Performance of Distributed Detection scheme with wideband and narrowband signals. The SNR values marked refer to those measured at the sensor where the signal strength is maximum.

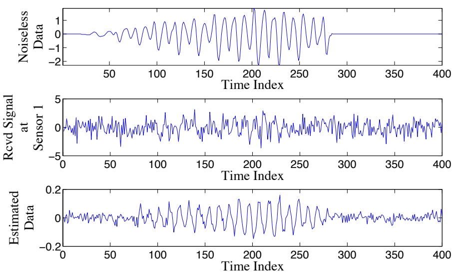

Fig. 5. Reconstruction of the pheasant's call. The uppermost figure shows the noiseless waveform. The middle figure shows the noisy waveform that was received at the first sensor. The bottom figure shows the reconstructed waveform which closely resembles the original

References

[1] R. Viswanathan and P. Varshney, "Distributed detection with multiple sensors I. Fundamentals," *Proceedings of the IEEE*, vol. 85, no. 1, pp. 54–63, 1997.

[2] R. Blum, S. Kassam, and H. Poor, "Distributed detection with multiple sensors I. Advanced topics," *Proceedings of the IEEE*, vol. 85, no. 1, pp. 64–79, 1997.

[3] D. Cochran, H. Gish, and D. Sinno, "A geometric approach to multiple-channel signal detection," *Signal Processing, IEEE Transactions on [see also Acoustics, Speech, and Signal Processing, IEEE Transactions on]*, vol. 43, no. 9, pp. 2049–2057, 1995.

[4] I. Ziskind and M. Wax, "Maximum likelihood localization of multiple sources by alternating projection," *Acoustics, Speech, and Signal Processing [see also IEEE Transactions on Signal Processing], IEEE Transactions on*, vol. 36, no. 10, pp. 1553–1560, 1988.

[5] J. Chen, R. Hudson, and K. Yao, "Maximum-likelihood source localization and unknown sensor location estimation for wideband signals in the near-field," *Signal Processing, IEEE Transactions on [see also Acoustics, Speech, and Signal Processing, IEEE Transactions on]*, vol. 50, no. 8, pp. 1843–1854, 2002.

[6] B. Ottersten, R. Roy, and T. Kailath, "Signal waveform estimation in sensor array processing," *Signals, Systems and Computers, 1989. Twenty-Third Asilomar Conference on*, vol. 2, 1989.

[7] H. Poor, "An Introduction to Signal Detection and Estimation," 1994.

[8] C. Van Loan, "The ubiquitous Kronecker product," *J. Comput. Appl. Math*, vol. 123, no. 1-2, pp. 85–100, 2000.

[9] I. Johnstone, "On the distribution of the largest eigenvalue in principal components analysis," *Ann. Statist*, vol. 29, no. 2, pp. 295–327, 2001.

[10] N. El Karoui, "On the largest eigenvalue of Wishart matrices with identity covariance when n, p and p/n tend to infinity," *Arxiv preprint math.ST/0309355*, 2003.

[11] C. Tracy and H. Widom, "Distribution Functions for Largest Eigenvalues and Their Applications," *Arxiv preprint math-ph/0210034*, 2002.

[12] S. Boyd and L. Vandenberghe, "Convex Optimization," 2004.

[13] M. Grant, S. Boyd, and Y. Ye, "CVX: Matlab Software for Disciplined Convex Programming," 2006.