

Collaborative Estimation in Dispersive Environments: A Frequency Domain Approach

Sriram Venkateswaran and Upamanyu Madhow

Department of ECE, University of California Santa Barbara, CA 93106

Email: {sriram, madhow}@ece.ucsb.edu

Abstract—We investigate the feasibility of reconstructing a signal recorded at multiple sensors through dispersive channels, with minimal prior information regarding the signal and the channels. We first demonstrate the efficacy of a parallelizable frequency domain algorithm which exploits the continuity of the channels across frequency, operating over small, overlapped frequency bins. The algorithm employs SVD-based estimates of the signal over each frequency bin (over which the channels are modeled as quadratic), and then stitches these estimates together based on consistency conditions across adjacent bins. While this algorithm gives excellent signal reconstruction over frequency bands that are one order of magnitude larger than the channel “coherence bandwidth,” fundamental limitations arise for larger frequency bands, in that it is possible to construct multiple explanations consistent with the recorded data.

I. INTRODUCTION

We investigate the fundamental problem of collaborative sensing of events regarding which we have limited prior knowledge. Without a prior model for the event signature, a single sensor has no means of distinguishing signal from noise. On the other hand, multiple sensors sensing a common event are expected to have correlated observations, which can potentially be pooled to obtain a reconstruction of the event signature. Specifically, we focus on the following setting, that is closely related to the classical problem of blind deconvolution: Each sensor observes a noisy version of a signal (the event signature) passed through a multipath channel, where the channels are different for different sensors. The signal and the multipath channels are *a priori* unknown, but we assume that there is a known upper bound on the channel delay spread. We would like to reconstruct the signal by pooling the observations from all the sensors, averaging out the noise and undoing the effects of multipath propagation. We would also like to understand whether there are fundamental ambiguities in reconstruction that cannot be resolved under our model.

Approach: We adopt a frequency domain approach, inspired by Orthogonal Frequency Division Multiplexing (OFDM) schemes used to combat dispersive channels in communication systems. The key to this approach is the observation that any dispersive channel can be approximated by a constant over a frequency band smaller than the channel coherence bandwidth. This allows us to efficiently estimate the signal in each band using a Singular Value Decomposition (SVD), thereby breaking down the difficult problem of blind deconvolution into a number of simpler subproblems. However, within each band, the signal can be estimated only up to complex scaling.

We resolve these scaling ambiguities in two stages – first, we refine the channel model and approximate it as a quadratic function of frequency within a small frequency band. We then exploit the continuity of the channels across adjacent frequency bands to “stitch” together the signal estimates across bands. We find that it is important to have significant overlap between adjacent bands for this approach to succeed; it is interesting to note that our inspiration for heavily overlapping bands, after obtaining poor performance with non-overlapping frequency bands as in OFDM, came from the front end of the human auditory system, which employs similar overlapping narrowband decomposition of the incoming signals [1].

Contributions: We devise and evaluate a two-stage algorithm for reconstructing the event signature. Within each frequency band, an alternating optimization algorithm is used for estimating the signal and the channels seen by each sensor, bootstrapping with a signal estimate obtained by first approximating the channels as constant. The signal (and channel) estimates are then combined across bands using a “left-to-right stitching” algorithm, based on enforcing consistency conditions on the signal and channel estimates in the overlap region for adjacent bands. Simulations and experiments on an acoustic testbed show that the algorithm gives excellent performance, but that there are occasional drops in correlation between the estimate and the true signal as the signal bandwidth gets large. We find that this is not an artifact of the algorithm, but is the result of fundamental limitations, in that, for large signal bandwidths (relative to the channel coherence bandwidth), multiple explanations for the recorded data can be constructed, while adhering to the assumed constraints on the channel delay spreads. However, due to space limitations, we omit detailed description of the procedure for constructing multiple explanations in this paper.

Related Work: We previously investigated collaborative sensing without a prior signal model in [2], but this work was for line of sight (LOS) channels from the source to the sensors. Source localization and implicit signal estimation is investigated in [3], but these results are also for a LOS channel model. The problem posed here falls within the broader framework of blind multichannel identification and equalization, on which there is a large body of literature (e.g., see [4], [5] for excellent reviews). Much of this work is not applicable to our setting, since it assumes prior knowledge of source statistics, such as oversampled second order statistics [6][7], or higher order moments [8]. However, prior work on

“deterministic subspace methods” [9] considers exactly the same scenario as the present paper, and prove that signal reconstruction is possible if the channels do not have common zeros. Unfortunately, the time domain algorithm in [9] does not scale to large datasets or large delay spreads, and, as the authors acknowledge, is sensitive to knowledge of channel length. To the best of our knowledge, the approach proposed in our paper is the first to provide a scalable solution to this problem (by breaking a complicated time domain problem into simple frequency domain subproblems), while requiring only coarse assumptions on channel length (to estimate the coherence bandwidth).

II. SYSTEM MODEL

We consider a network S sensors, where the source signature, denoted by $x[n]$, is recorded at the i th sensor after having passed through a dispersive channel $h_i[n]$ that is time-limited to P taps. Denoting the recorded samples by $y_i[n]$, we have,

$$y_i[n] = (x \star h_i)[n] + n_i[n], \quad n = 0, 1, 2, \dots, N - 1 \quad (1)$$

where $\star$ denotes the *linear* convolution operation between the sequences $\{x[n], n = 0, 1, \dots, N - 1\}$ and $\{h_i[n], n = 0, 1, \dots, P - 1\}$. The noise samples $n_i[n]$ are modeled as Gaussian random variables with variance σ^2 and are assumed to be independent across sensors and time. We denote the N -point DFT's of $y_i[n]$, $x[n]$ and $h_i[n]$ by $Y_i[k]$, $X[k]$ and $H_i[k]$ ($k = 0, 1, \dots, N - 1$), respectively. We use boldface notation to denote vectors containing samples of these signals in the time and frequency domains (e.g., $\mathbf{y}_i$ denotes the N -dimensional vector of samples at the i sensor, and $\mathbf{Y}_i$ the vector of corresponding DFT samples). Approximating the linear convolution in (1) by a circular convolution and taking the DFT of (1), we obtain

$$Y_i[k] = X[k]H_i[k] + N_i[k], \quad k = 0, 1, 2, \dots, N - 1 \quad (2)$$

where $N_i[k]$ are independent, complex Gaussian random variables with variance σ^2 .

Benefits of Frequency Domain Processing: The discrete frequency coefficients $H_i[k]$ are approximately constant over $L = B_{coh}T$ contiguous samples, where B_{coh} is the coherence bandwidth of the channel (in Hertz) and T is the recording duration (in seconds). Accordingly, we process the frequency domain samples of the recorded signal $Y_i[k]$ in blocks of L contiguous samples. With a $T = 4$ second recording and a coherence bandwidth $B_{coh} = 5$ Hz – numbers typical of an indoor acoustic setting – each block consists of only $L = 20$ samples. On the other hand, processing in the time domain involves substantially larger blocks: over the same recording interval, a sampling frequency $f_s = 16$ kHz results in 64000 samples. Thus, exploiting parallelism in the frequency domain breaks the problem into a number of subproblems of manageable dimensions, whose solutions then need to be “stitched” together. **Notation:** We denote the l^{th} element of a vector $\mathbf{U}$ by $\mathbf{U}[l]$. We process the recorded samples in frequency bands consisting of L contiguous indices, and index the frequency bands by b . If frequency band b corresponds to indices

$\{i_1, i_1 + 1, \dots, i_1 + L - 1\}$, then $\mathbf{U}[I_b]$ denotes the vector of corresponding samples from $\mathbf{U}$: $\mathbf{U}[I_b] = (\mathbf{U}[i_1] \ \mathbf{U}[i_1 + 1] \dots \mathbf{U}[i_1 + L - 1])$.

III. SIGNAL ESTIMATION ALGORITHM

We have already noted that even a complicated time domain channel exhibits a relatively simple structure – a quadratic function of frequency – over a small enough frequency band. Additionally, we note that the channel response varies continuously with frequency. We reconstruct the source from the recorded signals in two stages by exploiting these observations, as summarized below.

- *Stage 1:* We split the entire frequency range into small bands over which the channel response may be approximated by a quadratic. We bootstrap by approximating the channel as constant over band b , which allows us to employ an SVD to estimate the source in band b efficiently. This source estimate provides the starting point for an alternating optimization procedure that refines the estimates of the quadratically varying channels and the source signal.

The alternating optimization consists of multiple iterations of two basic steps:

- 1) Given an estimate of the source signal in band b , find the best estimates of channel responses to each sensor that are quadratic functions of frequency.
- 2) Given channel estimates to each sensor in band b that are quadratic functions of frequency, find the best estimate of the source within this band.

- *Stage 2:* Since the signal as well as channels are unknown, there is a scaling ambiguity in the signal estimate produced in stage 1. Thus, in order to reconstruct a signal with bandwidth larger than the channel coherence bandwidth, the estimates from different bands must be scaled consistently. This is accomplished by using overlap between successive bands and exploiting the continuity of the channel. Finally, we use the estimated scale factors along with the signal estimates from Stage 1 to reconstruct the source signal. We now provide the details.

A. Frequency domain channel model

At a frequency close to f_0 , say $f = f_0 + \delta$, the frequency response of a tapped delay line channel $h(t) = \sum_k \mu_k \delta(t - \tau_k)$ is given by,

$$H(f_0 + \delta) = \sum_k \mu_k e^{-j2\pi f_0 \tau_k} [1 - (j2\pi \tau_k) \delta - (2\pi^2 \tau_k^2) \delta^2 + \dots] \quad (3)$$

By choosing $|\delta| \ll 1/\tau_{max}$, where $\tau_{max} = \max_k |\tau_k|$ is the delay spread, cubic and higher powers of $\delta \tau_k$ can be ignored in the above expansion. Therefore, the channel can be approximated as a quadratic function of frequency over a small frequency band.

B. Stage 1: Estimation Within A Band

Consider a band b spanning the discrete frequencies $I_b = \{f_1 T, f_1 T + 1, f_1 T + 2, \dots, (f_1 + B_{coh})T\}$. The frequency domain samples at sensor i in this band are given by

$$Y_i[k] = H_i[k]X[k] + N_i[k], \quad k \in I_b.$$

Initial Guess: Approximating approximate the channel response to sensor i as a constant $\bar{H}_i$ within this band, we have

$$\mathbf{Y}_i[I_b] \approx \bar{H}_i \mathbf{X}[I_b] + \mathbf{N}[I_b], \quad (4)$$

where $\mathbf{N}[I_b] \sim \mathcal{CN}(0, \sigma^2 \mathbb{I})$. Since the processing in each band is identical, we drop the identity of the band from our notation in the following. Denoting the Maximum Likelihood (ML) estimates of the signal by $\hat{\mathbf{S}}$ and the channel for sensor i by $\hat{G}_i$, respectively, we have,

$$(\hat{\mathbf{S}}, \hat{G}_i) = \arg \min_{\mathbf{S}, G_i} \sum_{i=1}^S \|\mathbf{Y}_i[I_b] - G_i \mathbf{S}\|^2 \quad (5)$$

where S is the number of sensors. As shown in [2], the ML estimate of the signal $\hat{\mathbf{S}}$ is the eigenvector with the largest eigenvalue of the nonnegative definite matrix $\mathbf{M} = \sum_{i=1}^S \mathbf{Y}_i[I_b](\mathbf{Y}_i[I_b])^H$. The corresponding channel estimates are given by $\hat{G}_i = \hat{\mathbf{S}}^H \mathbf{Y}_i[I_b] / \|\hat{\mathbf{S}}\|^2$. These ML estimates are unique only up to scaling: $z\hat{\mathbf{S}}$ and $(1/z)\hat{G}_i$ give exactly the same value for the cost function in (5) for any nonzero complex scalar z .

Alternating Optimization: We now use the preceding signal estimate to refine the channel estimates, modeling their variation over the bin as quadratic. These refined channel estimates, in turn, yield an updated signal estimate.

Step (a) – Given signal, estimate channel: Based on our quadratic approximation, the i th channel is written as

$$\mathbf{G}_i[l] = A_i + B_i l + C_i l^2 \quad l = 0, 1, \dots, L-1 \quad (6)$$

Given a signal estimate $\hat{\mathbf{S}}$ in a frequency band, the corresponding ML estimate of the channel to the i th sensor is obtained by estimating the parameters (A_i, B_i, C_i) by solving the following linear least squares problem:

$$(\hat{A}_i, \hat{B}_i, \hat{C}_i) = \arg \min_{A_i, B_i, C_i} \sum_{l=0}^{L-1} \left| \mathbf{Y}_i[I_b, l] - \hat{\mathbf{S}}[l](A_i + B_i l + C_i l^2) \right|^2 \quad (7)$$

Step (b) – Given channel, estimate signal: Now, suppose we have channel estimates $\hat{\mathbf{G}}_i[l] = A_i + B_i l + C_i l^2$. Then the ML signal estimate $\hat{\mathbf{S}}[l]$ satisfies a “maximal ratio combining” rule:

$$\hat{\mathbf{S}}[l] = \frac{\sum_{i=1}^S \hat{\mathbf{G}}_i^*[l] \mathbf{Y}_i[I_b, l]}{\sum_{i=1}^S |\hat{\mathbf{G}}_i[l]|^2} \quad (8)$$

Output of Stage 1: After a few rounds of alternating optimization, the estimated signal and channel are approximately equal to scaled versions of the true signal and channel. Denoting the true signal and channels in band b by $\mathbf{X}[I_b]$ and $\mathbf{H}_i[I_b]$ and their estimates by $\hat{\mathbf{S}}_b$ and $\hat{\mathbf{G}}_{i,b}$, we have

$$z_b \hat{\mathbf{S}}_b \approx \mathbf{X}[I_b] \quad 1/z_b \hat{\mathbf{G}}_{i,b} \approx \mathbf{H}_i[I_b] \quad (9)$$

where z_b is an arbitrary complex scalar.




Fig. 1. Overlapping frequency bands

C. L-to-R Stitching Algorithm

We estimate the weights z_b by choosing adjacent bands to have significant overlap with one another, and enforcing consistency in the overlapped region.

Reconciling multiple estimates of the same quantity: Consider adjacent bands $b-1$ and b , as shown in Figure 1. Denote the frequencies common to bands $b-1$ and b by F_{common} i.e. $F_{common} = [b\Delta, (b-1)\Delta + B_{coh}]$ with $\Delta = (1-\chi)B_{coh}$. The parameter $\chi \in (0, 1)$ controls the amount of overlap between adjacent bands.

Denote the channel to sensor i within F_{common} by $\mathbf{H}_{i,common}$. We obtain two estimates of $\mathbf{H}_{i,common}$ from Stage 1: The first estimate is obtained from band $b-1$ by picking the *rightmost* χL entries of $\hat{\mathbf{G}}_{i,b-1}$. The chosen entries are shown by crossed squares in Figure 1 and we denote this estimate by $\hat{\mathbf{G}}_{i,b-1}^{right}$. The second estimate is obtained from band b by picking the *leftmost* χL entries of $\hat{\mathbf{G}}_{i,b}$ and is denoted by $\hat{\mathbf{G}}_{i,b}^{left}$. From equation (9), these estimates are related as,

$$(1/z_{b-1})\hat{\mathbf{G}}_{i,b-1}^{right} \approx \mathbf{H}_{i,common} \approx (1/z_b)\hat{\mathbf{G}}_{i,b}^{left} \quad (10)$$

Similarly, we denote the *leftmost* χL entries of $\hat{\mathbf{S}}_b$ by $\hat{\mathbf{S}}_b^{left}$ and the *rightmost* χL entries of $\hat{\mathbf{S}}_{b-1}$ by $\hat{\mathbf{S}}_{b-1}^{right}$ and obtain (via (9)),

$$z_{b-1}\hat{\mathbf{S}}_{b-1}^{right} \approx z_b\hat{\mathbf{S}}_b^{left} \quad (11)$$

We combine the constraints in (10) and (11) to obtain a single constraint of the form,

$$z_b \mathbf{u}_{b,b-1} \approx z_{b-1} \mathbf{u}_{b-1,b} \quad (12)$$

where the vectors $\mathbf{u}_{b,b-1}$ and $\mathbf{u}_{b-1,b}$ are given by $\mathbf{u}_{b,b-1} = (\hat{\mathbf{S}}_b^{left}, \hat{\mathbf{G}}_{1,b-1}^{right}, \dots, \hat{\mathbf{G}}_{S,b-1}^{right})$ and $\mathbf{u}_{b-1,b} = (\hat{\mathbf{S}}_{b-1}^{right}, \hat{\mathbf{G}}_{1,b}^{left}, \dots, \hat{\mathbf{G}}_{S,b}^{left})$.

We estimate the scale factors $\{z_b\}$ based on (12) in two stages. First, we estimate the “relative” scale factor $\gamma_{b,b-1} \triangleq z_b/z_{b-1}$ between bands b and $b-1$ as

$$\gamma_{b,b-1} = \frac{\mathbf{u}_{b,b-1}^H \mathbf{u}_{b-1,b}}{\|\mathbf{u}_{b,b-1}\|^2} \quad (13)$$

We then set $z_0 = 1$ (without loss of generality) and estimate the scale factors in other bands recursively:

$$z_b = z_{b-1} \gamma_{b,b-1} \quad b \geq 1$$

We call this the *L-to-R stitching algorithm* since the weights are estimated in a recursive fashion, starting from low frequencies (left end of the spectrum) and proceeding on to high frequencies (right end of the spectrum). Note that, if the signal energy in a given frequency band is too low, neither the signal nor the channel estimates will be reliable, which can severely disrupt the stitching procedure. We therefore identify bands with low signal energy and omit them from our processing, giving them a weight of zero when reconstructing the signal. Details are omitted due to lack of space.

IV. RESULTS

We now present experimental results from our indoor testbed, as well as simulation results, to quantify the performance of the L-to-R stitching algorithm.

A. Experimental Results

We recorded acoustic data using three Samson C03U USB microphones. We played the source signal from four different locations to emulate recordings with a larger number of microphones. We thus had 12 recordings (3 microphones $\times$ 4 source locations) of the source through dispersive channels. We used two types of source signals: (1) *Chirp signals*: they have the general form $x_{chirp}(t) = \cos\left(2\pi\left(f_0 + \frac{f_1 - f_0}{T}t\right)t\right)$, $0 \leq t \leq T$. We chose $f_0 = 1000$ Hz, $T = 4$ seconds and varied f_1 to characterize the effect of signal bandwidth on the reconstruction. Specifically, we chose $f_1 = 1050$ Hz (we call this signal Chirp50) and $f_1 = 1200$ Hz (we call this signal Chirp200) for signal bandwidths of 50 Hz and 200 Hz respectively. (2) *Sinusoidal signals*: These signals, denoted by Sines50 and Sines200, consist of sinusoids spaced 2 Hz apart for bandwidths of 50 Hz and 200 Hz respectively. We used a sampling frequency of 16 kHz to record the data.

Preprocessing: The low frequency components in the recorded signals contain substantial energy from background hum. Since this is registered at all sensors, it counts as “signal” rather than spatially uncorrelated noise. In order to control the source signal, we filter out the received signal energy in the bands from 0-950 Hz to eliminate the hum. We coarsely synchronize the recorded waveforms in a data-driven fashion, by matched filtering each of them against a reference (recorded signal at one of the sensors) and then shifting them in time based on the matched filter output. Finally, the recordings at one of the twelve “sensors” - corresponding to a particular source-microphone arrangement - had a very good correlation with the true signal, indicating that it had a near Line-of-Sight channel to the source. Since our goal is to understand the limits of recovering signals in the face of significant multipath, we exclude this sensor from further processing.

Parameters & Results: We choose $B_{coh} = 5$ Hz, the overlap between adjacent bands $\chi = 50\%$ and we declare a band to be good if its energy is greater than $(1/10)^{th}$ of the band with

Signal	ρ_{L-to-R}	ρ_{SVD}
Chirp50	0.97	0.84
Chirp200	0.82	0.7
Sines50	0.87	0.72
Sines200	0.68	0.56

TABLE I
RESULTS OF L-TO-R PROCESSING AND SINGLE TAP APPROXIMATION OF THE RECORDED SIGNALS.

maximum energy i.e. $\eta = 0.1$. Note that there is always a “global” delay ambiguity in the estimate: we can delay the signal by τ and correspondingly *advance* all the channels by τ to explain the received data. Therefore, we quantify the performance of the L-to-R Stitching algorithm by the normalized crosscorrelation, denoted by ρ_{L-to-R} , between the source $\mathbf{x}$ and its estimate $\hat{\mathbf{x}}$, maximized over all possible time-shifts τ of the estimate:

$$\rho_{L-to-R} = \max_{\tau} \frac{(\mathcal{D}^{\tau} \hat{\mathbf{x}})^T \mathbf{x}}{\|\mathbf{x}\| \|\hat{\mathbf{x}}\|} \quad (14)$$

where $\mathcal{D}$ is the delay operator that shifts the vector $\mathbf{x}$ by one sample. We note that $0 \leq \rho_{L-to-R} \leq 1$ with $\rho_{L-to-R} = 1$ indicating a perfect match between $\mathbf{x}$ and $\hat{\mathbf{x}}$. We compare the L-to-R output against a solution that ignores multipath and approximates the channel to consist only of a single tap. Since the optimal estimate with this approximation is given by an SVD of the received signals [2], we refer to this estimate as the “SVD estimate” and denote its correlation with the truth by ρ_{SVD} .

We compare these estimates in Table I and make the following observations: (1) The L-to-R algorithm consistently performs better than the SVD estimate. Therefore, accounting for the multipath channel and piecing together estimates from different bands yields a better reconstruction even over small signal bandwidths (~ 50 Hz). (2) We observe that the reconstruction is very good when the signal bandwidth is small and worsens as the bandwidth increases. For example, the reconstruction is nearly perfect when the source is Chirp50, with $\rho_{L-to-R} = 0.97$. In contrast, when the bandwidth of the chirp signal is increased to 200 Hz, ρ_{L-to-R} drops to 0.82.

We illustrate the gains provided by the L-to-R stitching algorithm visually, by plotting the true Chirp200 waveform, the corresponding recorded signals at four sensors and the estimated signal in Figure 2. The true chirp waveform, in the topmost plot, has a constant envelope. Signals recorded at four sensors are shown in subsequent plots, and they clearly do not have a constant envelope. The reconstructed waveform, shown in the last plot, exhibits significantly smaller variations in its envelope and resembles the signal to a greater degree.

To understand the drop in correlation with increasing bandwidth, we correlate the reconstructed waveforms over different “subbands” of width 50 Hz each, with the i th subband spanning the frequencies $[1000 + 50(i - 1), 1000 + 50i]$ Hz. From Table II, we see that the fit over subbands of width 50 Hz is very good (the chirp signal is reconstructed nearly perfectly over each 50 Hz-band with $\rho_{L-to-R} \approx 0.98$). However, we




Fig. 2. The topmost plot shows the true Chirp200 waveform, with a constant envelope. The following four plots show the recorded waveforms at different sensors. Notice that these waveforms undergo “deep fades” and no longer have a constant envelope. The final plot shows the reconstructed Chirp200 waveform, whose envelope shows lesser variation, illustrating the benefits of the L-to-R algorithm.

find that the delay between the reconstructed signal (over 50 Hz-bands) and the true signal in these bands varies across bands (see Table III). This naturally leads to the conjecture that these delay variations across 50 Hz bands cause the signal contributions from these bands to combine “incoherently”, thereby affecting the quality of the reconstruction over a larger band. We have shown that such delay variations are indeed the cause of fundamental ambiguities, and can be used to systematically construct multiple explanations of the recorded data. Due to space limitations, details of such constructions are omitted from the paper.

Signal	Band 1	Band 2	Band 3	Band 4
Chirp200	0.982	0.952	0.983	0.987
Sines200	0.882	0.859	0.869	0.847

TABLE II

FIT BETWEEN SOURCE AND ESTIMATE IN BANDS OF WIDTH 50 HZ IS VERY GOOD. BAND i SPANS THE FREQUENCIES $[1000 + 50(i - 1), 1000 + 50i]$ HZ.

	Band 1	Band 2	Band 3	Band 4
Chirp200	287	203	123	172
Sines200	320	214	156	164

TABLE III

DELAY BETWEEN THE TRUE SOURCE AND THE ESTIMATE OVER BANDS OF WIDTH 50 HZ (IN SAMPLES @ $f_s = 16$ KHZ). BAND i SPANS THE FREQUENCIES $[1000 + 50(i - 1), 1000 + 50i]$ HZ.

B. Simulation Results

To quantify the performance of the algorithm statistically, we simulate a setting with $S = 6$ sensors where the channel to each sensor consists of 15 taps. The taps are placed uniformly within a delay spread of $\tau_{max} = 20$ ms ($B_{coh} \approx 5$ Hz) and the tap values are uniformly distributed between -2 and 2. We add independent Gaussian noise so that the SNR is 5 dB and

choose the overlap between adjacent bins to be $\chi = 70\%$. We compute the *average* and *minimum* values of the correlation between the estimate and the truth over 100 trials and display the results in Table IV. The average performance of the L-to-R algorithm is very good and the trends in the results agree well with experiments. Additionally, we observe that the worst-case performance of the L-to-R algorithm, given by $\rho_{min,L-to-R}$ falls faster (from 0.94 to 0.81 over the same range of bandwidths). Multiple explanations for the recorded data are the reason for such occasional glitches; details are omitted due to lack of space.

Signal	$\rho_{av,L-R}$	$\rho_{av,SVD}$	$\rho_{min,L-R}$	$\rho_{min,SVD}$
Chirp50	0.977	0.925	0.939	0.722
Chirp200	0.942	0.713	0.827	0.535
Chirp300	0.925	0.664	0.807	0.462

TABLE IV

PERFORMANCE OF THE L-TO-R STITCHING ALGORITHM AND THE SVD ESTIMATE WITH CHIRP AND “RANDOM” SIGNALS OF VARYING BANDWIDTHS.

V. CONCLUSIONS

We have shown, both through experiments and simulations, that a frequency domain approach is effective for collaborative estimation of an unknown signal observed in an unknown dispersive environment. The correlation between the estimated signal and the true signal is excellent over bandwidths that are 10-20 times the channel coherence bandwidth. However, fundamental ambiguities are possible over larger bandwidths by the introduction of small delay differences across bands, which result in multiple combinations of signal and channels that can explain a given set of observations (systematic construction of multiple explanations is possible, but has not been discussed here due to lack of space). It remains an open issue as to whether and how additional knowledge about the signal and/or channels can be best leveraged to alleviate these ambiguities.

REFERENCES

- [1] S. Seneff, “Pitch and spectral estimation of speech based on auditory synchrony model,” in *Proc. ICASSP 1984*.
- [2] S. Venkateswaran and U. Madhow, “Distributed detection with a minimalistic signal model: A framework for exploiting correlated sensing,” in *Proc. ISIT 2008*.
- [3] 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*, Aug 2002.
- [4] L. Tong and S. Perreau, “Multichannel blind identification : from subspace to maximum likelihood methods,” *Proceedings of the IEEE*, 1998.
- [5] H. Liu, G. Xu, L. Tong, and T. Kailath, “Recent developments in blind channel equalization: From cyclostationarity to subspaces,” *Signal Processing*, 1996.
- [6] L. Tong, G. Xu, and T. Kailath, “Blind identification and equalization based on second-order statistics: A time domain approach,” *Information Theory, IEEE Transactions on*, 2002.
- [7] Y. Li and Z. Ding, “Blind channel identification based on second order cyclostationary statistics,” in *Proc. ICASSP 1993*.
- [8] O. Shalvi and E. Weinstein, “New criteria for blind deconvolution of nonminimum phase systems (channels),” *Information Theory, IEEE Transactions on*, 2002.
- [9] G. Xu, H. Liu, L. Tong, and T. Kailath, “A least-squares approach to blind channel identification,” *Signal Processing, IEEE Transactions on*, 1995.