

Distributed Sampling of Signals Linked by Sparse Filtering: Theory and Applications

Ali Hormati, *Student Member, IEEE*, Olivier Roy, *Student Member, IEEE*, Yue M. Lu, *Member, IEEE*, and Martin Vetterli, *Fellow, IEEE*

Abstract—We study the distributed sampling and centralized reconstruction of two correlated signals, modeled as the input and output of an unknown sparse filtering operation. This is akin to a Slepian-Wolf setup, but in the sampling rather than the lossless compression case. Two different scenarios are considered: In the case of *universal reconstruction*, we look for a sensing and recovery mechanism that works for *all* possible signals, whereas in what we call *almost sure reconstruction*, we allow to have a small set (with measure zero) of unrecoverable signals. We derive achievability bounds on the number of samples needed for both scenarios. Our results show that, *only* in the almost sure setup can we effectively exploit the signal correlations to achieve effective gains in sampling efficiency. In addition to the above theoretical analysis, we propose an efficient and robust distributed sampling and reconstruction algorithm based on annihilating filters. We evaluate the performance of our method in one synthetic scenario, and two practical applications, including the distributed audio sampling in binaural hearing aids and the efficient estimation of room impulse responses. The numerical results confirm the effectiveness and robustness of the proposed algorithm in both synthetic and practical setups.

Index Terms—Annihilating filter, compressed sensing, compressive sampling, distributed sampling, finite rate of innovation, iterative denoising, sparse reconstruction, Yule-Walker system.

I. INTRODUCTION

CONSIDER two signals that are linked by an unknown filtering operation, where the filter is sparse in the time domain. Such models can be used, e.g., to describe the correlation between the transmitted and received signals in an unknown multi-path environment. We sample the two signals in a distributed setup: Each signal is observed by a different sensor, which sends a certain number of *non-adaptive* and *fixed* linear measurements of that signal to a central decoder. We study how

Manuscript received February 04, 2009; accepted September 05, 2009. First published October 20, 2009; current version published February 10, 2010. This work was supported by the Swiss National Science Foundation under Grants 5005-67322 (NCCR-MICS) and 200020-103729. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Chong-Meng Samson See.

A. Hormati, O. Roy, and Y. M. Lu are with the School of Computer and Communication Sciences, Ecole Polytechnique Fédérale de Lausanne (EPFL), CH-1015 Lausanne, Switzerland (e-mail: ali.hormati@epfl.ch; olivier.roy@epfl.ch; yue.lu@epfl.ch).

M. Vetterli is with the School of Computer and Communication Sciences, Ecole Polytechnique Fédérale de Lausanne (EPFL), CH-1015 Lausanne, Switzerland. He is also with the Department of Electrical Engineering and Computer Sciences, University of California, Berkeley, CA 94720 USA (e-mail: martin.vetterli@epfl.ch).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TSP.2009.2034908

the correlation induced by the above model can be exploited to reduce the number of measurements needed for perfect reconstruction at the central decoder, but *without* any inter-sensor communication during the sampling process.

Our setup is conceptually similar to the Slepian-Wolf problem in distributed source coding [1]–[3], which consists of correlated sources to be encoded separately and decoded jointly. While communication between the encoders is precluded, correlation between the measured data can be taken into account as an effective means to reduce the amount of information transmitted to the decoder. The main difference between our work and this classical distributed source coding setup is that we study a *sampling* problem and hence are only concerned about the number of sampling measurements we need to take, whereas the latter is about *coding* and hence uses bits as its “currency”. From the sampling perspective, our work is closely related to the problem of distributed compressed sensing, first introduced in [4] (see also [5]–[10]). In that framework, jointly sparse data need to be reconstructed based on linear projections computed by distributed sensors. In [4], the authors proposed three joint-sparsity models for distributed signals, as well as efficient algorithms for signal recovery.

The first contribution of this paper is a novel correlation model for distributed signals. Instead of imposing any sparsity assumption on the signals themselves (as in [4]), we assume that the signals are linked by some unknown sparse filtering operation. Such models can be useful in describing the signal correlation in several practical scenarios (e.g., binaural audio recording). Under the sparse-filtering model, we introduce two strategies for the design of the sampling system: In the *universal* strategy, we seek to successfully recover *all* signals, whereas in what we call the *almost sure* strategy, we allow to have a small set (with measure zero) of unrecoverable signals. As the second contribution of our work, we establish the corresponding achievability bounds on the number of samples needed for the two strategies mentioned above. These bounds indicate that the sparsity of the filter can be useful only in the almost sure strategy.

Since the algorithms that achieve the aforementioned bounds are computationally prohibitive, we introduce, as our third contribution, a novel distributed sampling and reconstruction scheme that can recover the original signals in an efficient and robust way. As an intermediate step in the reconstruction process, our proposed algorithm employs the annihilating filter technique [11]–[13] to estimate the unknown sparse channel between the two signals. Several recent papers on channel estimation (e.g., [14] and [15]) have also explored the sparsity

of the channels, usually by using theories and techniques developed in compressed sensing. In this perspective, our proposed reconstruction algorithm can be viewed as a new approach to sparse channel estimation, based on annihilating filters. However, we emphasize that, rather than focusing on channel estimation, our main goal is to develop efficient distributed sampling schemes based on the signal correlation provided by the sparse filtering model.

The rest of this paper is organized as follows. After a precise definition of the considered model, we state in Section II a general formulation of the distributed sensing problem. In Section III-A, we demonstrate a somewhat surprising result: If one requires *all* possible signals to be reconstructed perfectly, then the aforementioned correlation between the observed signals cannot be exploited. In that case, the simple strategy of independently sending all coefficients of the distributed signals would be optimal. However, if one only considers the perfect recovery of *almost all* signals, then substantial gains in sampling efficiency can indeed be achieved. We derive an achievability bound for the almost sure reconstruction in Section III-B. Since the algorithm that attains the bound has combinatorial complexity, we propose in Section IV a slightly suboptimal, yet computationally efficient distributed algorithm based on annihilating filters. Moreover, we show how the proposed method can be made robust to model mismatch using an iterative procedure due to Cadzow [16]. In Section V, we discuss several possible extensions and generalizations of the modeling and recovery algorithms proposed in this paper. Finally, Section VI presents numerical experiments to illustrate the performance of our proposed scheme in both synthetic and practical scenarios. We conclude the paper in Section VII.

II. SIGNAL MODEL AND PROBLEM STATEMENT

A. Proposed Correlation Model

Consider two signals $x_1(t)$ and $x_2(t)$, where $x_2(t)$ can be obtained as a filtered version of $x_1(t)$. In particular, we assume that

$$x_2(t) = (x_1 * h)(t) \quad (1)$$

where $h(t) = \sum_{k=1}^K c_k \delta(t - t_k)$ is a stream of K Diracs with *unknown* delays $\{t_k\}_{k=1}^K$ and coefficients $\{c_k\}_{k=1}^K$. The above model characterizes the correlation between a pair of signals of interest in various practical applications. Examples include the correlation between transmitted and received signals under multipath propagation, with $h(t)$ representing the unknown channel, or the spatial correlation between signals recorded by two closely spaced microphones in a simple acoustic environment composed of a single source.

In this paper, we study a finite-dimensional discrete version of the above model. As shown in Fig. 1, we assume that the original continuous signal $x_1(t)$ is bandlimited to $[-\sigma, \sigma]$. Sampling $x_1(t)$ at uniform time interval T leads to a discrete sequence of samples $x_{s1}[n] \stackrel{\text{def}}{=} x_1(nT)$, where the sampling rate $1/T$ is set to be above the Nyquist rate σ/π . To obtain a finite-length

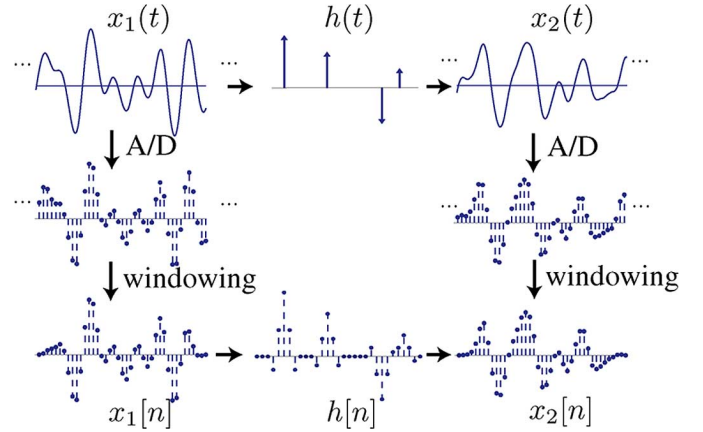


Fig. 1. The continuous-time sparse filtering operation and its discrete-time counterpart. The two bandlimited signals $x_1(t)$ and $x_2(t)$ are sampled above the Nyquist rate, followed by a smooth temporal windowing operation. Neglecting the windowing effect, we can approximately model the resulting finite-length signals $x_1[n]$ and $x_2[n]$ as the input and output of a discrete-time filtering operation. The discrete filter $h[n]$ is sparse as long as the sampling interval is fine enough.

signal, we subsequently apply a temporal window to the infinite sequence $x_{s1}[n]$ and get

$$x_1[n] \stackrel{\text{def}}{=} x_{s1}[n] w_N[n], \quad \text{for } n = 0, 1, \dots, N-1,$$

where $w_N[n]$ is a smooth temporal window (e.g., the Kaiser window [17]) of length N . It is easy to verify that the discrete Fourier transform (DFT) of the finite sequence $x_1[n]$ is

$$X_1[m] = \frac{1}{2\pi} (X_{s1} \circledast W_N) \left(e^{j \frac{2\pi m}{N}} \right), \quad (2)$$

for $m = \lfloor -N/2 \rfloor + 1, \dots, \lfloor N/2 \rfloor$, where $X_{s1}(e^{j\omega})$ and $W_N(e^{j\omega})$ are the discrete-time Fourier transforms of x_{s1} and w_N , respectively, and \circledast represents circular convolution. When N is large enough, we can omit the windowing effect, since $W_N(e^{j\omega})/(2\pi)$ approaches a Dirac function $\delta(\omega)$ as $N \rightarrow \infty$. It then follows from (2) that

$$X_1[m] \approx X_{s1} \left(e^{j \frac{2\pi m}{N}} \right) = \frac{1}{T} X_1 \left(\frac{2\pi m}{NT} \right)$$

where $X_1(\omega)$ is the continuous-time Fourier transform of $x_1(t)$, and the equality above is due to the standard sampling formula in the Fourier domain [18].

Applying the same procedure to $x_2(t)$ and using (1), we have

$$X_2[m] \approx \frac{1}{T} X_2 \left(\frac{2\pi m}{NT} \right) \approx X_1[m] H[m]$$

where

$$H[m] \stackrel{\text{def}}{=} \sum_{k=1}^K c_k e^{-j2\pi m t_k / (NT)}. \quad (3)$$

This relationship implies that, like the original continuous signals $x_1(t)$ and $x_2(t)$, the finite-length signals $x_1[n]$ and $x_2[n]$ can also be approximately modeled as the input and output of a *discrete-time* filtering operation, where the unknown filter




Fig. 2. The proposed distributed sensing setup. The i th sensor ($i = 1, 2$) provides an M_i -dimensional observation of the signal \mathbf{x}_i via a nonadaptive linear transform \mathbf{A}_i . The central decoder reconstructs the original vectors (as well as the sparse convolution kernel \mathbf{h}) based on the received measurements.

$H[n]$ —and in time domain $h[n]$ —contains all the information about the original continuous filter $h(t)$ ¹. In general, the location parameters $\{t_k\}$ in (3) can be arbitrary real numbers, and consequently, the discrete-time filter $h[n]$ is no longer sparse (see Fig. 1 for a typical impulse response of $h[n]$). However, when the sampling interval T is small enough, we can assume that the real-valued delays $\{t_k\}$ are close enough to the sampling grid, i.e., $t_k/T \approx n_k$ for some integers $\{n_k\}$. Under this assumption, the filter $h[n]$ becomes a sparse vector with K nonzero elements, represented by

$$h[n] = \sum_{k=1}^K c_k \delta[n - n_k].$$

We will follow this assumption² throughout the paper, and focus on the following model.

Definition 1 (Correlation Model): The signals of interest are two vectors $\mathbf{x}_1 = (x_1[0], \dots, x_1[N-1])^T$ and $\mathbf{x}_2 = (x_2[0], \dots, x_2[N-1])^T$, linked to each other through a circular convolution

$$x_2[n] = (x_1 \otimes h)[n] \quad \text{for } n = 0, 1, \dots, N-1 \quad (4)$$

where $\mathbf{h} = (h[0], \dots, h[N-1])^T \in \mathbb{R}^N$ is an unknown K -sparse vector, that is $\|\mathbf{h}\|_0 = K$.

Definition 2: The signal space \mathcal{X} is the set of all stacked vectors $\mathbf{x}^T \stackrel{\text{def}}{=} (\mathbf{x}_1^T, \mathbf{x}_2^T) \in \mathbb{R}^{2N}$ such that its components $\mathbf{x}_1, \mathbf{x}_2$ satisfy the correlation model given in Definition 1.

B. Distributed Sensing and Problem Statement

We consider the problem of sensing $\mathbf{x}^T = (\mathbf{x}_1^T, \mathbf{x}_2^T)$ in a *distributed* fashion, by two independent sensors taking linear measurements of \mathbf{x}_1 and \mathbf{x}_2 , respectively. The measurement matrices are *fixed* (i.e., they do not change with the input signals). As shown in Fig. 2, suppose that the i th sensor ($i = 1, 2$) takes M_i linear measurements of \mathbf{x}_i . We can write

$$\mathbf{y}_i = \mathbf{A}_i \mathbf{x}_i,$$

¹Note that in order to be unambiguous in the positions $\{t_k\}$, we need to ensure that $NT > \max_k \{t_k\}$.

²We introduce this assumption (i.e., $t_k/T = n_k$ for some $n_k \in \mathbb{Z}$) mainly for the simplicity it brings to the theoretical analysis in later parts of this paper. It is however not an inherent limitation of our work. In fact, the proposed sampling and recovery algorithm presented in Section IV can work with the case when the delays $\{t_k\}_{k=1}^K$ have arbitrary real values. We demonstrate this capability in Section VI-B, where we apply the proposed algorithm in estimating acoustic room impulse responses.

where $\mathbf{y}_i \in \mathbb{R}^{M_i}$ represents the vector of samples taken by the i th sensor, and \mathbf{A}_i is the corresponding sampling matrix. Considering the stacked vector $\mathbf{y}^T \stackrel{\text{def}}{=} (\mathbf{y}_1^T, \mathbf{y}_2^T)$, we have $\mathbf{y} = \mathbf{A}\mathbf{x}$, where

$$\mathbf{A} = \begin{bmatrix} \mathbf{A}_1 & \mathbf{0}_{M_1 \times N} \\ \mathbf{0}_{M_2 \times N} & \mathbf{A}_2 \end{bmatrix}. \quad (5)$$

Note that the block-diagonal structure of \mathbf{A} is due to the fact that \mathbf{x}_1 and \mathbf{x}_2 are sensed separately. This is in contrast to the *centralized* scenario in which \mathbf{x}_1 and \mathbf{x}_2 can be processed jointly and hence, the matrix \mathbf{A} can be arbitrary.

The measurements \mathbf{y}_1 and \mathbf{y}_2 are transmitted to a central decoder, which attempts to reconstruct the vector \mathbf{x} through some (possibly nonlinear) mapping $\psi: \mathbb{R}^{M_1+M_2} \mapsto \mathcal{X}$ as

$$\hat{\mathbf{x}} = \psi(\mathbf{y}).$$

By analogy to the Slepian-Wolf problem in distributed source coding [1], the natural questions to ask in the above sampling setup are the following:

- 1) What choices of sampling pairs (M_1, M_2) will allow us to reconstruct signals $\mathbf{x} \in \mathcal{X}$ from their samples?
- 2) What is the loss incurred by the distributed infrastructure in (5) over the centralized scenario in terms of the total number of measurements $M_1 + M_2$?
- 3) How to reconstruct the original signals from their samples in a computationally efficient way?

In what follows, we first answer the above questions in the case of *universal* reconstruction (Section III-A), where we want to recover *all* signals. We then consider *almost sure* reconstruction (Section III-B) in which case we allow to have a small set (with measure zero) of unrecoverable signals. In Section IV, we propose a robust and computationally efficient reconstruction algorithm based on annihilating filters.

III. BOUNDS ON ACHIEVABLE SAMPLING PAIRS

A. Universal Recovery

Let \mathbf{A}_1 and \mathbf{A}_2 be the sampling matrices used by the two sensors, and \mathbf{A} be the corresponding block-diagonal matrix as defined in (5). We first focus on finding those \mathbf{A}_1 and \mathbf{A}_2 such that every $\mathbf{x} \in \mathcal{X}$ is uniquely determined by its samples $\mathbf{A}\mathbf{x}$.

Definition 3 (Universal Achievability): We say a sampling pair (M_1, M_2) is achievable for universal reconstruction if there exist *fixed* measurement matrices $\mathbf{A}_1 \in \mathbb{R}^{M_1 \times N}$ and $\mathbf{A}_2 \in \mathbb{R}^{M_2 \times N}$ such that the set

$$\mathcal{B}(\mathbf{A}_1, \mathbf{A}_2) \stackrel{\text{def}}{=} \{\mathbf{x} \in \mathcal{X} : \exists \mathbf{x}' \in \mathcal{X} \text{ with } \mathbf{x} \neq \mathbf{x}' \text{ but } \mathbf{A}\mathbf{x} = \mathbf{A}\mathbf{x}'\} \quad (6)$$

is empty.

Intuition suggests that, due to the correlation between the vectors \mathbf{x}_1 and \mathbf{x}_2 , the minimum number of samples needed to perfectly describe all possible vectors \mathbf{x} can be made smaller than the total number of coefficients $2N$. The following proposition shows that, surprisingly, this is not the case.

Proposition 1: Let $\mathbf{x}_1 \in \mathbb{R}^N$ and $\mathbf{x}_2 \in \mathbb{R}^N$ be two signals that are related through an unknown K -sparse convolution kernel \mathbf{h} ,

and sampled separately. A sampling pair (M_1, M_2) is achievable for universal reconstruction if and only if $M_1 \geq N$ and $M_2 \geq N$.

Proof: The sufficiency of the above conditions is clear. To show its necessity, let us consider two stacked vectors $\mathbf{x}^T = (\mathbf{x}_1^T, \mathbf{x}_2^T)$ and $\mathbf{x}'^T = (\mathbf{x}'_1^T, \mathbf{x}'_2^T)$, each following the correlation model (4). They can be written under the form

$$\mathbf{x} = \begin{bmatrix} \mathbf{I}_N \\ \mathbf{C} \end{bmatrix} \mathbf{x}_1 \quad \text{and} \quad \mathbf{x}' = \begin{bmatrix} \mathbf{I}_N \\ \mathbf{C}' \end{bmatrix} \mathbf{x}'_1$$

where \mathbf{C} and \mathbf{C}' are two circulant matrices with vectors \mathbf{h} and \mathbf{h}' as the first column, respectively. It holds that

$$\mathbf{x} - \mathbf{x}' = \begin{bmatrix} \mathbf{I}_N & -\mathbf{I}_N \\ \mathbf{C} & -\mathbf{C}' \end{bmatrix} \begin{bmatrix} \mathbf{x}_1 \\ \mathbf{x}'_1 \end{bmatrix}. \quad (7)$$

Meanwhile, we can check that

$$\begin{aligned} \text{rank} \begin{bmatrix} \mathbf{I}_N & -\mathbf{I}_N \\ \mathbf{C} & -\mathbf{C}' \end{bmatrix} &= \text{rank} \begin{bmatrix} \mathbf{I}_N & \mathbf{0} \\ \mathbf{C} & \mathbf{C} - \mathbf{C}' \end{bmatrix} \\ &= N + \text{rank}(\mathbf{C} - \mathbf{C}'). \end{aligned} \quad (8)$$

Next, we show that for any K -sparse filter \mathbf{h} , we can always find another K -sparse filter \mathbf{h}' such that $\mathbf{C} - \mathbf{C}'$ is full rank. To see this, note that \mathbf{C} and \mathbf{C}' are circulant matrices, therefore they can be diagonalized by DFT matrices where the diagonal elements are the DFTs of the K -sparse filters \mathbf{h} and \mathbf{h}' , respectively. For any K -sparse filter \mathbf{h} , we can build a different K -sparse filter $h'[n] = h[n] + \alpha \delta[n - n_0]$ where n_0 is the position of any nonzero element of \mathbf{h} and $\alpha \neq -h[n_0]$. Since $\delta[n - n_0]$ has a flat spectrum, it insures that $\mathbf{C} - \mathbf{C}'$ is full rank. It follows from (8) that the matrix in (7) is of full rank too. Consequently, $\mathbf{x} - \mathbf{x}'$ in (7) can take any possible value in \mathbb{R}^{2N} . Now if the matrix \mathbf{A} is not full rank, we can always find two different vectors \mathbf{x} and \mathbf{x}' , such that their difference $\mathbf{x} - \mathbf{x}'$ is in the null space of \mathbf{A} , which implies that \mathbf{x} and \mathbf{x}' provide the same measurement vector. Hence, a necessary condition for the set (6) to be empty is that the block-diagonal matrix \mathbf{A} be a full rank matrix of size $M \times 2N$, with $M \geq 2N$. In particular, \mathbf{A}_1 and \mathbf{A}_2 must be full rank matrices of size $M_1 \times N$ and $M_2 \times N$, respectively, with $M_1, M_2 \geq N$. Note that, in the centralized scenario, the full rank condition would still require to take at least $2N$ measurements. ■

Remark: As a direct consequence of the above result for universal reconstruction, each sensor can simply process its signal independently without any loss of optimality. In particular, the simple strategy of sending all observed coefficients is optimal. Moreover, it is seen in the proof of Proposition 1 that there is no penalty associated with the distributed nature of the sampling setup. In other words, the total number of measurements cannot be made smaller than $2N$ even if the vectors \mathbf{x}_1 and \mathbf{x}_2 can be sampled jointly. The region of achievable sampling pairs for universal reconstruction is depicted as the shaded area in Fig. 3.

B. Almost Sure Recovery

As shown in Proposition 1, universal recovery is a rather strong requirement to satisfy since we have to take at least N




Fig. 3. Shaded area: achievable sampling region for universal reconstruction. Solid line: boundary of the sampling pairs achieved for almost sure reconstruction for K odd [any pair (M_1, M_2) above the line is achievable]. Dashed line: boundary of the sampling pairs achieved for almost sure reconstruction by the proposed algorithm based on annihilating filters (see Section IV for details).

samples at each sensor, with no chance to exploit the existing signal correlation. In many situations, however, it is sufficient to consider a weaker requirement, which aims at finding measurement matrices that permit the perfect recovery of *almost all* signals from \mathcal{X} . We will use the following definition in our discussion on the almost sure reconstruction.

Definition 4 (Nonsingular Probability Distribution): We say a probability distribution \mathcal{P} over \mathbb{R}^N is nonsingular if for any subset $\mathcal{S} \in \mathbb{R}^N$ with Lebesgue measure zero we have $\mathcal{P}(\mathcal{S}) = 0$.

For nonsingular distributions, the probability of signals living in a subspace with dimension less than N is zero. A typical example of a nonsingular distribution is a jointly Gaussian distribution with a nonsingular covariance matrix.

Definition 5 (Almost Sure Achievability): We say a sampling pair (M_1, M_2) is achievable for almost sure reconstruction if there exist *fixed* measurement matrices $\mathbf{A}_1 \in \mathbb{R}^{M_1 \times N}$ and $\mathbf{A}_2 \in \mathbb{R}^{M_2 \times N}$ such that the set $\mathcal{B}(\mathbf{A}_1, \mathbf{A}_2)$, as defined in (6), is of probability zero.

The above definition for almost sure recovery depends on the probability distributions of the signal \mathbf{x}_1 and the sparse filter \mathbf{h} . In our discussions in Section IV and Appendix A, we assume that the signal \mathbf{x}_1 follows a probability distribution such that the frequency components of \mathbf{x}_1 are nonzero with probability 1. Note that this condition is fairly mild and can be satisfied when \mathbf{x}_1 is drawn from any non-singular probability distribution over \mathbb{R}^N , or when \mathbf{x}_1 is sparse in a basis that is different from the standard Fourier basis. For the sparse filter \mathbf{h} , it is sufficient to assume that, given the locations of its nonzero coefficients, the values of these coefficients are drawn from a non-singular distribution over \mathbb{R}^K . The following proposition gives an achievability bound on the number of samples needed for almost sure reconstruction.

Proposition 2: Let $\mathbf{x}_1 \in \mathbb{R}^N$ and $\mathbf{x}_2 \in \mathbb{R}^N$ be two signals that are related through an unknown K -sparse convolution kernel \mathbf{h} ,

and sampled separately. A sampling pair (M_1, M_2) is achievable for almost sure reconstruction if

$$\begin{aligned} M_1 &\geq \min \{K + r, N\} \\ M_2 &\geq \min \{K + r, N\} \\ \text{and } M_1 + M_2 &\geq \min \{N + K + r, 2N\} \end{aligned} \quad (9)$$

where $r = 1 + \text{mod}(K, 2)$.

Proof: See Appendix A. ■

Remark: Proposition 2 shows that, in contrast to the universal scenario, the correlation between the signals by means of the sparse filtering provides a large improvement in sampling efficiency in the almost sure setup, especially when $K \ll N$. We show the boundary of the above achievable region as the solid line in Fig. 3.

IV. ALMOST SURE RECONSTRUCTION BASED ON ANNIHILATING FILTERS

We show in the proof of Proposition 2 (see Appendix A) that the algorithm which attains the bound in (9) is combinatorial in nature and, thus, computationally prohibitive. In the following, we propose a novel distributed sensing algorithm based on annihilating filters. This method, known as Prony's method in spectral estimation, belongs to the class of model-based parametric methods for high-resolution harmonic retrieval [11], [13], [19]. This algorithm needs effectively K more measurements with respect to the achievability region for the almost sure reconstruction but can substantially reduce the reconstruction complexity to $O(KN)$. We start our discussion with the noiseless case.

A. Noiseless Scenario

The proposed distributed sensing scheme is based on a frequency-domain representation of the input signals. Let us denote by $\mathbf{X}_1 \in \mathbb{C}^N$ and $\mathbf{X}_2 \in \mathbb{C}^N$ the DFTs of the vectors \mathbf{x}_1 and \mathbf{x}_2 , respectively. The circular convolution in (4) can be expressed as

$$\mathbf{X}_2 = \mathbf{H} \odot \mathbf{X}_1 \quad (10)$$

where $\mathbf{H} \in \mathbb{C}^N$ is the DFT of the filter \mathbf{h} and \odot denotes element-wise multiplication. Our approach consists of the following two main steps:

- 1) Finding the filter \mathbf{h} by sending the first $K + 1$ (1 real and K complex) DFT coefficients of \mathbf{x}_1 and \mathbf{x}_2 .
- 2) Sending the remaining frequency indices by sharing them among the two sensors.

We first show how a decoder can almost surely recover the unknown filter \mathbf{h} using only the first $K + 1$ DFT coefficients of \mathbf{x}_1 and \mathbf{x}_2 . This is achieved by using the annihilating filter approach, which is widely used in harmonic retrieval applications. In what follows, we focus on the main steps of the algorithm; more details can be found in [11] and [13].

The DFT coefficients of the filter \mathbf{h} are given by

$$H[m] = \sum_{k=1}^K c_k e^{-j \frac{2\pi}{N} n_k m} \quad \text{for } m = 0, 1, \dots, N - 1. \quad (11)$$

The sequence $H[m]$ is the sum of K complex exponentials, whose frequencies are determined by the positions $\{n_k\}$ of the

nonzero coefficients of the filter. It can be shown [11] that $H[m]$ can be "annihilated" by a filter $A(z)$ of degree K whose roots are of the form $e^{-j2\pi n_k/N}$ for $k = 1, 2, \dots, K$, i.e.,

$$A(z) = \prod_{k=1}^K (1 - e^{-j2\pi n_k/N} z^{-1}).$$

More specifically, in the spatial domain, the coefficients of this filter satisfy

$$(A * H)[m] = \sum_{i=0}^K A[i] H[m - i] = 0$$

or in matrix form

$$\begin{bmatrix} H[0] & H[-1] & \cdots & H[-K] \\ H[1] & H[0] & \cdots & H[-K+1] \\ \vdots & \vdots & \ddots & \vdots \\ H[K-1] & H[K-2] & \cdots & H[-1] \end{bmatrix} \begin{bmatrix} A[0] \\ A[1] \\ \vdots \\ A[K] \end{bmatrix} = 0. \quad (12)$$

The above matrix is of size $K \times (K + 1)$ and is built from $2K$ consecutive DFT coefficients. Moreover, it can be shown to be of rank K (see Appendix B) and hence, its null space is of dimension one. Therefore, the solution can be any vector in the null-space of the above matrix. Note that, due to the conjugate symmetry property, the coefficients of the matrix in (12) can be computed as

$$H[m] = \frac{X_2[m]}{X_1[m]} \quad \text{and} \quad H[-m] = H^*[m] \quad (13)$$

provided that $X_1[m]$ is nonzero for $m = 0, 1, \dots, K$. In other words, those signals \mathbf{x}_1 which have zero frequency components in the range $m = 0, 1, \dots, K$ are unrecoverable by the annihilating filter method. However, according to our assumption on the probability distribution of \mathbf{x}_1 , the probability of encountering such signals is zero. Once the coefficients of the annihilating filter have been obtained, it is simply a matter of computing its roots to retrieve the unknown positions $\{n_k\}$. The filter weights $\{c_k\}$ can then be recovered by means of the linear system of equations in (11).

Based on the above considerations, our distributed sensing scheme can be described as follows. Both sensors send the first³ $K + 1$ DFT coefficients of their signals to the decoder ($2K + 1$ real values each). They also transmit complementary subsets⁴ (in terms of frequency indexes) of the remaining DFT coefficients ($N - 2K - 1$ real values in total). This is illustrated in Fig. 4. Note that certain level of coordination, i.e., which Fourier coefficients to retain, is needed among the sensors. However, this type of coordination is fixed and can be set before deployment. The first $K + 1$ DFT coefficients allow us to almost surely reconstruct the filter \mathbf{h} . The missing frequency components of \mathbf{x}_1 (resp. \mathbf{x}_2) are then recovered from the available DFT coefficients of \mathbf{x}_2 (respectively, \mathbf{x}_1) using (10). Note that in order to

³Note that we could also use a pass-band frequency range but in that case, we need $2K$ complex measurements ($4K$ real), since we can no longer use the complex conjugate property of the Fourier transform to build $2K$ consecutive frequency data from $K + 1$ measurements.

⁴Two subsets A and B are called the complementary subsets of C if $A \cap B = \emptyset$ and $A \cup B = C$.

compute $X_1[m]$ from $X_2[m]$, the frequency component $H[m]$ should be nonzero. This is insured almost surely under our assumption that the nonzero elements of the filter \mathbf{h} are drawn from a nonsingular distribution in \mathbb{R}^K . We summarize the proposed method in Algorithm 1. In terms of achievability, we have thus shown the following result.

Algorithm 1: Sensing and Recovery Based on Annihilating Filters

- 1: Sensors 1 and 2 send the first $K + 1$ DFT coefficients of \mathbf{x}_1 and \mathbf{x}_2 , respectively. They also send complementary subsets of the remaining DFT coefficients.
 - 2: The decoder computes $2K$ consecutive DFT coefficients of \mathbf{h} from (13).
 - 3: The decoder retrieves the filter \mathbf{h} using the annihilating filter method.
 - 4: The decoder reconstructs \mathbf{x}_1 and \mathbf{x}_2 from (10) using \mathbf{h} and the remaining DFT coefficients of \mathbf{x}_1 and \mathbf{x}_2 .
-

Proposition 3: A sampling pair (M_1, M_2) is achievable for almost sure reconstruction using the proposed annihilating filter method if

$$\begin{aligned} M_1 &\geq \min \{2K + 1, N\} \\ M_2 &\geq \min \{2K + 1, N\} \\ \text{and } M_1 + M_2 &\geq \min \{N + 2K + 1, 2N\}. \end{aligned}$$

In contrast with the case of universal reconstruction, the total number of measurements can be reduced from $2N$ to $N + 2K + 1$, as depicted by the dashed line in Fig. 3. Although slightly oversampled compared to the bound for the almost sure setup given in Proposition 2, we obtain substantial improvement in terms of computational efficiency, as seen in the following.

The annihilating filter equation in (12) is a Yule-Walker system which is solvable in $O(K^2)$ operations [20]. To find the locations of the nonzero coefficients, the annihilating filter should be factorized. The roots of the filter can be found by the Chien search algorithm in $O(KN)$ operations [21]. In the last step, the weights of the nonzero coefficients of the filter are obtained by solving a Vandermonde system, which requires $O(K^2)$ operations [20]. Since $K \ll N$, the total number of operations needed in reconstructing the sparse filter is $O(KN)$.

B. Noisy Case

Noise or, more generally, model mismatch makes it difficult to directly apply the solution discussed above in practice. Adding robustness to the system requires sending additional measurements to the decoder.

Our first strategy for reconstruction in the noisy scenario is to use the total least squares (TLS) approach [19] to solve the system of equations in (12). In the TLS technique, it is assumed that the observation errors are on both sides of a system of equations, instead of just on the right hand side as in the least squares approach. This approach finds the solution by minimally perturbing both sides of a system of equations, until it is satisfied.

The TLS approach requires the singular value decomposition and the solution is estimated from the singular vectors [20]. In our system of equations in (12), the TLS approach consists of computing the singular value decomposition of the coefficient matrix and then, the solution is the singular vector which corresponds to the smallest singular value.

Algorithm 2: Cadzow Iterative Denoising

- 1: Build matrix \mathbf{T} of dimension $L \times (L + 1)$ of the form (12) from the measurements.
 - 2: Set ϵ to be a small constant.
 - 3: **while** $(\sigma_{K+1})/(\sigma_K) \geq \epsilon$ **do**
 - 4: Enforce rank K on \mathbf{T} by setting the $L - K + 1$ smallest singular values to zero.
 - 5: Enforce the Toeplitz form on \mathbf{T} by averaging the coefficients along the diagonals.
 - 6: **end while**
 - 7: Extract the denoised DFT coefficients from the first row and first column of \mathbf{T} .
-

To further improve robustness, we use an iterative method devised by Cadzow [16]. In our context, it can be summarized as follows. Sensor i transmits the first $L + 1$ DFT coefficients of x_i ($i = 1, 2$) with $L \geq K$. A matrix \mathbf{T} of dimension $L \times (L + 1)$ of the form (12) is then built from these measurements. In the noiseless case, this matrix has two key properties: (i) it has rank K (see Appendix B) and (ii) it is Toeplitz. In the noisy case, these two properties are not initially satisfied simultaneously, but can be subsequently enforced by alternatively performing the following two steps:

- i) Enforce rank K by setting the $L - K + 1$ smallest singular values of \mathbf{T} to zero.
- ii) Enforce the Toeplitz form on \mathbf{T} by averaging the coefficients along the diagonals.

The above procedure is guaranteed to converge to a matrix which exhibits the desired properties [16]. The iterations stop whenever the ratio of the $(K + 1)$ th singular value to the K th one, σ_{K+1}/σ_K , falls below a predetermined threshold (e.g., 10^{-6}). The denoised DFT coefficients are then extracted from the first row and first column, and used to solve the systems (11) and (12) using the TLS method. The method is summarized in Algorithm 2.

It may happen that the sparsity parameter K is not known *a priori* or the sparse filter is only approximately sparse, with fast decaying entries. In these cases, it is sufficient to estimate an upperbound \hat{K} for the effective sparsity parameter⁵. Then, the Cadzow denoising algorithm treats the remaining $N - \hat{K}$ least significant coefficients of the filter $h[n]$ as noise. In this case, the denoising performance depends on the energy of the \hat{K} significant coefficients with respect to the $N - \hat{K}$ nonsignificant ones.

⁵By effective sparsity parameter we mean the number of significant coefficients that retain most of the energy of the original filter $h[n]$. This is a model estimation problem which in general requires some *a priori* knowledge about the underlying structure which generates the sparse filter.




Fig. 4. Sensors 1 and 2 both send the first $K + 1$ DFT coefficients of their observations, but only complementary subsets of the remaining frequency components.

V. POSSIBLE EXTENSIONS

A. Generalizing the Model

In this paper, we consider the correlation model which consists of sparse filters in the time domain. Using the sampling theory for signals with finite rate of innovation [12], the proposed model can be extended to cases where the unknown filters are *piecewise polynomials* or *piecewise bandlimited* functions. Moreover, the model is extendable to filters which admit a sparse representation in an arbitrary basis. In that case, one can send samples which involve random frequency measurements [5] of the sparse filter. Then, the reconstruction uses techniques based on ℓ_1 minimization⁶ to recover the filter, with a complexity of $O(N^3)$. After finding the filter, the two sensors send complementary information on the remaining frequency indices of their signals to the decoder. Finally, another possible extension of the model is to consider general linear transforms (matrices) which have a sparse representation in a matrix dictionary [22].

B. The Multiframe Setup

Our discussion in this paper focuses on distributed sensing in a single frame setup, where only one “snapshot” is available for sensing and reconstruction. It is possible to consider the case where multiple consecutive frames are available and the sparse filters in different frames are inter-related through certain joint sparse models [4]. In this scenario, one can exploit the interrelation between the underlying K -sparse filters in different frames in the sensing and recovery architecture. This allows to either reduce the total number of measurements or, make the reconstruction more robust to noise and model mismatch. In [23], we showed how to efficiently exploit the interrelation between the sparse signals in different frames in the annihilating filter based reconstruction method. In the case of random frequency measurements followed by the ℓ_1 -based recovery method, the techniques proposed in [4] can be used to exploit redundancy across frames.

VI. NUMERICAL EXPERIMENTS AND APPLICATIONS

In this section, we present numerical results to demonstrate the effectiveness and robustness of the proposed sensing and recovery scheme based on annihilating filters. The simulations are divided into three parts. First, in Section VI-A, we run simulations on synthetic data with additive white Gaussian noise added to the sparse filter to evaluate the effect of model mismatch on the recovery algorithm. Then, in Section VI-B, we apply our proposed algorithm to estimate an acoustic room impulse response (RIR) generated by the image-source model [24], [25].

⁶One can also use greedy algorithms like orthogonal matching pursuit [7] for reconstruction.

Finally, in Section VI-C, we evaluate the performance of the algorithm in a distributed audio coding application, where the goal is to localize a sound source using two hearing aid microphones.

A. Synthetic Experiments

Assume that the signal \mathbf{x}_1 is of length $N = 256$ and the sparse filter has $K = 3$ or 6 nonzero coefficients. The elements of the signal \mathbf{x}_1 and the nonzero coefficients of the filter are chosen to be i.i.d. random with distribution $\mathcal{N}(0, 1)$. The positions of the nonzero coefficients of \mathbf{h} are chosen uniformly from $[1, \dots, N]$ at random.

To evaluate the effect of model mismatch on the filter model, we add independent white Gaussian noise to all N coefficients of the filter \mathbf{h} and get

$$\mathbf{h}^n = \mathbf{h} + \mathbf{n}, \quad \mathbf{n} \sim \mathcal{N}(0, \sigma^2 I_{N \times N}).$$

Then, the signal observed by the second sensor is given by

$$\mathbf{x}_2^n = \mathbf{x}_1 * \mathbf{h}^n.$$

We define

$$\text{SNR}_h = \frac{\|\mathbf{h}\|_2}{\|\mathbf{h}^n - \mathbf{h}\|_2}.$$

In this simulation, we let the first sensor send the whole signal \mathbf{x}_1 , and let the second sensor transmit the first $L + 1$ DFT coefficients of \mathbf{x}_2^n to the decoder, for different values of L . The decoder first estimates the true unknown K -sparse filter \mathbf{h} , call it $\hat{\mathbf{h}}$, and then computes an estimate of the signal of the second sensor as

$$\hat{\mathbf{x}}_2 = \mathbf{x}_1 * \hat{\mathbf{h}}.$$

To evaluate the quality of the reconstruction, we calculate the normalized mean-square error (MSE) as the ℓ_2 norm of the difference between the ideal noiseless signal of the second sensor $\mathbf{x}_2 = \mathbf{x}_1 * \mathbf{h}$ and the reconstructed one, $\hat{\mathbf{x}}_2$, divided by the energy of \mathbf{x}_2 , i.e.

$$\text{MSE} = \frac{\|\hat{\mathbf{x}}_2 - \mathbf{x}_2\|_2}{\|\mathbf{x}_2\|_2}.$$

The results are averaged over 50 000 realizations. Fig. 5(a) shows the normalized MSE as a function of SNR_h using the TLS approach, with or without the Cadzow denoising procedure introduced in Section IV. The different sets of curves correspond to different oversampling factors $L = K + 1, K + 4, K + 12$ (from top to bottom).

We observe that the normalized MSE can be significantly reduced by sending just a few more measurements than the minimum required (i.e., $2K + 1$). Besides, the gain provided by the Cadzow denoising procedure increases as the number of transmitted coefficients increases. In Fig. 5(b), we show the reconstruction error as a function of the signal length N . The parameters are set to be $K = 6, L = 36, \text{SNR}_h = 25$ dB, with N taking values from the set $\{128, 256, 384, 512, 768, 1024\}$. The results shown in the figure clearly exhibits the effectiveness of the Cadzow denoising algorithm on the reconstruction performance in the presence of model mismatch.

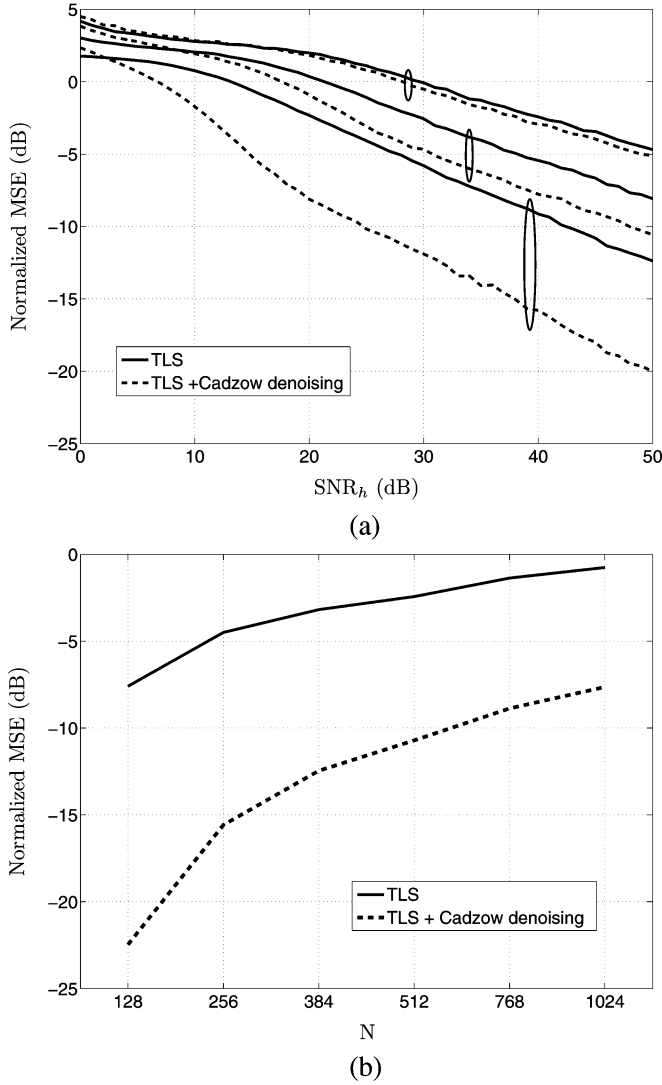


Fig. 5. Normalized MSE of the reconstructed signals. The filter coefficients are perturbed by additive white Gaussian noise. The recovery algorithm is the annihilating filter method introduced in Section IV, with or without the Cadzow denoising step. (a) $N = 256$, $K = 3$, $L = 4, 7, 15$ (from top to bottom) (b) $K = 6$, $L = 36$, $\text{SNR}_h = 25$ dB and different values for the signal length N .

B. Room Impulse Response Estimation

In this section, we use the proposed algorithm to estimate the room impulse response (RIR) in an acoustic environment. The same setup can also be relevant in an ultrawideband localization scenario in which one can sample the echo at much slower speed than the Nyquist rate and still get very good localization information.

This experiment is a bit different from our original setup. Instead of recovering two signals \mathbf{x}_1 and \mathbf{x}_2 , our goal here is to estimate the room impulse response \mathbf{h} by sending a signal and collecting a set of noisy observations. In our original setup, this is equivalent to the case where the decoder has full access to the first signal \mathbf{x}_1 , and then recovers the second signal \mathbf{x}_2 and the underlying sparse filter. In this simulation, we consider two sources of noise to simulate the imperfections on both the transmitter and receiver sides.

In our experiment, the acoustic room impulse response is generated by the standard image-source model (ISM) [24], [25].

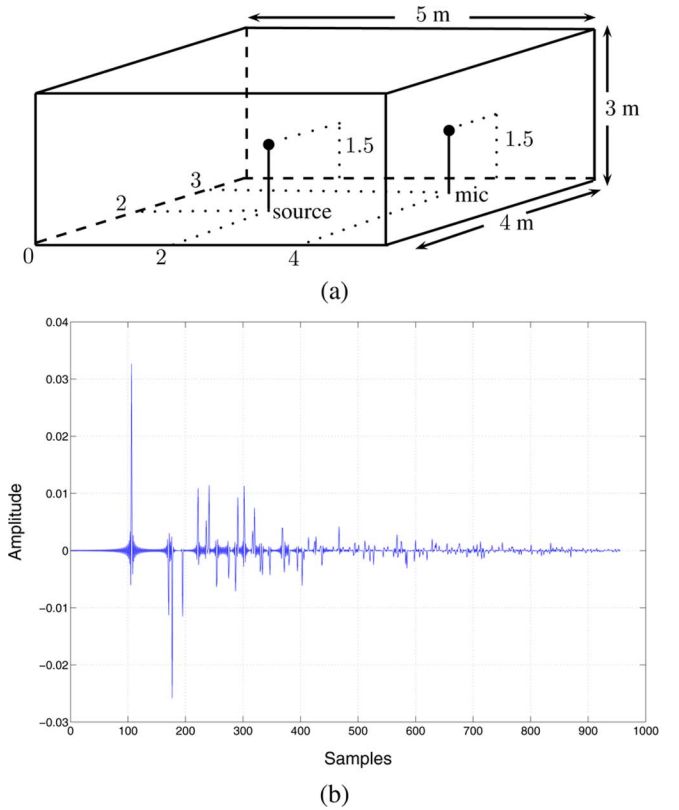


Fig. 6. Synthesizing the room impulse response using the image-source model [24], [25]. (a) Room setup. (b) Synthesized RIR using the image-source model.

This scheme provides realistic impulse responses which can be used to generate signals that would effectively be recorded in the considered environment. To produce the RIR, we used the code provided by E. A. Lehmann, available online at <http://www.watri.org.au/~ericl>. Our setup is shown in Fig. 6(a). The room has dimensions $x_r = 5$, $y_r = 4$, and $z_r = 3$ m. The reflection coefficients of all six walls are set to 0.6. The sound source is located at coordinates $x_s = 2$, $y_s = 2$, and $z_s = 1.5$ m. We put the microphone at coordinates $x_m = 4$, $y_m = 3$, and $z_m = 1.5$ m. The speed of sound is set to $c = 340$ m/s and the sampling frequency is $f_s = 16$ kHz. The impulse response is computed until the time index where its overall energy content has decreased by 20 dB, which sets the value of N to 956. Also, the RIR synthesis algorithm makes use of fractional delay filters, which allow the representation of noninteger delays for all acoustic reflections. The synthesized RIR \mathbf{h} is shown in Fig. 6(b) and consists of 150 reflections.

The setup for sending and receiving signals is shown in Fig. 7. We build the signal \mathbf{x}_1 using i.i.d. elements from the Gaussian distribution $\mathcal{N}(0, 1)$. In order to simulate the imperfections in sending this sequence through the loudspeaker, we add white Gaussian noise to \mathbf{x}_1 to reach an $\text{SNR}_{(1)}$ of 25 dB. Consequently, the transmitted signal is given by

$$\mathbf{x}_1^t = \mathbf{x}_1 + \mathbf{n}_1, \quad \text{SNR}_{(1)} = \frac{\|\mathbf{x}_1\|_2}{\|\mathbf{n}_1\|_2}.$$

Moreover, after convolving \mathbf{x}_1^t with the RIR filter \mathbf{h} , we add to the result another white Gaussian noise sequence \mathbf{n}_2 , which is

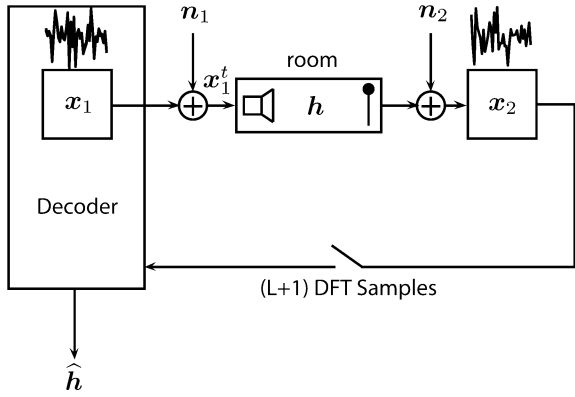


Fig. 7. The setup for sending and receiving signals in order to estimate the acoustic room impulse response.

TABLE I
THE AVERAGE ERROR (NORMALIZED TO THE SAMPLING PERIOD) IN ESTIMATING THE POSITIONS OF THE 7 STRONGEST REFLECTIONS, INDICATED BY NUMBERS IN FIG. 8. THE MATCHING ALGORITHM IS BASED ON AN EXHAUSTIVE SEARCH. THE BLACK SQUARE INDICATES THAT THE CORRESPONDING DELAY COULD NOT BE FOUND. THE RESULTS ARE AVERAGED OVER 500 TRIALS

	1	2	3	4	5	6	7
$L = 48$	0.161	2.280	2.796	■	0.635	1.740	1.951
$L = 84$	0.063	0.281	0.594	0.392	0.245	0.228	1.327
$L = 120$	0.033	0.064	0.123	0.118	0.161	0.102	1.145
$L = 156$	0.028	0.040	0.081	0.082	0.083	0.086	0.669

independent of n_1 and used to model the imperfections on the receiver side. Consequently, the microphone observes

$$\mathbf{x}_2 = \mathbf{x}_1^t * \mathbf{h} + \mathbf{n}_2$$

with

$$\text{SNR}_{(2)} = \frac{\|\mathbf{x}_1^t * \mathbf{h}\|_2}{\|\mathbf{n}_2\|_2}$$

set to 25 dB.

In the reconstruction phase, we assume that the first sequence \mathbf{x}_1 is fully available to the recovery system (note that a noisy version of \mathbf{x}_1 is convolved with the RIR filter \mathbf{h}). To estimate the RIR, we set $K = 12$ and use different oversampling factors $L = 48, 84, 120$, and 156 on \mathbf{x}_2 . Since the true reflection delays are real valued, we set the root finding part in the annihilating filter method to look for roots at machine precision. In this way, our algorithm is able to find the true delays even if they do not lie on the sampling grid. The original and the reconstructed RIRs are shown in Fig. 8. It can be verified visually that taking more measurements will result in better reconstruction quality.

Table I summarizes the average error (normalized to the sampling period) in estimating the positions of the seven strongest reflections of the original RIR (see Fig. 8), using the annihilating filter method. We do an exhaustive search to match the set of original and reconstructed spikes. First, we separately pick the seven strongest among each of the two sets. Then, from the two lists, we pick two spikes that are closest in time. If their time difference is smaller than a predefined threshold (which is set to

10 samples here), they are matched and removed from the lists. This process is continued until either all the seven strongest original spikes are matched or the minimum time difference between the remaining spikes in the two lists are above the threshold. The errors in positions are normalized to convert the absolute error (in seconds) to the error in the number of samples. The results are averaged over 500 trials. This table clearly shows the positive effect of oversampling on the performance of the recovery algorithm.

C. Binaural Hearing Aids Setup

In this experiment, we consider the signals recorded by two hearing aids mounted on the left and right ears of the user [26], [27]. We assume that the signals of the two hearing aids are related through a filtering operation. We refer to this filter as binaural filter. In the presence of a single source in far field, and neglecting reverberations and the head-shadow effect [28], the signal recorded at hearing aid 2 is simply a delayed version of the one observed at hearing aid 1. Hence, the binaural filter can be assumed to have sparsity factor $K = 1$. In the presence of reverberations and head shadowing, the filter from one microphone to the other is no longer sparse which introduces model mismatch. Despite this model mismatch, it is assumed that the main binaural features can still be well captured with a filter having only a few non-zero coefficients. Thus, the transfer function between the two received signals should be sparse, with the main peak indicating the desired relative delay.

A single sound source, located at distance d from the head of a KEMAR mannequin, moves back and forth between two angles α_{\min} and α_{\max} . The angular speed of the source is denoted as ω . The sound is recorded by the microphones of the two hearing aids, located at the ears of the mannequin. We want to retrieve the binaural filter between the two hearing aids at hearing aid 1, from limited data transmitted by hearing aid 2. Then, the main peak of the binaural filter indicates the relative delay between the two received signals, which can be used to localize the source. In addition to that, retrieving this binaural filter also allows the hearing aid 1 to find out the received signal at hearing aid 2. The setup is shown in Fig. 9. To get the head related transfer functions (HRTF) at each position of the sound source, we used the CIPIC HRTF database⁷, which is a public-domain database of high-spatial-resolution HRTF measurements [29]. The database includes 2500 measurements of head-related impulse responses for 45 different subjects.

In order to track the binaural impulse response, we process the received signals frame by frame in a DFT filter bank architecture [18], [30]. In this way, we calculate the near-sparse filter in each frame with a few measurements, considering it as a 1-sparse filter. We use a weighted overlap-add (WOLA) filter bank, which allows for an efficient realization of a DFT filter bank to obtain the short-time Fourier transform of the input signal.

The parameters of the simulation setup is as follows. The sound source is located at the distance $d = 1$ meter in front of the head, at elevation 0. It travels between $\alpha_{\min} = -45^\circ$ and $\alpha_{\max} = 45^\circ$ with the angular speed of $\omega = 18$ deg/sec. The total

⁷<http://interface.cipic.ucdavis.edu>




Fig. 8. The original RIR along with the reconstructed responses with different number of measurements in a typical trial. The parameter L is set to 156, 120, 84, and 48 from top to bottom and $N = 956$. White Gaussian noise is added to the two signals $x_1[n]$ and $x_2[n]$ to get signal to noise ratio of 25 dB in each one. The recovery algorithm reconstructs $K = 12$ Diracs of the RIR. The numbered spikes are the 7 strongest ones considered in Table I.

length of the sound is 10 seconds. In the frame by frame processing, the HRTF for the left and right ears are extracted from the database. Then, each frame of length 1024 is windowed by a Hanning window and then filtered by the impulse responses to generate the received signals for the left and right ears. The overlap between the frames is set to half of the window size.

On the reconstruction phase, we use the WOLA filter bank architecture. In our implementation, the sampling frequency is set to 16 kHz. The analysis window is a Hanning window of size 896, zero-padded with 64 zeros on both sides, to give a total length of 1024 samples. In each frame, we estimate the underlying impulse response between the two ears, using the distributed algorithm presented in Section IV. Since we are interested in the relative delay between the two signals, we set $K = 1$, i.e., we are looking for just one Dirac. To overcome the model mismatch in the system, we send more measurements than the minimum rate and use the Cadzow's iterative de-noising followed by the TLS method. The parameter L in our algorithm is selected from the set $L = 5, 10, 15, 20, 25$, which is much smaller than the frame length. This allows the left ear to localize the sound source in each frame by looking at the position of the main peak in the corresponding binaural filter. Moreover, it can reconstruct the signal received by the right ear using the frame-by-frame reconstruction approach. The original and the reconstructed sound data are available at <http://rr.epfl.ch/sparse-filtering>.

Note that if the angular speed of the source is small compared to the frame length, the binaural filter does not change




Fig. 9. Audio experiment setup. (a) A sound source travels at a distance of d meters in front of the head. It moves back and forth with angular speed ω between the horizontal angles α_{\min} and α_{\max} . There are two microphones (hearing aids) at the two ears which record the received signals. (b) Angular position of the sound source with respect to time.

too rapidly over several frames. Therefore, we can make the reconstruction more robust to noise and model mismatch by averaging the measurements along multiple frames. We used an




Fig. 10. A typical example of the binaural filter impulse response and its one-tap sparse approximation. (a) Original version. (b) Sparse approximation obtained by Cadzow iterative denoising + TLS algorithm. The parameters are $K = 1$ and $L = 20$. We observe that the non-zero coefficient is located at the position of the filter tap with the largest magnitude.

exponentially weighted averaging filter, with the averaging constant equal to $\alpha = 200$ msec. However, setting large values for α degrades the tracking capability of the algorithm.

Fig. 10 shows, as an example, the binaural filter and its approximation as a one-tap filter, estimated by the reconstruction algorithm, using $L = 20$ measurements. Because of the reverberations and the shadowing effects, the binaural filter is not exactly sparse. Moreover, the frame-by-frame processing further distorts the binaural filter. That is why we need to over-sample, in order to get a good approximation for the position of the main peak.

Fig. 11 demonstrates the localization performance of the algorithm. Fig. 11(a) shows the evolution of the original binaural impulse response over time. Fig. 11(b)–11(f) exhibits the sparse approximation to the filter, obtained by the proposed algorithm, using different number of measurements. This clearly demonstrates the effect of the over-sampling factor on the robustness of the reconstruction algorithm.

When multiple sources are concurrently active, the binaural filter is no longer sparse. Therefore, our scheme can not be applied directly. A possible solution is to apply the proposed scheme on a subband basis [31], where the assumption is that




Fig. 11. Tracking the binaural impulse response with a 1-tap filter using different number of measurements. The quality of estimating the main delay improves significantly by increasing the number of measurements. Each column in the image corresponds to the binaural impulse response at the time mentioned on the x axis. (a) Original binaural filter. (b)–(f) Tracking the evolution of the main peak with different values of the oversampling factor L . (a) Original. (b) $L = 5$. (c) $L = 10$. (d) $L = 15$. (e) $L = 20$. (f) $L = 25$.

only one source is active in each subband. However, the conjugate symmetry property can no longer be used to obtain twice as many consecutive frequency coefficients. Moreover, the oversampling ratio must be reduced to keep a constant total number of coefficients. We leave this for future work.

VII. CONCLUSION

We investigated the task of recovering distributed signals at a central decoder based on linear measurements obtained by distributed sensors, using a fixed linear measurement structure. Our main focus was on the scenario where the two signals are related through a sparse time-domain filtering operation. Under this model, we proposed universal and almost sure reconstruction strategies. We derived the achievability bounds on the minimum number of samples needed in each strategy and showed that, for universal recovery, one can not use the sparsity of the filter to reduce the number of measurements. On the other hand, by means of a computationally demanding recovery algorithm, the number of measurements in the almost sure setup is decreased effectively by $N - K$. Since $K \ll N$, this is about a two times reduction in the number of measurements. To overcome the high complexity of the recovery algorithm, we proposed a slightly

TABLE II
SUMMARY OF THE MAIN RESULTS

	sparsity	measurements	complexity
universal	K	$2N$	$O(1)$
almost sure	K	$N + K + 1 + \text{mod}(K, 2)$	combinatorial
annihilating filter	K	$N + 2K + 1$	$O(KN)$

sub-optimal yet efficient and robust distributed sensing and recovery scheme based on annihilating filters. The proposed approach needs effectively K more measurements compared to the theoretical achievability bound but its complexity is of $O(KN)$. The results are summarized in Table II. Our future work will focus on pursuing more practical applications of the proposed algorithm, e.g., in binaural hearing aids [26], [27].

APPENDIX A
THE ACHIEVABILITY BOUND FOR ALMOST SURE
RECONSTRUCTION

The goal of this Appendix is to show the achievability result for the almost sure reconstruction as stated in Proposition 2. To this end, we first state some definitions and lemmas.

Definition 6 (Spark [32]): The spark of a matrix \mathbf{A} , denoted as $\mathcal{S}(\mathbf{A})$, is defined as the largest integer L , such that every subgroup of L columns from \mathbf{A} are linearly independent.

The following lemma shows that, using a measurement matrix with Spark at least $K + 1$, we can almost surely reconstruct a K -sparse signal with only $K + 1$ measurements. Our proof follows the main ideas of the proof sketched in [4], but with a few additional details.

Lemma 1 (Sparse Reconstruction): Let $\mathbf{h} \in \mathbb{R}^N$ be a K -sparse vector with the nonzero coefficients chosen from a nonsingular distribution in \mathbb{R}^K . Furthermore, let \mathbf{A} be a matrix with N columns, satisfying $\mathcal{S}(\mathbf{A}) \geq K + 1$. Then almost surely, we can reconstruct the sparse vector \mathbf{h} from the measurements $\mathbf{y} = \mathbf{A}\mathbf{h}$.

Proof: Assume, without loss of generality, that the nonzero coefficients of \mathbf{h} lie on the first K indices $\mathcal{I} = \{1, 2, \dots, K\}$. The measurement vector \mathbf{y} lives in the span of the columns of \mathbf{A} indexed by \mathcal{I} ; we denote these columns by $\mathcal{A}_{\mathcal{I}}$ and define

$$\mathcal{A}_{\mathcal{I}} \triangleq \text{Span}\{\mathbf{A}_{\mathcal{I}}\}.$$

The combinatorial search algorithm goes through all the possible $\binom{N}{K}$ subspaces and checks to which one the measurement vector \mathbf{y} belongs. Let \mathcal{I}' be any other set of indices different from \mathcal{I} with cardinality K . Since $\mathcal{S}(\mathbf{A}) \geq K + 1$, the dimension of the sum of the two subspaces $\mathcal{A}_{\mathcal{I}}$ and $\mathcal{A}_{\mathcal{I}'}$ satisfies

$$\dim(\mathcal{A}_{\mathcal{I} \cup \mathcal{I}'}) \geq K + 1. \quad (14)$$

It can be easily verified that

$$\dim(\mathcal{A}_{\mathcal{I} \cup \mathcal{I}'}) = \dim(\mathcal{A}_{\mathcal{I}}) + \dim(\mathcal{A}_{\mathcal{I}'}) - \dim(\mathcal{A}_{\mathcal{I}} \cap \mathcal{A}_{\mathcal{I}'}).$$

Meanwhile, since $\dim(\mathcal{A}_{\mathcal{I}}) = \dim(\mathcal{A}_{\mathcal{I}'}) = K$, we have

$$\dim(\mathcal{A}_{\mathcal{I}} \cap \mathcal{A}_{\mathcal{I}'}) < K.$$

Since the columns of $\mathbf{A}_{\mathcal{I}}$ are linearly independent and the nonzero elements of \mathbf{h} have a nonsingular probability distribution, the probability that the measurement vector \mathbf{y} belongs to any low dimensional subspace of $\mathcal{A}_{\mathcal{I}}$ is zero. To see this, take any low dimensional subspace S' of $\mathcal{A}_{\mathcal{I}}$, and let \mathbf{v} be a nonzero vector in the null space of S' . We have

$$\mathbf{v}^T \mathbf{y} = \mathbf{v}^T \mathbf{A}_{\mathcal{I}} \mathbf{h}_{\mathcal{I}} = 0$$

where $\mathbf{h}_{\mathcal{I}}$ denotes the subvector of \mathbf{h} with indices given in the set \mathcal{I} . This shows that the nonzero elements of \mathbf{h} which generate \mathbf{y} live in $\text{span}^{\perp}\{\mathbf{v}\}$, a proper subspace of \mathbb{R}^K . Since the nonzero elements of \mathbf{h} have a nonsingular probability distribution, these nonzero elements have probability zero which in turn, indicates that the probability of $\mathbf{y} \in S'$ is zero.

Therefore, this intersection set of dimension less than K has total probability zero in $\mathcal{A}_{\mathcal{I}}$. Since the number of subspaces are finite, the total probability of the set of points in $\mathcal{A}_{\mathcal{I}}$ which also live in any other subspace $\mathcal{A}_{\mathcal{I}'}$ is zero. Thus, there is a one-to-one relationship between the measurements \mathbf{y} and the sparse vector \mathbf{h} almost surely. In general, the recovery algorithm should search through all $\binom{N}{K}$ possible subspaces and check to which one the measurement vector \mathbf{y} belongs. Therefore, the recovery algorithm is combinatorial in nature and computationally prohibitive. Note that if the filter \mathbf{h} lies in the intersection of two subspaces (which has probability zero under the conditioned discussed), the filter is not recoverable uniquely. ■

Now we are ready to prove the achievability result stated in Proposition 2. We divide the sensing and recovery architecture into two parts:

- I. The two sensors send frequency measurements to the decoder, necessary to find the underlying K -sparse filter \mathbf{h} . For fixed values of N and K , the structure of this part does not change by choosing different values for the number of measurements M_1 and M_2 in the achievability region.
- II. The sensors transmit complementary subsets of the DFT coefficients necessary to construct the two signals \mathbf{x}_1 and \mathbf{x}_2 . The sensors adjust their share on this part of linear measurements in order to reach the required rates M_1 and M_2 .

We show the results of Proposition 2 in two separate situations, corresponding to whether K is an odd or even number.

a) K is odd:

In part I, define the matrix \mathbf{M} of size $((K + 1)/(2) + 1) \times N$ as

$$\mathbf{M}_{ij} = \begin{cases} 1 & i = j, j \in \{1, \dots, \frac{K+1}{2} + 1\} \\ 0 & \text{otherwise.} \end{cases}$$

The two sets of measurements in part I, denoted as $\mathbf{y}_1^{(1)}$ and $\mathbf{y}_2^{(1)}$, are given by

$$\begin{aligned} \mathbf{y}_1^{(1)} &= \mathbf{M}\mathbf{F}\mathbf{x}_1 \\ \mathbf{y}_2^{(1)} &= \mathbf{M}\mathbf{F}\mathbf{x}_2 \end{aligned} \quad (15)$$

in which \mathbf{F} is the DFT matrix of size $N \times N$. The number of *real* measurements in (15) is $K + 2$ for each sensor,

composed of 1 real DC component plus $(K + 1)/2$ complex frequency components.

In order to reconstruct the sparse filter \mathbf{h} , the central decoder computes the consecutive frequency components of \mathbf{h} in the index range $\mathcal{N} = \{-(K + 1)/(2), \dots, (K + 1)/(2)\}$. This is done by dividing the two sets of measurements in (15) index by index and making use of the complex conjugate property of the Fourier transform of real sequences. Note that this division is allowed as long as the denominator is not zero. Since a zero frequency component forces \mathbf{x}_1 to lie on a hyperplane in \mathbb{R}^N , the frequency components are nonzero almost surely for the signal \mathbf{x}_1 drawn from a nonsingular distribution in \mathbb{R}^N (More precisely, we only require that the frequency components of \mathbf{x}_1 be nonzero in the range \mathcal{N} .)

In the matrix form, we obtain

$$\mathbf{H} = \mathbf{W}\mathbf{h}$$

where \mathbf{H} is the Fourier transform of \mathbf{h} in the range given by \mathcal{N} . The matrix \mathbf{W} is a complex valued generalized Vandermonde system [33] of size $(K + 2) \times N$ which consists of $K + 2$ consecutive rows of the DFT matrix, corresponding to the indices in the set \mathcal{N} , i.e.

$$W_{pq} = e^{-j2\pi(p-(K+1)/(2)-1)(q-1)/N}$$

for $p \in \{1, \dots, K + 2\}$ and $q \in \{1, \dots, N\}$. Thus, by the properties of generalized Vandermonde matrices, we get

$$S(\mathbf{W}) = K + 2.$$

Note that, by assumption, the nonzero elements of \mathbf{h} satisfy a nonsingular distribution in \mathbb{R}^K . Therefore, by Lemma 1, it is possible to find the K -sparse filter \mathbf{h} from $\mathbf{y}_1^{(1)}$ and $\mathbf{y}_2^{(1)}$ almost surely.

In the second part of the sensing and recovery architecture, the sensors send complementary subsets of the remaining frequency information of their signals \mathbf{x}_1 and \mathbf{x}_2 to the decoder. By having access to the frequency information of one of the signals \mathbf{x}_i , $i = 1, 2$, the decoder uses the point-wise relationship

$$X_2[m] = X_1[m]H[m] \quad (16)$$

to get the frequency content of the other sequence, with the help of the known filter \mathbf{h} . Note that since the nonzero elements of \mathbf{h} satisfy a nonsingular distribution in \mathbb{R}^K , the frequency components of the filter are all nonzero almost surely. Therefore, it is possible to find $X_2[m]$ from $X_1[m]$ or *vice versa*. The two sensors share the rest of the frequency indices to reach any point on the achievability region. Note that (16) also allows the decoder to calculate the real (imaginary) part of $X_2[m]$ by knowing the imaginary (real) part of the other sequence $X_1[m]$ and *vice versa*. This is achieved by equating the real and imaginary parts in (16) and solving a linear system of equations.

In this way, each sensor sends $K + 2$ real measurements in part I for the decoder to find the underlying sparse filter \mathbf{h} almost surely. Moreover, in part II, the sensors send complementary information to the decoder. This indicates that the sensors can choose the sampling pairs given by $K + 2 \leq M_1 \leq N$ and $M_2 = K + 2 + N - M_1$. Therefore, the algorithm described earlier attains any point on the achievability region given in Proposition 2.

b) K is even:

The same measurement and reconstruction techniques apply. In this case, however, we are able to find the K -sparse filter \mathbf{h} from $K + 1$ real measurements, instead of $K + 2$. This corresponds to the matrix \mathbf{M} of size $((K/2) + 1) \times N$, given by

$$M_{ij} = \begin{cases} 1 & i = j, j \in \{1, \dots, \frac{K}{2} + 1\} \\ 0 & \text{otherwise.} \end{cases}$$

Therefore, when K is even, we have

$$\begin{aligned} M_1 &\geq \min\{K + 1, N\} \\ M_2 &\geq \min\{K + 1, N\} \\ \text{and } M_1 + M_2 &\geq \min\{N + K + 1, 2N\}. \end{aligned}$$

APPENDIX B

DERIVATION OF THE RANK PROPERTY

We investigate the rank property of the Toeplitz matrices constructed by a signal which is composed by a sum of K exponentials. We show that whenever the number of rows and columns of this matrix is larger than or equal to K , its rank is K .

Consider a sequence \mathbf{h} of length N with K nonzero coefficients $\{\alpha_i\}_{i=1}^K$ at positions $\{n_i\}_{i=1}^K$. The $L + L' + 1$ consecutive frequency samples of \mathbf{h} (starting from index $P - L$) are given by

$$H[m] = \sum_{i=1}^K \alpha_i \omega_i^m, m = (P - L), \dots, (P + L'),$$

where $\omega_i \triangleq e^{-j2\pi n_i/N}$ and $L, L' \geq K - 1$. We build the Toeplitz matrix \mathbf{H} of size $(L' + 1) \times (L + 1)$ as follows:

$$\mathbf{H} = \begin{bmatrix} H[P] & H[P-1] & \dots & H[P-L] \\ H[P+1] & H[P] & \dots & H[P-L+1] \\ \vdots & \vdots & \ddots & \vdots \\ H[P+L'] & H[P+L'-1] & \dots & H[P+L'-L] \end{bmatrix}.$$

To show that the rank of \mathbf{H} is K , we can write it as the product

$$\mathbf{H} = \mathbf{UCV},$$

where the three matrices are given by

$$\mathbf{U}_{(L'+1) \times K} = \begin{bmatrix} \omega_1^P & \omega_2^P & \dots & \omega_K^P \\ \omega_1^{P+1} & \omega_2^{P+1} & \dots & \omega_K^{P+1} \\ \vdots & \vdots & \ddots & \vdots \\ \omega_1^{P+L'} & \omega_2^{P+L'} & \dots & \omega_K^{P+L'} \end{bmatrix}$$

$$\mathbf{C}_{K \times K} = \begin{bmatrix} \alpha_1 & 0 & \cdots & 0 \\ 0 & \alpha_2 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \alpha_K \end{bmatrix}$$

and

$$\mathbf{V}_{K \times (L+1)} = \begin{bmatrix} 1 & \omega_1^{-1} & \cdots & \omega_1^{-L} \\ 1 & \omega_2^{-1} & \cdots & \omega_2^{-L} \\ \vdots & \vdots & \ddots & \vdots \\ 1 & \omega_K^{-1} & \cdots & \omega_K^{-L} \end{bmatrix}.$$

The matrices \mathbf{U} and \mathbf{V} are Vandermonde matrices of rank K . Therefore, since the diagonal elements of \mathbf{C} are nonzero, the matrix \mathbf{H} is of rank K .

REFERENCES

- [1] D. Slepian and J. K. Wolf, "Noiseless coding of correlated information sources," *IEEE Trans. Inf. Theory*, vol. 19, pp. 471–480, Jul. 1973.
- [2] T. Berger, "Multiterminal source coding," *Proc. Lectures Notes Presented at CISM Summer School Inf. Theory Approach to Commun.*, Jul. 1977.
- [3] Y. Oohama, "Gaussian multiterminal source coding," *IEEE Trans. Inf. Theory*, vol. 43, pp. 1912–1923, Nov. 1997.
- [4] D. Baron, M. B. Wakin, M. F. Duarte, S. Sarvotham, and R. G. Baraniuk, "Distributed Compressed Sensing," Elect. Comput. Eng. Dep., Rice University, Houston, TX, Tech. Rep. ECE-0612, Dec. 2006.
- [5] E. J. Candès, J. Romberg, and T. Tao, "Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information," *IEEE Trans. Inf. Theory*, vol. 52, pp. 489–509, Feb. 2006.
- [6] D. L. Donoho, "Compressed sensing," *IEEE Trans. Inf. Theory*, vol. 52, pp. 1289–1306, Apr. 2006.
- [7] J. A. Tropp and A. Gilbert, "Signal recovery from random measurements via orthogonal matching pursuit," *IEEE Trans. Inf. Theory*, vol. 53, pp. 4655–4666, Dec. 2007.
- [8] W. Bajwa, J. Haupt, A. Sayeed, and R. Nowak, "Compressive wireless sensing," in *Proc. Int. Conf. Inf. Process. Sens. Netw.*, Apr. 2006, pp. 134–142.
- [9] M. Rabbat, J. Haupt, A. Singh, and R. Nowak, "Decentralized compression and redistribution via randomized gossiping," in *Proc. Int. Conf. Inf. Process. Sens. Netw.*, Apr. 2006, pp. 51–59.
- [10] W. Wang, M. Garofalakis, and K. Ramchandran, "Distributed sparse random projections for refinable approximation," in *Proc. Int. Conf. Inf. Process. Sens. Netw.*, Apr. 2007, pp. 331–339.
- [11] P. Stoica and R. Moses, *Introduction to Spectral Analysis*. Englewood Cliffs, NJ: Prentice-Hall, 2005.
- [12] M. Vetterli, P. Marziliano, and T. Blu, "Sampling signals with finite rate of innovation," *IEEE Trans. Signal Process.*, vol. 50, pp. 1417–1428, Jun. 2002.
- [13] T. Blu, P. Dragotti, M. Vetterli, P. Marziliano, and L. Coulot, "Sparse sampling of signal innovations," *IEEE Signal Process. Mag.*, vol. 25, no. 2, pp. 31–40, 2008.
- [14] W. U. Bajwa, J. Haupt, G. Raz, and R. Nowak, "Compressed channel sensing," in *Proc. 42nd Ann. Conf. Inf. Sci. Syst. (CISS 08)*, Princeton, NJ, Mar. 2008, pp. 5–10.
- [15] J. Paredes, G. Arce, and Z. Wang, "Ultra-Wideband compressed sensing: Channel estimation," *IEEE J. Sel. Topics Signal Process.*, vol. 1, pp. 383–395, Oct. 2007.
- [16] J. A. Cadzow, "Signal enhancement—A composite property mapping algorithm," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 36, pp. 49–67, Jan. 1988.
- [17] A. V. Oppenheim and R. W. Schaffer, *Discrete-Time Signal Processing*. Englewood Cliffs, NJ: Prentice-Hall, 1989.
- [18] R. E. Crochiere and L. R. Rabiner, *Multirate Digital Signal Processing*. Englewood Cliffs, NJ: Prentice-Hall, 1983.
- [19] B. Rao and K. Arun, "Model based processing of signals: A state space approach," *Proc. IEEE*, vol. 80, pp. 283–309, Feb. 1992.
- [20] G. H. Golub and C. F. Van Loan, *Matrix Computations*. Baltimore, MD: Johns Hopkins Univ. Press, 1989.
- [21] R. T. Chien, "Cyclic decoding procedure for the Bose-Chaudhuri-Hocquenghem codes," *IEEE Trans. Inf. Theory*, vol. 10, pp. 357–363, Oct. 1964.
- [22] G. E. Pfander, H. Rauhut, and J. Tanner, "Identification of matrices having a sparse representation," *IEEE Trans. Signal Process.*, vol. 56, pp. 5376–5388, Nov. 2008.
- [23] A. Hormati and M. Vetterli, "Distributed compressed sensing: Sparsity models and reconstruction algorithms using annihilating filter," in *Proc. ICASSP*, Apr. 2008, pp. 5141–5144.
- [24] J. B. Allen and D. A. Berkley, "Image method for efficiently simulating small-room acoustics," *J. Acoust. Soc. Amer.*, vol. 65, pp. 943–950, Apr. 1979.
- [25] E. A. Lehmann and A. M. Johansson, "Prediction of energy decay in room impulse responses simulated with an image-source model," *J. Acoust. Soc. Amer.*, vol. 124, pp. 269–277, July 2008.
- [26] O. Roy, "Distributed signal processing for binaural hearing aids," Ph.D. thesis, Ecole Polytechnique Fédérale de Lausanne, Lausanne, Switzerland, Oct. 2008.
- [27] O. Roy and M. Vetterli, "Rate-Constrained collaborative noise reduction for wireless hearing aids," *IEEE Trans. Signal Process.*, vol. 57, no. 2, pp. 645–657, 2009.
- [28] J. Blauert, *Spatial Hearing: The Psychophysics of Human Sound Localization*. Cambridge, MA: MIT Press, 1997.
- [29] V. R. Algazi, R. O. Duda, D. M. Thompson, and C. Avendano, "The CIPIC HRTF database," in *Proc. IEEE WASPAA*, Oct. 2001, vol. 100, pp. 99–102.
- [30] M. Vetterli and J. Kovačević, *Wavelets and Subband Coding*. Englewood Cliffs, NJ: Prentice-Hall, 1995.
- [31] C. Faller, "Parametric coding of spatial audio," Ph.D. dissertation, Ecole Polytechnique Fédérale de Lausanne, Lausanne, Switzerland, Jul. 2004.
- [32] D. L. Donoho and M. Elad, "Maximal sparsity representation via ℓ_1 minimization," in *Proc. Natl. Acad. Sci.*, Mar. 2003, vol. 100, pp. 2197–2202.
- [33] F. R. Gantmacher, *Matrix Theory*. Providence, RI: AMS Chelsea, 2000.



Ali Hormati (S'06) received the B.Sc. and M.Sc. degrees in communication systems from Sharif University of Technology, Tehran, Iran, in 2001 and 2003, respectively.

From 2003 to 2006, he worked as a communication systems engineer in industry on the communication protocols of software radio systems. From March to September 2006, he attended the Predoc program in communication systems at the Ecole Polytechnique Fédérale de Lausanne (EPFL), Switzerland. Since October 2006, he has been a

Research Assistant with the Audiovisual Communications Laboratory, EPFL. He has been a Teaching Assistant for classes in digital signal processing and advanced signal processing. His research interests include sampling theory, compressed sensing, inverse problems and acoustic tomography.




Olivier Roy (S'05) received the M.S. degree in communication systems from the Ecole Polytechnique Fédérale de Lausanne (EPFL), Lausanne, Switzerland, in 2004, and the Ph.D. degree from the same institution in 2008.

He is now a Postdoctoral Researcher with the Karmanos Cancer Institute, Detroit, MI. From March to September 2002, he was a software engineer with Logitech, Fremont, CA, where he worked on Internet video streaming. From August 2003 to March 2004, he was a research intern with the Institute

for Telecommunications Research, University of South Australia, Mawson Lakes. There, he investigated information-theoretic aspects of multiantenna communication systems. From June to August 2007, he was with Phonak, a hearing aid company based in Stäfa, Switzerland, where he developed noise reduction schemes for binaural hearing aids. From October 2004 to November 2008, he was a research assistant with the Audiovisual Communications Laboratory, EPFL. His research interests are in distributed estimation and coding, microphone array processing, inverse problems, and acoustic tomography.

Mr. Roy received a Best Student Paper Award at the Australian Communication Theory Workshop in 2004.



Yue M. Lu (S'04–M'07) received the B.Eng. and M.Eng. degrees in electrical engineering from Shanghai Jiao Tong University, China, in 1999 and 2002, respectively. He received the M.Sc. degree in mathematics and the Ph.D. degree in electrical engineering from the University of Illinois at Urbana-Champaign in 2007.

He was a Research Assistant with the University of Illinois at Urbana-Champaign, and has worked for Microsoft Research Asia, Beijing, China, and Siemens Corporate Research, Princeton, NJ. He is now with the Audiovisual Communications Laboratory, Ecole Polytechnique Fédérale de Lausanne (EPFL), Switzerland. His research interests include signal processing for sensor networks; the theory, constructions, and applications of multiscale geometric representations for multidimensional signals; image and video processing; and sampling theory.

Dr. Lu received the Most Innovative Paper Award of IEEE International Conference on Image Processing (ICIP) in 2006 for his paper (with M. N. Do) on the construction of directional multiresolution image representations, and the Student Paper Award of IEEE ICIP in 2007.



Martin Vetterli (F'95) received the Dipl. El.-Ing. degree from ETH Zurich (ETHZ), Switzerland, in 1981, the M.S. degree from Stanford University, Stanford, CA, in 1982, and the Doctorat ès Sciences degree from EPFL Lausanne (EPFL), Switzerland, in 1986.

He was a Research Assistant at Stanford University and EPFL, and has worked for Siemens and AT&T Bell Laboratories. In 1986, he joined Columbia University, New York, where he was last an Associate Professor of Electrical Engineering and codirector of the Image and Advanced Television

Laboratory. In 1993, he joined the University of California at Berkeley, where he was a Professor with the Department of Electrical Engineering and Computer Sciences until 1997, and now holds an Adjunct Professor position. Since 1995, he has been a Professor of Communication Systems at EPFL Lausanne, Switzerland, where he chaired the Communications Systems Division (1996/1997), and heads the Audiovisual Communications Laboratory. From 2001 to 2004, he directed the National Competence Center in Research on mobile information and communication systems. He is also a Vice-President for International Affairs at EPFL since October 2004. He has held visiting positions with ETHZ (1990) and Stanford University (1998). His research interests include sampling, wavelets, multirate signal processing, computational complexity, signal processing for communications, digital image/video processing, joint source/channel coding, and signal processing for sensor networks.

Dr. Vetterli is a fellow of the IEEE, a fellow of the ACM, and a member of SIAM. He is on the editorial boards of *Applied and Computational Harmonic Analysis*, *The Journal of Fourier Analysis and Application*, and the IEEE JOURNAL ON SELECTED TOPICS IN SIGNAL PROCESSING. He received the Best Paper Award of EURASIP in 1984 for his paper on multidimensional subband coding, the Research Prize of the Brown Boverly Corporation (Switzerland) in 1986 for his doctoral thesis, the IEEE Signal Processing Society's Senior Paper Awards in 1991 in 1996, and 2007 (for papers with D. LeGall and K. Ramchandran, and Marziliano and Blu, respectively). He won the Swiss National Latsis Prize in 1996, the SPIE Presidential award in 1999, and the IEEE Signal Processing Technical Achievement Award in 2001. He was a member of the Swiss Council on Science and Technology until December 2003. He was a plenary speaker at various conferences (e.g., IEEE ICIP, ICASSP, ISIT). He is the coauthor of books with J. Kovačević (*Wavelets and Subband Coding*), with P. Prandoni (*Signal Processing for Communications*), and with J. Kovačević and V. K. Goyal (*The World of Fourier and Wavelets*). He has published about 120 journal papers on a variety of topics in signal/image processing and communications, holds a dozen patents, and is an ISI highly cited researcher in engineering.