

Raking the Cocktail Party

Ivan Dokmanic, *Student Member, IEEE*, Robin Scheibler, *Student Member, IEEE*, and Martin Vetterli, *Fellow, IEEE*

Abstract—We present the concept of an acoustic rake receiver—a microphone beamformer that uses echoes to improve the noise and interference suppression. The rake idea is well-known in wireless communications; it involves constructively combining different multipath components that arrive at the receiver antennas. Unlike spread-spectrum signals used in wireless communications, speech signals are not orthogonal to their shifts. Therefore, we focus on the spatial structure, rather than the temporal. Instead of explicitly estimating the channel, we create correspondences between early echoes in time and image sources in space. These multiple sources of the desired and the interfering signal offer additional spatial diversity that we can exploit in the beamformer design. We present several “intuitive” and optimal formulations of acoustic rake receivers, and show theoretically and numerically that the rake formulation of the maximum signal-to-interference-and-noise beamformer offers significant performance boosts in terms of noise and interference suppression. Beyond signal-to-noise ratio, we observe gains in terms of the *perceptual evaluation of speech quality* (PESQ) metric for the speech quality. We accompany the paper by the complete simulation and processing chain written in Python. The code and the sound samples are available online at <http://lcav.github.io/AcousticRakeReceiver/>.

Index Terms—Acoustic rake receiver, beamforming, echo sorting, interference cancellation, noise suppression, perceptual evaluation of speech quality (PESQ), room impulse response.

I. INTRODUCTION

RAKE receivers take advantage of the multipath propagation, instead of trying to mitigate it. The basic idea of the rake receivers, which are commonly used in wireless communications, is to coherently add the multipath components and thus increase the effective signal-to-noise ratio (SNR). The original scheme was developed for single-input-single-output systems [1], and it was later extended to arrays of antennas [2], [3] that exploit spatial diversity. Thus multipath components that are not resolvable with a single antenna because they arrive at similar times, become resolvable because they arrive from different directions.

In spite of the success of the rake receivers in wireless communications, the principle has not received significant attention in room acoustics. Nevertheless, constructive use of echoes in

rooms to improve beamforming has been mentioned in the literature [4]–[6]. In particular, the term *acoustic rake receiver* (ARR) was used in the SCENIC project proposal [4].

The list of ingredients for ARR in room acoustics is similar as in wireless communications: a wave (acoustic instead of electromagnetic) propagates in space; reflections and scattering cause the wave to arrive at the receiver through multiple paths in addition to the direct path, and these multipath components all contain the source waveform.

The main difference is that in room acoustics we do not get to design the input signal. Spreading sequences used in CDMA are designed to be near-orthogonal to their shifts and orthogonal between different users, which facilitates the multipath channel estimation; such orthogonality is not exhibited by speech. Moreover, speech segments are very long with respect to the time delay between two consecutive echoes, and they are *a priori* unknown at the receiver.

On the contrary, there are no significant differences in terms of the spatial structure. If we know where the echoes are coming from, we can design spatial processing algorithms—for example beamformers—that use multiple copies of the same signal arriving from different directions.

Imagine first that we know the room geometry. Then, if we localize the source, we can predict where its echoes will come from using simple geometric rules [7], [8]. Localizing the direct signal in a reverberant environment is a well-understood problem [9]. What is more, we do not need to know the room shape in detail—locations of the most important reflectors (ceiling, floor, walls) suffice to localize the major echoes. In many cases this knowledge is readily available from the floor plans or measurements. In ad-hoc deployments, the room geometry may be difficult to obtain. If that is the case, we can first perform a calibration step to learn it. An appealing method to infer the room geometry is by using sound, as was demonstrated recently [10]–[13].

We may still be able to take advantage of the echoes without estimating the room geometry. Note that we are not after the room geometry itself; rather, we only need to know where the early echoes are coming from. Echoes can be seen as signals emitted by *image sources*—mirror images of the true source across reflecting walls [7]. Knowing where the echoes are coming from is equivalent to knowing where the image sources are.

Image source localization can be solved, for example, by *echo sorting* as described in [13]. Alternatively, O'Donovan, Duraiswami and Zotkin [5] propose to use an *audio camera* with a large number of microphones to find the images. Once the image sources are localized (in a calibration phase or otherwise), we can predict their movement using geometrical rules, as discussed in Section V. Thus, the acoustic raking is

Manuscript received July 15, 2014; revised January 15, 2015; accepted March 06, 2015. Date of publication March 23, 2015; date of current version nulldate. This work was supported by the ERC Advanced Investigators Grant: Sparse Sampling: Theory, Algorithms, and Applications SPARSAM no. 247006, and a Google Doctoral Fellowship. The guest editor coordinating the review of this manuscript and approving it for publication was XXXXX XXXXXX.

The authors are with the LCAV-EPFL, 1015 Lausanne, Switzerland (e-mail: ivan.dokmanic@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/JSTSP.2015.2415761

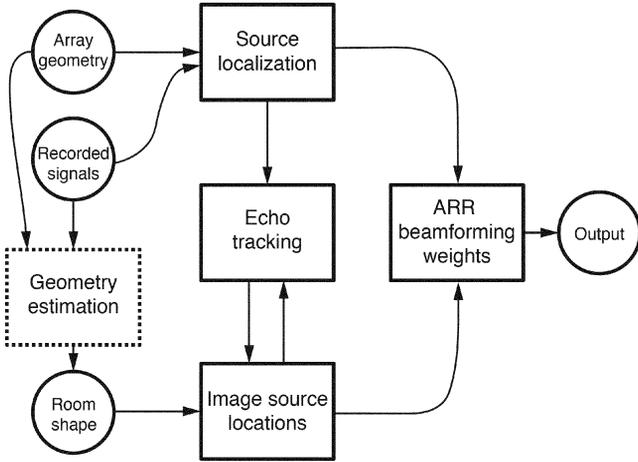


Fig. 1. A block diagram for acoustic rake receivers. In this paper, we focus on ARR beamforming weight computation, and we briefly discuss echo tracking and image source localization. The geometry estimation block is optional (room geometry could be known in advance), hence the dashed box.

a multi-stage process comprising image source localization, image source tracking, and beamforming weight computation. The complete block diagram of an acoustic rake receiver is shown in Fig. 1.

It is interesting to note the analogy between the ARR and the human auditory perception. It is well established that the early echoes improve speech intelligibility [14], [15]. In fact, adding energy in the form of early echoes (approximately within the first 50 ms of the room impulse response (RIR)) is equivalent to adding the same energy to the direct sound [14]. This observation suggests new designs for indoor beamformers, with different choices of performance measures and reference signals. A related discussion of this topic is given by Habets and co-authors [16], who examine the tradeoff between dereverberation and denoising in beamforming. In addition to the standard SNR, we propose to use the useful-to-detrimental ratio (UDR), first defined by Lochner and Burger [15], and used by Bradley, Sato and Picard [14]. We generalize UDR to a scenario with interferers, defining it as the ratio of the direct and early reflection energy to the energy of the noise and interference.

ARRs focus on the early part of the RIR, trying to concentrate the energy contained in the early echoes. In that regard, there are similarities between ARR and channel shortening [17], [18]. Channel shortening produces filters that are much better behaved than complete inversion, e.g., by the multiple-input-output-theorem (MINT) [19], [20]. Nevertheless, it is still tacitly assumed that we know the acoustic impulse responses between the sources and the microphones. In contrast to channel shortening, as well as other methods assuming this knowledge [19], [21], we never attempt the difficult task of estimating the impulse responses. Our task is simpler: we only need to detect the early echoes, and lift them to 3D space as image sources.

A. Main Contributions and Limitations

We introduce the acoustic rake receiver (ARR) as the echo-aware microphone beamformer. We present several formulations with different properties, and analyze their behavior theoretically and numerically. The analysis shows that ARR

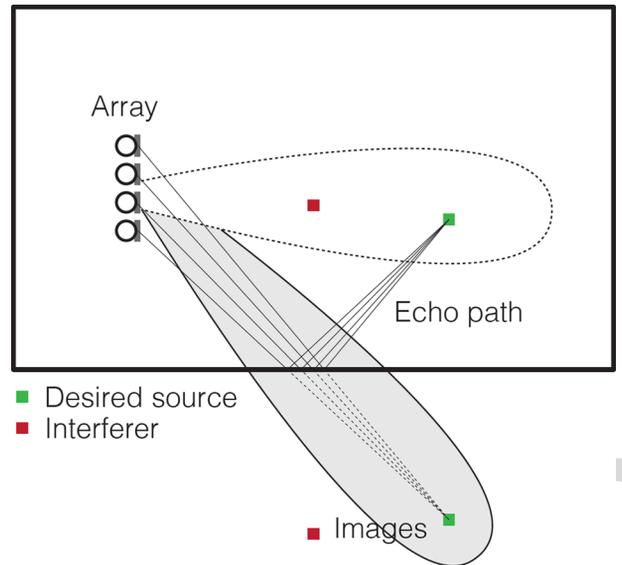


Fig. 2. Listening behind an interferer by listening to echoes (illustration). A beam directed towards the desired, green source will necessarily pick up the red interferer's signal. Acoustic rake receiver avoids it by beamforming towards the echoes of the desired source.

lead to significantly improved SNR and interference cancellation when compared with standard beamformers that only extract the direct path. ARR can suppress interference in cases when conventional beamforming is bound to fail, for example when an interferer is occluding the desired source (an illustration is given in Fig. 2; for a sneak-peak of real beam patterns, fast-forward to Fig. 7). We present optimal formulations that outperform the earlier delay-and-sum (DS) approaches [6], especially when interferers are present. Significant gains are observed not only in terms of signal-to-interference-and-noise ratio (SINR) and UDR, but also in terms of perceptual evaluation of speech quality (PESQ) [22].

The raking microphone beamformers are particularly well-suited to extracting the desired speech signal in the presence of interfering sounds, in part because they can focus on echoes of the desired sound and cancel the echoes of the interfering signals. The analogous human capacity to focus on a particular acoustic stimulus while not perceiving other, unwanted sounds is called the *cocktail party effect* [23]. This paper's title was inspired by that analogy.

We design and apply the ARR in the frequency domain. Frequency domain formulation is simple and concise; it allows us to focus on objective gains from acoustic raking. Time-domain designs [24] offer better control over the impulse responses of the beamforming filters, but they are out of the scope of this paper. For a recent time-domain approach to ARR, see [25].

Let us also mention some limitations of our results. For clarity, the numerical experiments are presented in a 2D "room", and as such are directly applicable to planar (e.g., linear or circular) arrays. Extension to 3D arrays is straightforward. We do not discuss robust formulations that address uncertainties in the array calibration. Microphones are assumed to be ideally omni-directional with a flat frequency response. Except for Section V, we assume that the locations of the image sources are known. We explain how to find the image sources

when the room geometry is either known or unknown, but we do not provide a deep overview of the geometry estimation techniques. To this end, we suggest a number of references for the interested reader. We consider the walls to be flat-fading; in reality, they are frequency selective. We do not discuss the estimation of various covariance matrices [26].

The results in this paper are reproducible. Python (NumPy) [27] code for the beamforming routines, for the STFT processing engine, and to generate the figures and the sound samples is available online at <http://lcav.github.io/AcousticRakeReceiver/>.

B. Paper Outline

In Section II we explain the notation and the signal model used in the paper. A brief overview of the relevant beamforming techniques and performance analysis is given in Section III. We formulate the acoustic rake receiver in Section IV, and we present a theoretical and numerical analysis of the corresponding beamformers. Section V explains how to locate the image sources, and comments on localizing the direct source. Numerical experiments are presented in Section VI.

II. NOTATION AND SIGNAL MODEL

We denote all matrices by bold uppercase letters, for example \mathbf{A} , and all vectors by bold lowercase letters, for example \mathbf{x} . The Hermitian transpose of a matrix or a vector is denoted by $(\cdot)^H$, as in \mathbf{A}^H , and the Euclidean norm of a vector by $\|\cdot\|$, that is, $\|\mathbf{x}\| \stackrel{\text{def}}{=} (\mathbf{x}^H \mathbf{x})^{1/2}$.

Suppose that the desired source of sound is at the location \mathbf{s}_0 in a room. Sound from this source arrives at the microphones located at $\{\mathbf{r}_m\}_{m=1}^M$ via the direct path, but also via the echoes from the walls. The echoes can be replaced by the image sources—mirror images of the true sources across the corresponding walls—according to the image source model [7], [8]. An important consequence is that instead of modeling the source of the desired or the interfering signal as a single point in a room, we can model it as a collection of points in free space. A more detailed discussion of the image source model is given in Section V.

Denote the signal emitted by the source $\tilde{x}[n]$ (e.g., the speech signal). Then all the image sources emit $\tilde{x}[n]$ as well, and the signals from the image sources reach the microphones with the appropriate delays. In our application, the essential fact is that the echoes correspond to image sources. We denote the image source positions by \mathbf{s}_k , $1 \leq k \leq K$, where K denotes the largest number of image sources considered. Note that we do not care about the sequence of walls that generates \mathbf{s}_k , nor do we care about how many walls are in the sequence. For us, all \mathbf{s}_k are simply additional sources of the desired signal. The described setup is illustrated Fig. 3.

Suppose further that there is an interferer at the location \mathbf{q}_0 (for simplicity, we consider only a single interferer). The interferer emits the signal $\tilde{z}[n]$, and its image sources emit $\tilde{z}[n]$ as well. Similarly as for the desired source, we denote by \mathbf{q}_k , $1 \leq k \leq K'$ the positions of the interfering image sources, with K' being the largest number of interfering image source considered. The model mismatch (e.g., the image sources of high orders and the late reverberation) and the noise are absorbed in the term $\tilde{n}_m[n]$.

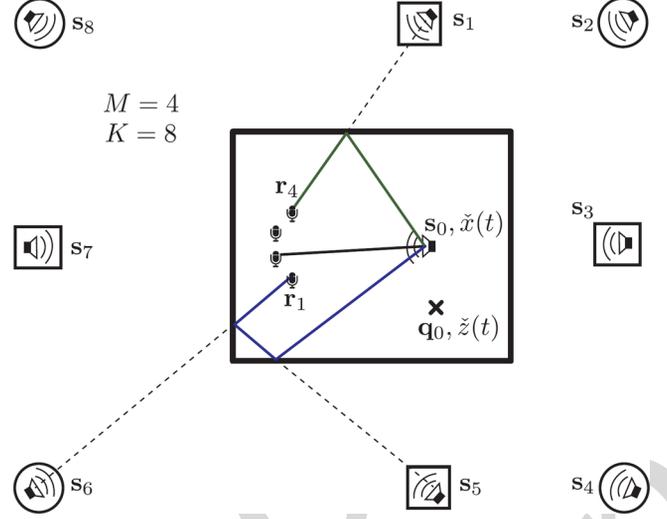


Fig. 3. Illustration of the notation and concepts. Echoes of the desired signal emitted at \mathbf{s}_0 can be modeled as a direct sound coming from the image sources of \mathbf{s}_0 . Two generations of image sources are illustrated: first ($\mathbf{s}_1, \mathbf{s}_3, \mathbf{s}_5, \mathbf{s}_7$) and second ($\mathbf{s}_2, \mathbf{s}_4, \mathbf{s}_6, \mathbf{s}_8$), as well as the corresponding sound rays for \mathbf{s}_5 and \mathbf{s}_6 . The interferer is located at \mathbf{q}_0 (its image sources are not shown), and the microphones are located at $\mathbf{r}_1, \dots, \mathbf{r}_4$.

The signal received by the m th microphone is then a sum of convolutions

$$\tilde{y}_m[n] = \sum_{k=0}^K (\tilde{a}_m(\mathbf{s}_k) * \tilde{x})[n] + \sum_{k=0}^{K'} (\tilde{a}_m(\mathbf{q}_k) * \tilde{z})[n] + \tilde{n}_m[n], \quad (1)$$

where $\tilde{a}_m(\mathbf{s}_k)$ denotes the impulse response of the channel between the source located at \mathbf{s}_k and the m th microphone—in this case a delay and a scaling factor.

We design and analyze all the beamformers in the frequency domain. That is, we will be working with the DTFT of the discrete-time signal \tilde{x} ,

$$x(e^{j\omega}) \stackrel{\text{def}}{=} \sum_{n \in \mathbb{Z}} \tilde{x}[n] e^{-j\omega n}. \quad (2)$$

In practical implementations, we use the discrete-time short-time Fourier transform (STFT). More implementation details are given in Section VI.

Using these notations, we can write the signal picked up by the m th microphone as

$$y_m(e^{j\omega}) = \sum_{k=0}^K a_m(\mathbf{s}_k, \Omega) x(e^{j\omega}) + \sum_{k=0}^{K'} a_m(\mathbf{q}_k, \Omega) z(e^{j\omega}) + n_m(e^{j\omega}), \quad (3)$$

where $n_m(e^{j\omega})$ models the noise and other errors, and $a_m(\mathbf{s}_k, \Omega)$ denotes the m th component of the steering vector for the source \mathbf{s}_k . The steering vector is the Fourier transform of the continuous version of the impulse response $\tilde{a}(\mathbf{s}_k)$, evaluated at the frequency Ω . The discrete-time frequency ω and the continuous-time frequency Ω are related as $\omega = \Omega T_s$, where

T_s is the sampling period. The steering vector is then simply $\mathbf{a}(\mathbf{s}_k, \Omega) = [a_m(\mathbf{s}_k, \Omega)]_{m=0}^{M-1}$.

We can write out the entries of the steering vectors explicitly for a point source in free space. They are given as the appropriately scaled free-space Green's functions for the Helmholtz equation [28],

$$a_m(\mathbf{s}_k, \Omega) = \frac{\alpha_k}{4\pi\|\mathbf{r}_m - \mathbf{s}_k\|} e^{-j\kappa\|\mathbf{r}_m - \mathbf{s}_k\|}, \quad (4)$$

where we define the wavenumber as $\kappa \stackrel{\text{def}}{=} \Omega/c$, and α_k is the attenuation corresponding to \mathbf{s}_k .

Using vector notation, the microphone signals can be written concisely as

$$\mathbf{y}(e^{j\omega}) = \mathbf{A}_s(e^{j\omega})\mathbf{1}x(e^{j\omega}) + \mathbf{A}_q(e^{j\omega})\mathbf{1}z(e^{j\omega}) + \mathbf{n}(e^{j\omega}), \quad (5)$$

where $\mathbf{A}_s(e^{j\omega}) \stackrel{\text{def}}{=} [\mathbf{a}(\mathbf{s}_1, \Omega), \dots, \mathbf{a}(\mathbf{s}_K, \Omega)]$, $\mathbf{A}_q(e^{j\omega}) \stackrel{\text{def}}{=} [\mathbf{a}(\mathbf{q}_1, \Omega), \dots, \mathbf{a}(\mathbf{q}_{K'}, \Omega)]$, and $\mathbf{1}$ is the all-ones vector. From here onward, we suppress the frequency dependency of the steering vectors and the beamforming weights to reduce the notational clutter.

III. BEAMFORMING PRELIMINARIES

Microphone beamformers combine the outputs of multiple microphones in order to achieve spatial selectivity, thereby suppressing noise and interference [29]. In the frequency domain, a beamformer forms a linear combination of the microphone outputs to yield the output u . That is,

$$u = \mathbf{w}^H \mathbf{y} = \mathbf{w}^H \mathbf{A}_s \mathbf{1}x + \mathbf{w}^H \mathbf{A}_q \mathbf{1}z + \mathbf{w}^H \mathbf{n}, \quad (6)$$

where the vector $\mathbf{w} \in \mathbb{C}^M$ contains the beamforming weights.

The weights \mathbf{w} are often selected so that they optimize some design criterion. Common examples of beamformers are the delay-and-sum (DS) beamformer, minimum-variance-distortionless-response (MVDR) beamformer, maximum-signal-to-interference-and-noise (Max-SINR) beamformer, and minimum-mean-squared-error (MMSE) beamformer. In this paper we discuss the rake formulation of the DS and the Max-SINR beamformers; for completeness, we first describe the non-raking variants.

A. Delay-and-Sum Beamformer

DS is the simplest and often quite effective beamformer [29]. Assume that we want to listen to a source at \mathbf{s} . Then we form the DS beamformer by compensating the propagation delays from the source \mathbf{s} to the microphones \mathbf{r}_m ,

$$u_{\text{DS}} = \frac{1}{M} \sum_{m=0}^{M-1} y_m e^{j\kappa\|\mathbf{r}_m - \mathbf{s}\|} \quad (7)$$

$$\approx \frac{x}{4\pi\|\bar{\mathbf{r}} - \mathbf{s}\|} + \frac{1}{M} \sum_{m=0}^{M-1} n_m, \quad (8)$$

where $\bar{\mathbf{r}} = 1/M \sum_{m=0}^{M-1} \mathbf{r}_m$ denotes the center of the array. The beamforming weights can be read out from (7) as

$$\mathbf{w}_{\text{DS}} = \frac{\mathbf{a}(\mathbf{s})}{\|\mathbf{a}(\mathbf{s})\|}, \quad (9)$$

where we used the definition of y_m (5) and the definition of the steering vector (4). We can see from (8) that if $\mathbf{n} \sim \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I}_M)$, then the output noise is distributed according to $\mathcal{N}(0, \sigma^2/M)$, that is, we obtain an M -fold decrease in the noise variance at the output with respect to any reference microphone.

B. Maximum Signal-to-Interference-and-Noise Ratio Beamformer

The SINR is an important figure of merit used to assess the performance of ARR, and to compare it with the standard non-raking beamformers. It is computed as the ratio of the power of the desired output signal to the power of the undesired output signal. The desired output signal is the output signal due to the desired source, while the undesired signal is the output signal due to the interferers and noise.

For a desired source at \mathbf{s} and an interfering source at \mathbf{q} we can write

$$\text{SINR} \stackrel{\text{def}}{=} \frac{\mathbb{E}|\mathbf{w}^H \mathbf{a}(\mathbf{s})x|^2}{\mathbb{E}|\mathbf{w}^H (\mathbf{a}(\mathbf{q})z + \mathbf{n})|^2} = \sigma_x^2 \frac{\mathbf{w}^H \mathbf{a}(\mathbf{s})\mathbf{a}(\mathbf{s})^H \mathbf{w}}{\mathbf{w}^H \mathbf{K}_{nq} \mathbf{w}}, \quad (10)$$

where \mathbf{K}_{nq} is the covariance matrix of the noise and the interference.

It is compelling to pick \mathbf{w} that maximizes the SINR (10) [29]. The maximization can be solved by noting that the rescaling of the beamformer weights leaves the SINR unchanged. This means that we can minimize the denominator subject to numerator being an arbitrary constant. The solution is given as

$$\mathbf{w}_{\text{SINR}} = \frac{\mathbf{K}_{nq}^{-1} \mathbf{a}_s}{\mathbf{a}_s^H \mathbf{K}_{nq}^{-1} \mathbf{a}_s}. \quad (11)$$

Using the definition (10), we can derive the SINR for the Max-SINR beamformer as

$$\text{SINR} = \sigma_x^2 \mathbf{a}_s^H \mathbf{K}_{nq}^{-1} \mathbf{a}_s. \quad (12)$$

Because \mathbf{K}_{nq}^{-1} is a Hermitian symmetric positive definite matrix, it has an eigenvalue decomposition as $\mathbf{K}_{nq}^{-1} = \mathbf{U}^H \mathbf{\Lambda} \mathbf{U}$, where \mathbf{U} is unitary, and $\mathbf{\Lambda}$ is diagonal with positive entries. We can write $\mathbf{a}_s^H \mathbf{K}_{nq}^{-1} \mathbf{a}_s = (\mathbf{U} \mathbf{a}_s)^H \mathbf{\Lambda} (\mathbf{U} \mathbf{a}_s)$. Because $\|\mathbf{U} \mathbf{a}_s\|^2 = \|\mathbf{a}_s\|^2$, and because $\mathbf{\Lambda}$ is positive, increasing $\|\mathbf{a}_s\|^2$ typically leads to an increased SINR, although we can construct counterexamples. This will be important when we discuss the SINR gain of the Rake-Max-SINR beamformer.

IV. ACOUSTIC RAKE RECEIVER

In this section, we present several formulations of the ARR. The Rake-DS beamformer is a straightforward generalization of the conventional DS beamformer. The one-forcing beamformer implements the idea of steering a fixed beam power towards every image source, while trying to minimize the interference and noise. The Rake-Max-SINR and Rake-Max-UDR beamformers optimize the corresponding performance measures; we show in Section VI that the Rake-Max-SINR beamforming performs best in terms of all evaluation criteria except, as expected, in terms of UDR.

A. Delay-and-Sum Raking

If we had access to every echo separately (*i.e.*, not summed with all the other echoes), we could align them all to maximize the performance gain. Unfortunately, this is not the case: each microphone picks up the convolution of speech with the impulse response, which is effectively a sum of overlapping echoes of the speech signal. If we only wanted to extract the direct path, we would use the standard DS beamformer (9). To build the Rake-DS receiver, we create a DS beamformer for every image source, and average the outputs,

$$\frac{1}{K+1} \sum_{k=0}^K \frac{\alpha'_k}{M} \sum_{m=0}^{M-1} y_m e^{j\kappa \| \mathbf{r}_m - \mathbf{s}_k \|}, \quad (13)$$

where $\alpha'_k \stackrel{\text{def}}{=} \alpha_k / (4\pi \| \mathbf{r}_m - \mathbf{s}_k \|)$. We read out the beamforming weights from (13) as

$$\mathbf{w}_{\text{R-DS}} = \frac{1}{\| \sum_k \mathbf{a}(\mathbf{s}_k) \|} \sum_{k=0}^K \mathbf{a}(\mathbf{s}_k) = \frac{\mathbf{A}_s \mathbf{1}}{\| \mathbf{A}_s \mathbf{1} \|}, \quad (14)$$

where we chose the scaling in analogy with (9) (scaling of the weights does not alter the output SINR). It can be seen that this is just a scaled sum of the steering vectors for each image source.

B. One-Forcing Raking

A different approach, based on intuition, is to design a beamformer that listens to all K image sources with the same power, and at the same time minimizes the noise and interference energy:

$$\begin{aligned} & \underset{\mathbf{w} \in \mathbb{C}^M}{\text{minimize}} \mathbb{E} \left| \sum_{k=0}^{K'} \mathbf{w}^H \mathbf{a}(\mathbf{q}_k) z + \mathbf{w}^H \mathbf{n} \right|^2 \\ & \text{subject to } \mathbf{w}^H \mathbf{a}(\mathbf{s}_k) = 1, \forall 0 \leq k \leq K. \end{aligned} \quad (15)$$

Alternatively, we may choose to null the interfering source and its image sources. Both cases are an instance of the standard linearly-constrained-minimum-variance (LCMV) beamformer [30]. Collecting all the steering vectors in a matrix, we can write the constraint as $\mathbf{w}^H \mathbf{A}_s = \mathbf{1}^T$. The solution can be found in closed form as

$$\mathbf{w}_{\text{OF}} = \mathbf{K}_{nq}^{-1} \mathbf{A}_s (\mathbf{A}_s^H \mathbf{K}_{nq}^{-1} \mathbf{A}_s)^{-1} \mathbf{1}_M. \quad (16)$$

A few remarks are in order. First, with M microphones, it does not make sense to increase K beyond M , as this results in more constraints than degrees of freedom. Second, using this beamformer is a bad idea if there is an interferer along the ray through the microphone array and any of image sources.

As with all LCMV beamformers, adding linear constraints uses up degrees of freedom that could otherwise be used for noise and interference suppression. It is better to let the “beamformer decide” or “the beamforming procedure decide” on how to maximize a well-chosen cost function; one such procedure is described in the next subsection.

C. Max-SINR Raking

The main workhorse of the paper is the Rake-Max-SINR. We compute the weights so as to maximize the SINR, taking into account the echoes of the desired signal, and the echoes of the interfering signal,

$$\underset{\mathbf{w} \in \mathbb{C}^M}{\text{maximize}} \frac{\mathbb{E} \left| \sum_{k=0}^K \mathbf{w}^H \mathbf{a}(\mathbf{s}_k) x \right|^2}{\mathbb{E} \left| \sum_{k=0}^{K'} \mathbf{w}^H \mathbf{a}(\mathbf{q}_k) z + \mathbf{w}^H \mathbf{n} \right|^2}. \quad (17)$$

The logic behind this expression can be summarized as follows: we present the beamforming procedure with a set of good sources, whose influence we aim to maximize at the output, and with a set of bad sources, whose power we try to minimize at the output. Interestingly, this leads to the standard Max-SINR beamformer with a structured steering vector and covariance matrix. Define the combined noise and interference covariance matrix as

$$\mathbf{K}_{nq} \stackrel{\text{def}}{=} \mathbf{K}_n + \sigma_z^2 \left(\sum_{k=0}^{K'} \mathbf{a}(\mathbf{q}_k) \right) \left(\sum_{k=0}^{K'} \mathbf{a}(\mathbf{q}_k) \right)^H, \quad (18)$$

where \mathbf{K}_n is the covariance matrix of the noise term, and σ_z^2 is the power of the interferer at a particular frequency.

Then the solution to (17) is given as

$$\mathbf{w}_{\text{R-SINR}} = \frac{\mathbf{K}_{nq}^{-1} \mathbf{A}_s \mathbf{1}}{\mathbf{1}^H \mathbf{A}_s^H \mathbf{K}_{nq}^{-1} \mathbf{A}_s \mathbf{1}}. \quad (19)$$

It is interesting to note that when $\mathbf{K}_{nq} = \sigma^2 \mathbf{I}_M$ (*e.g.*, no interferers and iid noise), the Rake-Max-SINR beamformer reduces to $\mathbf{A}_s \mathbf{1} / \| \mathbf{A}_s \mathbf{1} \|$, which is exactly the Rake-DS beamformer. This is analogous to the non-raking DS beamformer (9).

D. Max-UDR Raking

Finally, it is interesting to investigate what happens if we choose the weights that optimize the perceptually motivated UDR [14], [15]. The UDR expresses the fact that adding early reflections (up to 50 ms in the RIR) is as good as adding the energy to the direct sound in terms of speech intelligibility. The useful signal is a *coherent* sum of the direct and early reflected speech energy, so that

$$\text{UDR} = \frac{\mathbb{E} \sum_{k=0}^K | \mathbf{w}^H \mathbf{a}(\mathbf{s}_k) x |^2}{\mathbb{E} \left| \sum_{k=0}^{K'} \mathbf{w}^H \mathbf{a}(\mathbf{q}_k) z + \mathbf{w}^H \mathbf{n} \right|^2}, \quad (20)$$

In applications K is rarely large enough to cover all the reflections occurring within 50 ms, simply because it is too optimistic to assume we know all the corresponding image sources. Therefore, (20) typically underestimates the UDR.

Alas, because (20) is specified in the frequency domain, it is challenging to control whether the reflections in the numerator arrive before or after the direct sound. Nevertheless, it is interesting to analyze it as it provides a basis for future work on

time-domain raking formulations, and a meaningful evaluation of the raking algorithms presented in this paper.

To compute the Rake-Max-UDR weights, we solve the following program

$$\underset{\mathbf{w} \in \mathbb{C}^M}{\text{maximize}} \frac{\mathbb{E} \sum_{k=0}^K \left| \mathbf{w}^H \mathbf{a}(\mathbf{s}_k) x \right|^2}{\left| \mathbb{E} \left[\mathbf{w}^H \sum_{k=0}^{K'} \mathbf{a}(\mathbf{q}_k) z + \mathbf{w}^H \mathbf{n} \right]^2 \right|}. \quad (21)$$

This amounts to maximizing a particular generalized Rayleigh quotient,

$$\frac{\mathbf{w}^H \mathbf{A}_s \mathbf{A}_s^H \mathbf{w}}{\mathbf{w}^H \mathbf{K}_{nq} \mathbf{w}}. \quad (22)$$

Assuming that \mathbf{K}_{nq} has a Cholesky decomposition as $\mathbf{K}_{nq} = \mathbf{C}^H \mathbf{C}$ we can rewrite the quotient (22) as

$$\frac{\tilde{\mathbf{w}}^H (\mathbf{C}^{-1})^H \mathbf{A}_s \mathbf{A}_s^H \mathbf{C}^{-1} \tilde{\mathbf{w}}}{\tilde{\mathbf{w}}^H \tilde{\mathbf{w}}}, \quad (23)$$

where $\tilde{\mathbf{w}} \stackrel{\text{def}}{=} \mathbf{C} \mathbf{w}$. The maximum of this expression is

$$\lambda_{\max}((\mathbf{C}^{-1})^H \mathbf{A}_s \mathbf{A}_s^H \mathbf{C}^{-1}), \quad (24)$$

where $\lambda_{\max}(\cdot)$ denotes the largest eigenvalue of the matrix in the argument. The maximum is achieved by the corresponding eigenvector $\tilde{\mathbf{w}}_{\max}$. Then the optimal weights are given as

$$\mathbf{w}_{\text{R-UDR}} = \mathbf{C}^{-1} \tilde{\mathbf{w}}_{\max}. \quad (25)$$

E. SINR Gain From Raking

Intuitively, if we have multiple sources of the desired signal scattered in space, and we account for it in the design, we should do at least as well as when we ignore the image sources. Let us see how large the gain can be for the Rake-Max-SINR beamformer. We have that

$$\text{SINR} = \sigma_x^2 (\mathbf{A}_s \mathbf{1})^H \mathbf{K}_{nq}^{-1} (\mathbf{A}_s \mathbf{1}). \quad (26)$$

Intuitively, the larger the norm of $\mathbf{A}_s \mathbf{1}$, the better the SINR (as \mathbf{K}_{nq} is positive). To explicitly see if there is any gain in using the acoustic rake receiver, we should compare the standard Max-SINR beamformer with the Rake-Max-SINR, *e.g.*, we should evaluate

$$\frac{(\sum_k \mathbf{a}(\mathbf{s}_k))^H \mathbf{K}_{nq}^{-1} (\sum_k \mathbf{a}(\mathbf{s}_k))}{\mathbf{a}(\mathbf{s}_0)^H \mathbf{K}_{nq}^{-1} \mathbf{a}(\mathbf{s}_0)}. \quad (27)$$

One possible interpretation of (27) is that we ask whether the steering vectors $\mathbf{a}(\mathbf{s}_k)$ sum coherently or they cancel out.

To answer this, assume that \mathbf{s}_k , $0 \leq k \leq K$ are the desired sources (true and image), and let $\beta \stackrel{\text{def}}{=} \sum_{k=1}^K (\alpha_k / \alpha_0)^2$, where α_k is the strength of the source \mathbf{s}_k received by the array. Then

$$\mathbb{E} \left(\left\| \sum_{k=0}^K \mathbf{a}(\mathbf{s}_k) \right\|^2 \right) \approx (1 + \beta) \mathbb{E}(\|\mathbf{a}(\mathbf{s}_0)\|^2), \quad (28)$$

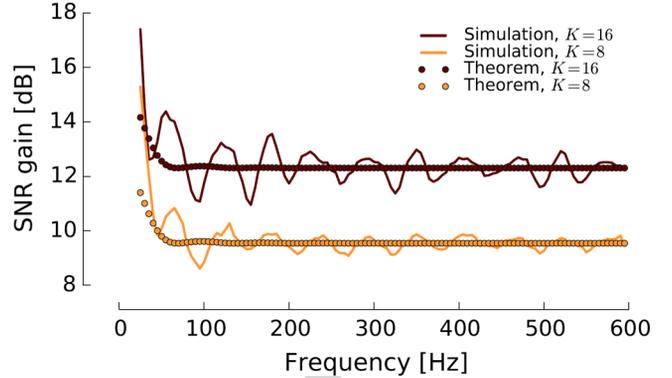


Fig. 4. Comparison of the simulated SNR gains and the theoretical prediction from Theorem 1 for $K = 8$, and $K = 16$. The theoretical prediction of the gain is $10 \log_{10}(8 + 1) \approx 9.54$ for $K = 8$, and $10 \log_{10}(16 + 1) \approx 12.30$ for $K = 16$.

that is, we can expect an increase in the output SINR approximately by a factor of $(1 + \beta)$ when using the Rake-Max-SINR beamformer. This statement is made precise in Theorem 1 in the Appendix. It holds when \mathbf{K}_{nq} has eigenvalues of similar magnitude, which is typically not the case in the presence of interferers. However, we show in Section VI that with interferers present, the gains actually increase.

A couple of remarks are in order:

- 1) This result is in expectation; it says that on average, the SINR will increase by a factor of $(1 + \beta)$. In the worst case, the steering vectors $\mathbf{a}(\mathbf{s}_k)$ can even cancel out so that the SINR decreases. But the numerical experiments suggest that this is very rare in practice, and we can on the other hand observe large gains.
- 2) We see that summing the phasors in $a_m(\mathbf{s}_k)$ behaves as a two-dimensional random walk. It is known that the root-mean-square distance of a 2D random walk from the origin after n steps is \sqrt{n} [31].
- 3) Due to the far-field assumption in Theorem 1, the attenuations α_k are assumed to be independent of the microphones; in reality they do depend on the source locations. However, they also depend on a number of additional factors, for example wall attenuations and radiation patterns of the sources. Therefore, for simplicity, we consider them to be independent. One can verify that this assumption does not change the described trend.

It is reassuring to observe the behavior suggested by (28) in practice. Fig. 4 shows the comparison of the prediction by Theorem 1 with the SNR gains observed in simulated rooms. In this case, we are comparing the pure SNR gain for white noise, without interferers. To generate Fig. 4, we randomized the location of the source inside the rectangular room. For simplicity we fixed the signal power as received by the microphones to the same value for all the image sources, so that the expected gain is $K + 1$ in the linear scale. The curves agree near-perfectly with the prediction of Theorem 1.

V. FINDING AND TRACKING THE ECHOES

Thus far we assumed that the locations of the image sources are known. In this section we briefly describe some methods to localize them when they are *a priori* unknown. We assume that

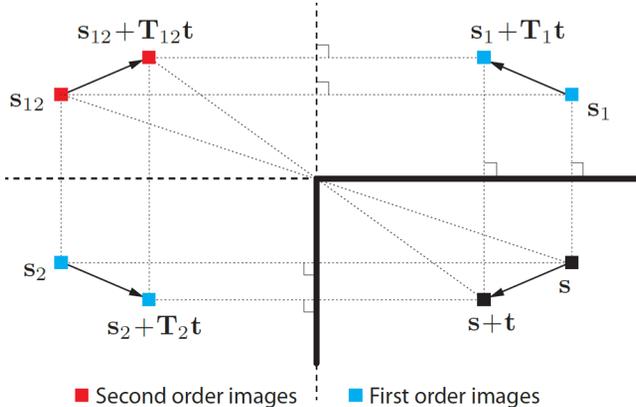


Fig. 5. Illustration of image source tracking in rectangular geometries.

we can localize the true source, or at least one image source. Combined with the knowledge of the room geometry, this suffices to find the locations of other image sources [32].

A. Known Room Geometry

In many cases, for example for fixed deployments, the room geometry is known. This knowledge could be obtained at the time of the deployment, or from blueprints. In most indoor environments, we encounter a large number of planar reflectors. These reflectors correspond to image sources. With reference to Fig. 5, we can easily compute the image source locations [7] (we note that the image source model is not limited to right angle geometries [8]).

Suppose that the real source is located at \mathbf{s} . Then the image source with respect to wall i is computed as,

$$\mathbf{im}_i(\mathbf{s}) = \mathbf{s} + 2\langle \mathbf{p}_i - \mathbf{s}, \mathbf{n}_i \rangle \mathbf{n}_i, \quad (29)$$

where i indexes the wall, \mathbf{n}_i is the outward normal associated with the i th wall, and \mathbf{p}_i is any point belonging to the i th wall. Analogously, we compute image sources corresponding to higher-order reflections,

$$\mathbf{im}_j(\mathbf{im}_i(\mathbf{s})) = \mathbf{im}_i(\mathbf{s}) + 2\langle \mathbf{p}_j - \mathbf{im}_i(\mathbf{s}), \mathbf{n}_j \rangle \mathbf{n}_j. \quad (30)$$

The above expressions are valid regardless of the dimensionality, concretely in 2D and 3D.

B. Acoustic Geometry Estimation

When the room geometry is not known, it is possible to estimate it using the same array that we use for beamforming. Recently a number of different methods appeared in the literature that propose to use sound to estimate the shape of a room. For example, in [10] the authors use a dictionary of wall impulse responses recorded with a particular array. In [11] the authors use tools from projective geometry together with the Hough transform to estimate the room geometry. In [13] the authors derive an *echo sorting* mechanism that finds the image sources, from which the room geometry is then derived.

C. Without Estimating the Room Geometry

To design an ARR, we do not really need to know how the room looks like; we only need to know where the major echoes are coming from. One possible approach is to locate the image

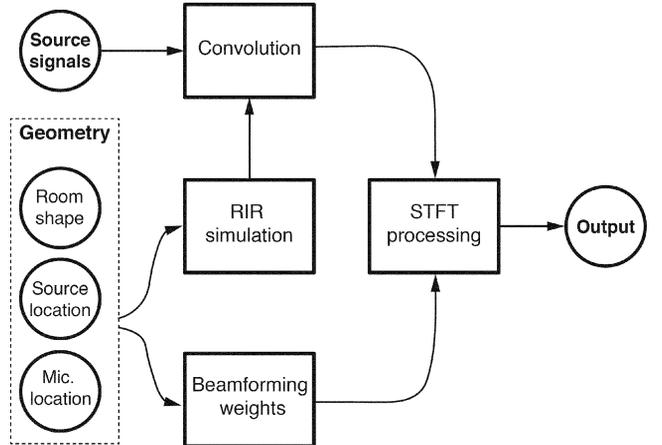


Fig. 6. Block diagram of the simulation setup used for numerical experiments.

sources in the initial calibration phase, and then track their movement by tracking the true source.

We propose a tracking rule that leverages the knowledge of the displacement of the true source. Again with reference to Fig. 5, we can state the following simple proposition.

Proposition 1: Suppose that the room has only right angles so that the walls are parallel with the coordinate axes. Let the source move from \mathbf{s} to $\mathbf{s} + \mathbf{t}$. Then any image source \mathbf{s}_k , moves to a point given by

$$\mathbf{s}_k + \mathbf{T}\mathbf{t}, \quad (31)$$

where $\mathbf{T} = \text{diag}(\pm 1, \mp 1)$ for odd generations, and $\mathbf{T} = \pm \mathbf{I}_2$ for even generations.

Proof: The proof follows directly from the figure. The displacement of the true source is the same as the displacement of the image source, passed through a series of reflections. Reflection matrices are diagonal matrices with ± 1 on the diagonal, and determinant equal to -1 , hence the result. \square

The usefulness of this proposition is that it gives us a tool to track the image sources even when we do not know the room geometry (as long as it has right angles). A possible use scenario is to start with a calibration procedure with a controlled source, and perform the echo sorting to find multiple image sources. Then if possible, we assign to each image source a generation (this is in fact a by-product of echo sorting), or we try different hypotheses using Proposition 1, and choose the one that maximizes the output SINR.

VI. NUMERICAL EXPERIMENTS

In this section, we validate the described theoretical results through numerical experiments. First, we analyze the beam-patterns produced by the ARR; second, we evaluate the SINR for various beamformers as a function of the number of image sources used in weight computation; and third, we evaluate the PESQ metric [22]. Finally, we show spectrograms that reveal visually the improved interferer and noise suppression achieved by the ARR.

A. Simulation Setup

We use a simple room acoustic framework written in Python, which relies on Numpy and Scipy for matrix computations [27].

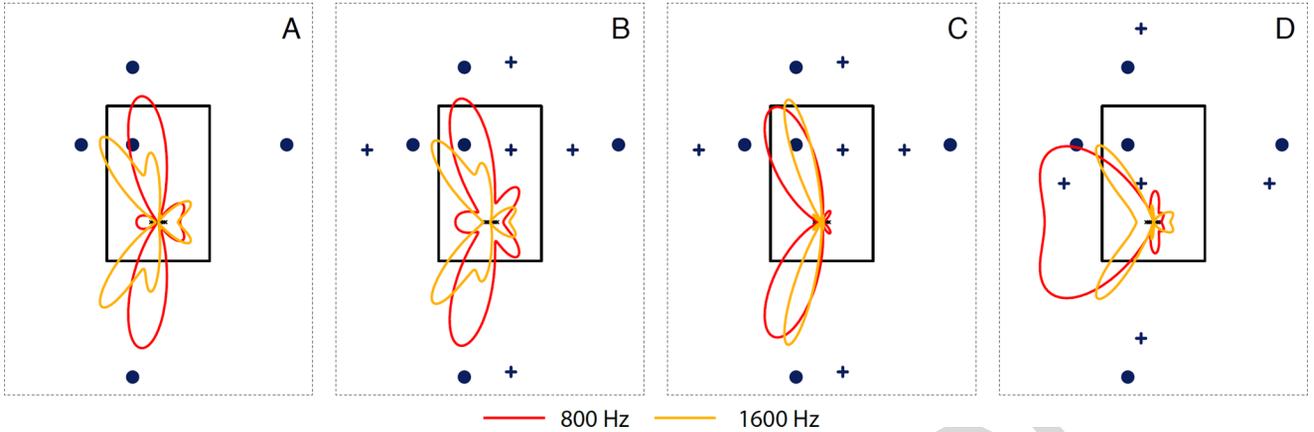


Fig. 7. Beam patterns in different scenarios. The rectangular room is 4 by 6 metres and contains a source of interest (\bullet) and an interferer ($+$). The first-order image sources are also displayed. The weight computation of the beamformer includes the direct source and the first order image sources of both desired source and interferer (when applicable). (A) Rake-Max-SINR, no interferer, (B) Rake-Max-SINR, one interferer, (C) Rake-Max-UDR, one interferer, (D) Rake-Max-SINR, interferer is in direct path.

We limit ourselves in this paper to 2D geometry and rectangular rooms. In all experiments, the sampling frequency F_s was set to 8 kHz. An overview of the simulation setup is shown in Fig. 6.

Starting from the room geometry and the positions of the sources and microphones, we first compute the locations of all images sources up to ten generations. The reflectivity of the walls is fixed to 0.9. The RIR between the source \mathbf{s}_0 and the microphone \mathbf{r}_m is convolved with an ideal low-pass filter in the continuous domain and then sampled at the sampling frequency F_s ,

$$\tilde{a}_m(\mathbf{s}_0)[n] = \sum_{k=0}^K \frac{\alpha_k}{4\pi\|\mathbf{r}_m - \mathbf{s}_k\|} \text{sinc}\left(n - F_s \frac{\|\mathbf{r}_m - \mathbf{s}_k\|}{c}\right), \quad (32)$$

where K is the number of image sources considered. We choose the limits of n such that the cardinal sine in (32) decays sufficiently to avoid artifacts. The discrete signals from all sound sources are then convolved with their respective RIRs, and added together to obtain the m th microphone's signal.

We use the discrete-time STFT processing with a frame size of $L = 4096$ samples, 50% overlap and zero padding on both sides of the signal by $L/2$, and compute the beamforming in the frequency domain. A real fast Fourier transform of size $2L$ and a Hann window are used in the analysis. By exploiting the conjugate symmetry of the real FFT we only need to compute $L + 1$ beamforming weights, one for every positive frequency bin. The length L is dictated by the length of the beamforming filters in the time-domain and was set empirically so as to avoid any aliasing in the filter responses. We synthesize the output signal using the conventional overlap-add method [33].

B. Results

1) *Beampatterns*: We first inspect the beampatterns produced by the Rake-Max-SINR and Rake-Max-UDR beamformers for different source-interferer placements. We consider a 4 m \times 6 m rectangular room with a source of interest at (1 m, 4.5 m) and a linear microphone array centered at (2 m, 1.5 m), parallel to the x -axis. Spacing between the microphones was set to 8 cm. In Fig. 7, we show the beampatterns for four different configurations of the source and the interferer. We consider a

scenario without an interferer, one with an interferer placed favorably at (2.8 m, 4.3 m), and finally one where the interferer is placed half-way between the desired source and the array, at (1.5 m, 3 m).

The last scenario is the least favorable. Interestingly, we can observe that the Rake-Max-SINR beampattern adjusts by completely ignoring the direct path, and steering the beam towards the echoes of the source of interest. This is validating the intuition that we can “hear behind an interferer by listening for the echoes”. Note that such a pattern cannot be achieved by a beamformer that only takes into account the direct path. We further note that, while the beampatterns only show the magnitude of the beamformer's response, the phase plays an important role with multiple sources present.

2) *SINR Gains From Raking*: In the experiments in this subsection, we set the power of the desired source and of the interferer to be equal, $\sigma_x^2 = \sigma_z^2 = 1$. The noise covariance matrix is set to $10^{-3} \cdot \mathbf{I}_M$. We use a circular array of $M = 12$ microphones with a diameter of 30 cm, and randomize the position of the desired source and the interferer inside the room. The resulting curves show median performance out of 20000 runs.

Fig. 8 shows output SINR for different beamformers. The one-forcing beamformer is left out because it performs poorly in terms of SINR, as predicted in the earlier discussion. Clearly, the Rake-Max-SINR beamformer outperforms all others. The output SINR for beamformers using only the direct path (Max-SINR and DS) remains approximately constant. The UDR is plotted against the number of image sources for various beamformers in Fig. 9. Even though the Rake-Max-UDR beamformer performs well in terms of the two measures, its output is perceptually displeasing due to audible pre-echoes; in informal listening tests, the Rake-Max-SINR beamformer did not produce such artifacts. It is interesting to note that the Rake-Max-SINR also performs well in terms of the UDR. Similar SINR gains to those in Fig. 8 are observed in Fig. 10 over a range of frequencies. It is therefore justified to extrapolate the results at one frequency in Fig. 8 to the wideband SINR.

3) *Evaluation of Speech Quality*: We complement the informal listening tests and the evaluation of SINR and UDR with extensive simulations to assess the improvement in speech

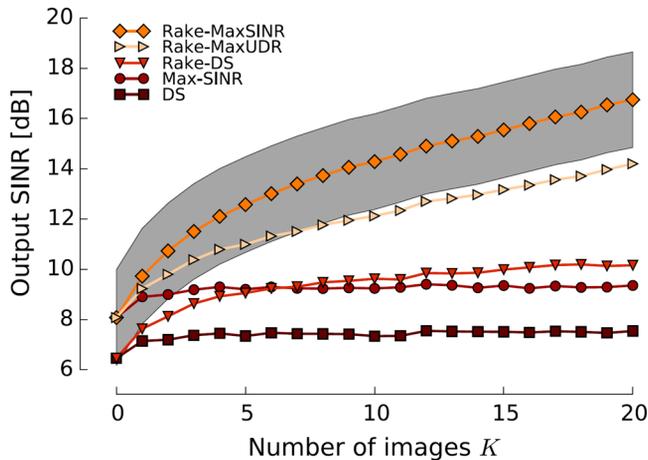


Fig. 8. Median output SINR plotted against the number of image sources used in the design for different beamformers, at a frequency $f = 1$ kHz. The shaded area contains the Rake-Max-SINR output SINRs for 50% of the 20000 Monte Carlo runs.

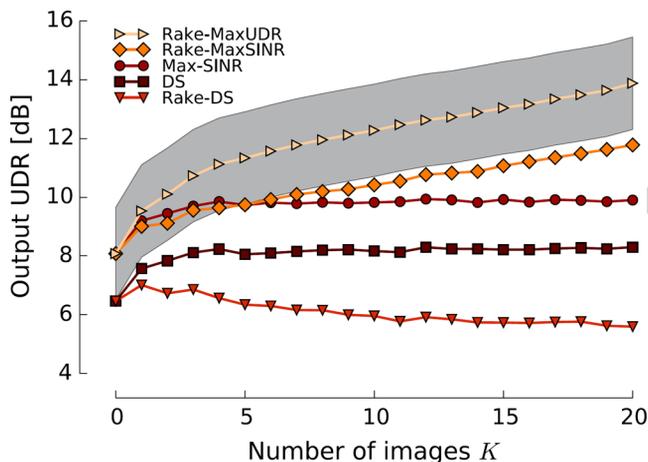


Fig. 9. Median output UDR plotted against the number of image sources used in the design for different beamformers, at a frequency $f = 1$ kHz. The shaded area contains the Rake-Max-UDR output UDRs for 50% of the 20000 Monte Carlo runs.

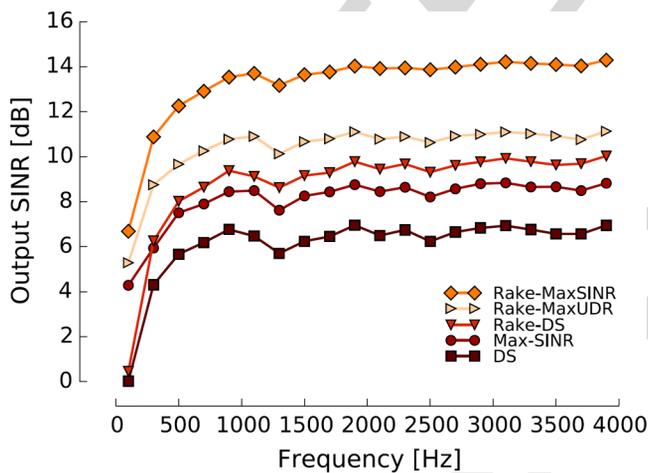


Fig. 10. Output SINR as a function of frequency for different beamformers and $K = K' = 10$. The curves show the average of 20000 runs, with averaging performed in the dB domain.

quality achieved by acoustic raking. We simulate a room with two sources—a desired source and an interferer—and compare the outputs of the Rake-DS, Rake-Max-SINR, and

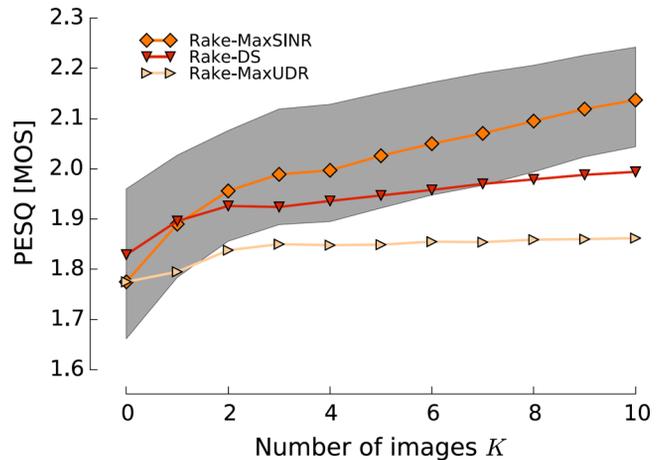


Fig. 11. Perceptual quality in MOS, evaluated using PESQ, as a function of the number of image sources used K . The lower limit of the ordinates is set to the median MOS of the degraded signal before processing, as measured at center of the array. The shaded area contains the Rake-Max-SINR output for 50% of the 10000 Monte Carlo runs.

Rake-Max-UDR as a function of the number of image sources used to design the beamformers.

The same number of image sources is used for the desired source and the interferer ($K = K'$). The performance metric used is the PESQ [22]. In particular, we use the reference implementation described by the ITU P.862 Amendment 2 [34]. PESQ compares the reference signal with the degraded signal and predicts the perceptual quality of the latter as it would be measured by the mean opinion score (MOS) value, on a scale from 1 to 4.5.

We consider the same room and microphone array setting as before (see Fig. 7(a)). The desired and the interfering sources are placed uniformly at random in a rectangular area with lower left corner at (1 m, 2.5 m) and upper right corner at (3 m, 5 m). To limit the experimental variation, the speech samples attributed to the sources are fixed throughout the simulation. The two sources start reproducing speech at the same time and approximately overlap for the total duration of the speech samples. The signals are normalized to have the same power at the source, and we add white Gaussian noise to the microphone signals, with power chosen so that the SNR of the direct sound for the desired source is 20 dB at the center of the microphone array. All signals are high-pass filtered with a cut-off frequency of 300 Hz. The reference for all PESQ results is the direct path of the desired source as measured at the center of the array (2 m, 1.5 m).

The median PESQ measure of 10000 Monte Carlo runs, given in raw MOS, is shown in Fig. 11. The median PESQ of the degraded signal measured at the center of the array before processing was found to be 1.6 raw MOS. When only the direct sound is used (*i.e.*, $K = 0$), all three beamformers yield the same improvement of about 0.2 raw MOS. We observe that Rake-DS is slightly outperforms the other beamformers. Using any number of echoes in addition to the direct sound results in larger MOS for all beamformers. When more than one image source is used, the Rake-Max-SINR beamformer always yields the largest MOS, with up to 0.5 MOS gain when using 10 images sources.

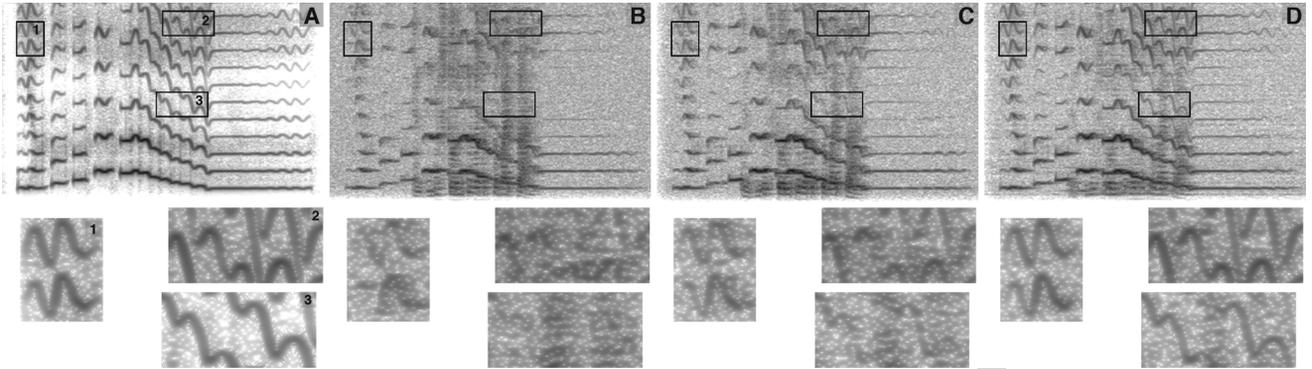


Fig. 12. Comparison of the conventional Max-SINR and Rake-Max-SINR beamformer on a real speech sample. Spectrograms of (A) clean signal of interest, (B) signal corrupted by an interferer and additive white Gaussian noise at the microphone input, outputs of (C) conventional Max-SINR and (D) Rake-Max-SINR beamformers. Time naturally goes from left to right, and frequency increases from zero at the bottom up to $F_s/2$. To highlight the improvement of Rake-Max-SINR over Max-SINR, we blow-up three parts of the spectrograms in the lower part of the figure. The boxes and the corresponding part of the original spectrogram are numbered in (A). The numbering is the same but omitted in the rest of the figure for clarity.

It is worth mentioning that in the beamformer design, we do not assume that we know the spectrum of the source or the interferer—we design as if it was flat. Thus the interferer acts as a strong source of colored, spatially correlated, non-stationary noise, spectrally mismatched with the designed beamformer. There is another source of model mismatch: while the RIRs were computed using hundreds of image sources, we use only up to ten to design the beamformers.

4) *Spectrograms and Sounds Samples*: Finally, we present the spectrograms for a scenario where we want to focus on a singer in the presence of interfering speech. We consider the same room, source, interferer, and microphone geometry as in Fig. 7(b).

The source signal is a passage by a female opera singer (Fig. 12(a)), with strongly pronounced harmonics; the interfering signal is a male speech extract. The two signals are normalized to have unit maximum amplitude. We add white Gaussian noise to the microphone signals with power such that the SNR of the direct sound of the desired source is 20 dB at the center of the microphone array. All signals are high-pass filtered with a cut-off frequency of 300 Hz. The Rake-Max-SINR beamformer weights are computed using the direct source and three generations of image sources for both the desired sound source (singing) and the interferer (speech).

The output of the conventional Max-SINR beamformer (Fig. 12(c)) is compared to that of the Rake-Max-SINR (Fig. 12(d)). We can observe from the spectrogram that the Rake-Max-SINR reduces very effectively the power of the interfering signal at all frequencies, but particularly in the mid to high range. This is true even when the interferer overlaps significantly with the desired signal. Informal listening tests confirm that the Rake-Max-SINR maintains high quality of the desired signal while strongly reducing the interference. The Rake-Max-UDR beamformer provides good interference suppression, but it produces audible pre-echoes that render it unsuitable for speech processing applications. The sound clips can be found online along the code.

VII. CONCLUSION

We investigated the concept of acoustic rake receivers—beamformers that use echoes. Unlike earlier related

work, we presented optimal formulations that outperform the delay-and-sum style approaches by a large margin. This is especially true in the presence of interferers, hence the title “Raking the Cocktail Party”. We demonstrate theoretically that the ARR improves the SINR, and the numerical simulations agree well with these predictions.

Beyond theoretical and numerical evaluations of the performance measures, we demonstrated in informal listening tests the improved interference suppression by the ARR. Moreover, extensive simulation determined that the ARR improves the subjective quality, as predicted by PESQ, proportionally to the number of image sources used. A particularly illustrative example is when the interferer is occluding the desired source—the optimal ARR takes care of this simply by listening to the echoes.

Perhaps the most important aspect of ongoing work is the design of robust formulations of ARRs. This may involve various heuristics, as well as combinatorial optimization due to the discrete nature of image sources. We expect that the raking beamformers described in this paper inherit the robustness properties of their classical counterparts. For example, the Rake-DS beamformer is likely to be more robust to array calibration errors than the Rake-Max-SINR beamformer. Furthermore, we expect that taking the image source perspective makes various ARRs more robust to errors in source locations than the schemes that assume the knowledge of the RIR.

Another line of ongoing work investigates the time-domain formulations of the ARRs, with some initial results already available [25]. Time-domain formulations offer better control over whether the echoes appear before or after the direct sound. This provides a more natural framework for optimizing perceptually motivated performance measures such as UDR.

APPENDIX THEOREM 1

We note that the theorem is stated for a linear array, but the described behavior is universal.

Theorem 1: Assume that there are $K + 1$ sources located at $\mathbf{s}_k = r_k[\cos \theta_k \sin \theta_k]^T$ where $\theta_k \sim \mathcal{U}(0, 2\pi)$ and $r_k \sim \mathcal{U}(a, b)$ are all independent, for some $0 < a < b$ such that the far-field assumption holds. Let \mathbf{A}_s collect the corresponding

TABLE II
SUMMARY OF BEAMFORMERS.

Acronym	Description	Beamforming Weights
DS	Align delayed copies of signal at the microphone	$\mathbf{w}_{\text{DS}} = \mathbf{a}(\mathbf{s})/\ \mathbf{a}(\mathbf{s})\ $
Max-SINR	$\max. \mathbf{w}^H \mathbf{a}_s \mathbf{a}_s^H \mathbf{w} / (\mathbf{w}^H \mathbf{K}_{nq} \mathbf{w})$	$\mathbf{w}_{\text{SINR}} = \mathbf{K}_{nq}^{-1} \mathbf{a}_s / (\mathbf{a}_s^H \mathbf{K}_{nq}^{-1} \mathbf{a}_s)$
Rake-DS	Weighted average of DS beamformers over image sources	$\mathbf{w}_{\text{R-DS}} = \mathbf{A}_s \mathbf{1} / \ \mathbf{A}_s \mathbf{1}\ $
Rake-OF	$\min. \mathbb{E} \left \sum_{k=0}^{K'} \mathbf{w}^H \mathbf{a}(\mathbf{q}_k) z + \mathbf{w}^H \mathbf{n} \right ^2, \text{ s.t. } \mathbf{w}^H \mathbf{A}_s = \mathbf{1}^T$	$\mathbf{w}_{\text{OF}} = \mathbf{K}_{nq}^{-1} \mathbf{A}_s (\mathbf{A}_s^H \mathbf{K}_{nq}^{-1} \mathbf{A}_s)^{-1} \mathbf{1}_M$
Rake-Max-SINR	$\max. \mathbb{E} \left \sum_{k=0}^K \mathbf{w}^H \mathbf{a}(\mathbf{s}_k) x \right ^2 / \mathbb{E} \left \sum_{k=0}^{K'} \mathbf{w}^H \mathbf{a}(\mathbf{q}_k) z + \mathbf{w}^H \mathbf{n} \right ^2$	$\mathbf{w}_{\text{R-SINR}} = \mathbf{K}_{nq}^{-1} \mathbf{A}_s \mathbf{1} / (\mathbf{1}^H \mathbf{A}_s^H \mathbf{K}_{nq}^{-1} \mathbf{A}_s \mathbf{1})$
Rake-Max-UDR	$\max. \mathbb{E} \sum_{k=0}^K \mathbf{w}^H \mathbf{a}(\mathbf{s}_k) x ^2 / \mathbb{E} \left \mathbf{w}^H \sum_{k=0}^{K'} \mathbf{a}(\mathbf{q}_k) z + \mathbf{w}^H \mathbf{n} \right ^2$	$\mathbf{w}_{\text{R-UDR}} = \mathbf{C}^{-1} \tilde{\mathbf{w}}_{\max} ((\mathbf{C}^{-1})^H \mathbf{A}_s \mathbf{A}_s^H \mathbf{C}^{-1})$

TABLE I
[AUTHOR: Please cite Tables 1 and 2 in the paper] SUMMARY OF NOTATION.

Symbol	Meaning
M	Number of microphones
\mathbf{r}_m	Location of the m th microphone
\mathbf{s}_0	Location of the desired source
\mathbf{s}_i	Location of the i th image of the desired source ($i \geq 1$)
\mathbf{q}_0	Location of the interfering source
\mathbf{q}_i	Location of the i th image of the interfering source ($i \geq 1$)
$x(e^{j\omega})$	Spectrum of the sound from the desired source
$z(e^{j\omega})$	Spectrum of the sound from the interfering source
$\mathbf{w}(e^{j\omega})$	Vector of beamformer weights
K	Number of considered desired image sources
K'	Number of considered interfering image sources
$a_m(\mathbf{s})$	m th component of the steering vector for a source at \mathbf{s}
y_m	Signal picked up by the m th microphone
$\ \cdot\ $	Euclidean norm, $\ \mathbf{x}\ = (\sum x_i ^2)^{1/2}$.

steering vectors for a uniform linear microphone array. Then $\mathbb{E} \|\mathbf{A}_s \mathbf{1}\|^2 \geq (1 + \beta) \mathbb{E} \|\mathbf{a}(\mathbf{s}_0)\|^2$, where $\beta = \sum_{k=1}^K (\alpha_k / \alpha_0)^2$, and α_k are attenuations of the steering vectors, assumed independent from the source locations. In fact, $\mathbb{E} (\|\mathbf{A}_s \mathbf{1}\|^2) = (1 + \beta) \mathbb{E} (\|\mathbf{a}(\mathbf{s}_0)\|^2) + O(1/\Omega^3)$.

Proof: Thanks to the far-field assumption, we can decompose the steering vector into a factor due to the array, and a phase factor due to different distances of different image sources. We have that

$$a_m = (\mathbf{A}_s \mathbf{1})_m = \sum_{k=0}^K \alpha_k e^{-j\kappa m d \sin \theta_k} e^{-j\Omega \delta_k / c}, \quad (33)$$

where d is the microphone spacing and $\kappa \stackrel{\text{def}}{=} \Omega / c$. Without loss of generality we assume that $\delta_k \sim \mathcal{U}(a, b)$. We can further write

$$\begin{aligned} \mathbb{E} |a_m|^2 &= \mathbb{E} \left[\left(\sum_{k=0}^K \alpha_k e^{-j\kappa m d \sin \theta_k} e^{-j\kappa \delta_k} \right) \right. \\ &\quad \times \left. \left(\sum_{\ell=0}^K \alpha_\ell e^{j\kappa m d \sin \theta_\ell} e^{j\kappa \delta_\ell} \right) \right] \\ &= \sum_{k=0}^K \alpha_k^2 + \sum_{k \neq \ell=0}^K \alpha_k \alpha_\ell \mathbb{E} \left[e^{j\kappa m d (\sin \theta_\ell - \sin \theta_k)} e^{j\kappa (\delta_\ell - \delta_k)} \right]. \end{aligned} \quad (34)$$

Invoking the independence for $k \neq \ell$, we compute the above expectation as

$$\begin{aligned} &\mathbb{E} \left[e^{j\kappa m d (\sin \theta_\ell - \sin \theta_k)} e^{j\kappa (\delta_\ell - \delta_k)} \right] \\ &= \frac{2J_0^2(m d \kappa) [1 - \cos(\Delta \kappa)]}{(\Delta \kappa)^2}, \quad (35) \end{aligned}$$

where J_0 denotes the Bessel function of the first kind and zeroth order and $\Delta \stackrel{\text{def}}{=} b - a$.

Plugging this back into (34), we obtain

$$\mathbb{E} |a_m|^2 = \sum_{k=0}^K \alpha_k^2 \left(1 + C \frac{2J_0^2(m d \kappa) [1 - \cos(\Delta \kappa)]}{(\Delta \kappa)^2} \right), \quad (36)$$

where $C = \sum_{k \neq \ell} \alpha_k \alpha_\ell / \sum_k \alpha_k^2$.

Because $|J_0(z)| \leq \sqrt{2/(\pi z)} + O(|z|^{-1})$ ([35], Eq. 9.2.1), we see that the expression in brackets is $1 + O(\Omega^{-3})$. Rewriting

$$\sum_{k=0}^K \alpha_k^2 = \frac{1}{M} \mathbb{E} \|\mathbf{a}(\mathbf{s}_0)\|^2 \left(1 + \sum_{k=1}^K \left(\frac{\alpha_k}{\alpha_0} \right)^2 \right) \quad (37)$$

concludes the proof. \square

ACKNOWLEDGMENT

We thank the anonymous reviewers for a number of constructive suggestions that have improved the manuscript.

REFERENCES

- [1] R. Price and P. E. Green, "A communication technique for multipath channels," *Proc. IRE. IEEE*, vol. 46, no. 3, pp. 555–570, Mar. 1958.
- [2] A. F. Naguib, "Space-time receivers for CDMA multipath signals," in *Proc. IEEE ICC. IEEE*, Montreal, QC, Canada, 1997, pp. 304–308.
- [3] B. H. Khalaj, A. Paulraj, and T. Kailath, "2D RAKE receivers for CDMA cellular systems," in *Proc. IEEE GLOBECOM*, 1994, pp. 400–404.
- [4] P. Annibale, F. Antonacci, P. Bestagini, A. Brutti, A. Canclini, L. Cristoforetti, J. Filos, E. Habets, W. Kellerman, K. Kowalczyk, A. Lombard, E. Mabande, D. Markovic, P. Naylor, and M. Omologo, "The SCENIC project: Space-time audio processing for environment-aware acoustic sensing and rendering," in *Proc. 131st Conv. Audio Eng. Soc.*, New York, NY, USA, 2011, Audio Engineering Society.
- [5] A. E. O'Donovan, R. Duraiswami, and D. N. Zotkin, "Automatic matched filter recovery via the audio camera," in *Proc. IEEE ICASSP. IEEE*, Dallas, TX, USA, 2010, pp. 2826–2829.
- [6] E.-E. Jan, P. Svaizer, and J. L. Flanagan, "Matched-filter processing of microphone array for spatial volume selectivity," in *Proc. IEEE ISCAS*, 1995, vol. 2, pp. 1460–1463.

- [7] J. B. Allen and D. A. Berkley, "Image method for efficiently simulating small-room acoustics," *J. Acoust. Soc. Amer.*, vol. 65, no. 4, pp. 943–950, 1979.
- [8] J. Borish, "Extension of the image model to arbitrary polyhedra," *J. Acoust. Soc. Amer.*, vol. 75, no. 6, pp. 1827–1836, 1984.
- [9] D. B. Ward, E. A. Lehmann, and R. C. S. Williamson, "Particle filtering algorithms for tracking an acoustic source in a reverberant environment," *IEEE Trans. Audio, Speech, Lang. Process.*, vol. 11, no. 6, pp. 826–836, Nov. 2003.
- [10] F. Ribeiro, D. A. Florencio, D. E. Ba, and C. Zhang, "Geometrically constrained room modeling with compact microphone arrays," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 20, no. 5, pp. 1449–1460, Jul. 2012.
- [11] F. Antonacci, J. Filos, M. R. P. Thomas, E. A. P. Habets, A. Sarti, P. A. Naylor, and S. Tubaro, "Inference of room geometry from acoustic impulse responses," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 20, no. 10, pp. 2683–2695, Oct. 2012.
- [12] I. Dokmanic, Y. M. Lu, and M. Vetterli, "Can one hear the shape of a room: The 2-D polygonal case," in *Proc. IEEE ICASSP*, Prague, Czech Republic, 2011, pp. 321–324.
- [13] **[AUTHOR: Please provide page range]** I. Dokmanic, R. Parhizkar, A. Walther, Y. M. Lu, and M. Vetterli, "Acoustic echoes reveal room shape," *Proc. Nat. Acad. Sci.*, vol. 110, no. 30, Jun. 2013.
- [14] J. S. Bradley, H. Sato, and M. Picard, "On the importance of early reflections for speech in rooms," *J. Acoust. Soc. Amer.*, vol. 113, no. 6, p. 3233, 2003.
- [15] J. Lochner and J. F. Burger, "The influence of reflections on auditorium acoustics," *J. Sound Vibr.*, vol. 1, no. 4, pp. 426–454, 1964.
- [16] E. Habets, J. Benesty, I. Cohen, S. Gannot, and J. Dmochowski, "New insights into the MVDR beamformer in room acoustics," *IEEE Trans. Audio, Speech, Lang. Process.*, vol. 18, no. 1, pp. 158–170, Jan. 2010.
- [17] M. R. P. Thomas, N. D. Gaubitch, and P. A. Naylor, "Application of channel shortening to acoustic channel equalization in the presence of noise and estimation error," in *Proc. IEEE WASPAA*, New Paltz, NY, USA, 2011, pp. 113–116.
- [18] W. Zhang, E. Habets, and P. A. Naylor, "On the use of channel shortening in multichannel acoustic system equalization," in *Proc. IWAENC*, Tel Aviv, Israel, 2010.
- [19] M. Miyoshi and Y. Kaneda, "Inverse filtering of room acoustics," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 36, no. 2, pp. 145–152, Feb. 1988.
- [20] K. Furuya, "Noise reduction and dereverberation using correlation matrix based on the multiple-input/output inverse-filtering theorem (MINT)," in *Proc. Int. Workshop on HSC*, Kyoto, Japan, 2001, pp. 59–62.
- [21] J. Benesty, J. Chen, Y. A. Huang, and J. Dmochowski, "On microphone-array beamforming from a MIMO acoustic signal processing perspective," *IEEE Trans., Audio, Speech, Lang. Process.*, vol. 15, no. 3, pp. 1053–1065, Mar. 2007.
- [22] A. W. Rix, J. G. Beerends, M. P. Hollier, and A. P. Hekstra, "Perceptual evaluation of speech quality (PESQ)—A new method for speech quality assessment of telephone networks and codecs," in *Proc. IEEE ICASSP*, Salt Lake City, UT, USA, 2001, pp. 749–752.
- [23] S. Haykin and Z. Chen, "The cocktail party problem," *Neural Comput.*, vol. 17, no. 9, pp. 1875–1902, 2005.
- [24] W. Herbordt and W. Kellermann, "Adaptive beamforming for audio signal acquisition," in *Adaptive Signal Processing*. Berlin/Heidelberg, Germany: Springer, 2003, pp. 155–194.
- [25] **[AUTHOR: Please update status]** R. Scheibler, I. Dokmanic, R., and M. Vetterli, "Raking echoes in the time domain," in *Proc. IEEE ICASSP*, 2015, accepted for publication.
- [26] B. D. Carlson, "Covariance matrix estimation errors and diagonal loading in adaptive arrays," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 24, no. 4, pp. 397–401, Jul. 1988.
- [27] T. E. Oliphant, "Python for scientific computing," *Comput. Sci. Eng.*, vol. 9, no. 3, pp. 10–20, 2007.
- [28] D. Duffy, *Green's Functions With Applications*. Boca Raton, FL, USA: Chapman & Hall/CRC, 2001.
- [29] B. D. Van Veen and K. M. Buckley, "Beamforming: A versatile approach to spatial filtering," *IEEE ASSP Mag.*, vol. 5, no. 2, pp. 4–24, Apr. 1988.
- [30] O. L. I. Frost, "An algorithm for linearly constrained adaptive array processing," in *Proc. IEEE. IEEE*, 1972, pp. 926–935.
- [31] W. McCrea and F. Whipple, "Random paths in two and three dimensions," in *Proc. Roy. Soc., Edinburgh*, U.K., 1940, vol. 60, pp. 281–298.
- [32] O. Öçal, I. Dokmanic, and M. Vetterli, "Source localization and tracking in non-convex rooms," in *Proc. IEEE ICASSP*, 2014, pp. 1429–1433.
- [33] J. J. Shynk, "Frequency-domain and multirate adaptive filtering," *IEEE Signal Process. Mag.*, vol. 9, no. 1, pp. 14–37, Jan. 1992.
- [34] "ITU-T P.862 Amendment 2, Reference Implementations and Conformance Testing for ITU-T Recs P.862, P.862.1 and P.862.2," 11, 2005, p. 862.
- [35] M. Abramowitz and I. A. Stegun, *Handbook of Mathematical Functions*. New York, NY, USA: Courier Dover, 1972.



ICASSP 2011. In 2014 he received a Google Ph.D. Fellowship.



acoustic beamforming.



Ivan Dokmanic is a Ph.D. candidate in Audiovisual Communications Laboratory (LCAV) at EPFL (expected graduation in May 2015). His interests include inverse problems, audio and acoustics, signal processing for sensor arrays/networks, and fundamental signal processing. Previously he was a Teaching Assistant at the University of Zagreb, a Codec Developer for MainConcept AG, and a Digital Audio Effects Designer for Little Endian Ltd. For his work on room shape reconstruction using sound he received the best student paper award at

Robin Scheibler (S'07) received the B.Sc. and M.Sc. in communications systems from Ecole Polytechnique Fédérale de Lausanne (EPFL), Switzerland, in 2006 and 2009, respectively. From 2009 to 2010, he was with IBM Research–Zürich. After this, he was with the NEC Media Information Processing group in Kawasaki, Japan, from 2011 to 2012. He then returned to EPFL where he is working towards the Ph.D. degree at the Laboratory for Audiovisual Communications. His research interests lie in algorithmic signal processing, speech processing, and

Martin Vetterli received his engineering degree from ETH Zurich, his M.S. from Stanford, and his Doctorate from EPF Lausanne. He was on the faculty of Columbia University and UC Berkeley before joining EPFL. He is currently President of the National Research Council of the Swiss NSF. He works on signal processing and communications, e.g., wavelet theory and applications, sparse sampling, joint source–channel coding and sensor networks. He received the best paper awards of IEEE Signal Processing Society (1991, 1996, 2006), and the IEEE Signal Processing Technical and Society Awards (2002, 2010). He is a Fellow of IEEE, ACM, and EURASIP. He is the coauthor of the books *Wavelets and Subband Coding* (1995), *Signal Processing for Communications* (2008), and *Foundations of Signal Processing* (2014). He is a highly cited researcher in engineering (Thomson ISI Web of Science, 2007, 2014).

Raking the Cocktail Party

Ivan Dokmanic, *Student Member, IEEE*, Robin Scheibler, *Student Member, IEEE*, and Martin Vetterli, *Fellow, IEEE*

Abstract—We present the concept of an acoustic rake receiver—a microphone beamformer that uses echoes to improve the noise and interference suppression. The rake idea is well-known in wireless communications; it involves constructively combining different multipath components that arrive at the receiver antennas. Unlike spread-spectrum signals used in wireless communications, speech signals are not orthogonal to their shifts. Therefore, we focus on the spatial structure, rather than the temporal. Instead of explicitly estimating the channel, we create correspondences between early echoes in time and image sources in space. These multiple sources of the desired and the interfering signal offer additional spatial diversity that we can exploit in the beamformer design. We present several “intuitive” and optimal formulations of acoustic rake receivers, and show theoretically and numerically that the rake formulation of the maximum signal-to-interference-and-noise beamformer offers significant performance boosts in terms of noise and interference suppression. Beyond signal-to-noise ratio, we observe gains in terms of the *perceptual evaluation of speech quality* (PESQ) metric for the speech quality. We accompany the paper by the complete simulation and processing chain written in Python. The code and the sound samples are available online at <http://lcav.github.io/AcousticRakeReceiver/>.

Index Terms—Acoustic rake receiver, beamforming, echo sorting, interference cancellation, noise suppression, perceptual evaluation of speech quality (PESQ), room impulse response.

I. INTRODUCTION

RAKE receivers take advantage of the multipath propagation, instead of trying to mitigate it. The basic idea of the rake receivers, which are commonly used in wireless communications, is to coherently add the multipath components and thus increase the effective signal-to-noise ratio (SNR). The original scheme was developed for single-input-single-output systems [1], and it was later extended to arrays of antennas [2], [3] that exploit spatial diversity. Thus multipath components that are not resolvable with a single antenna because they arrive at similar times, become resolvable because they arrive from different directions.

In spite of the success of the rake receivers in wireless communications, the principle has not received significant attention in room acoustics. Nevertheless, constructive use of echoes in

rooms to improve beamforming has been mentioned in the literature [4]–[6]. In particular, the term *acoustic rake receiver* (ARR) was used in the SCENIC project proposal [4].

The list of ingredients for ARR in room acoustics is similar as in wireless communications: a wave (acoustic instead of electromagnetic) propagates in space; reflections and scattering cause the wave to arrive at the receiver through multiple paths in addition to the direct path, and these multipath components all contain the source waveform.

The main difference is that in room acoustics we do not get to design the input signal. Spreading sequences used in CDMA are designed to be near-orthogonal to their shifts and orthogonal between different users, which facilitates the multipath channel estimation; such orthogonality is not exhibited by speech. Moreover, speech segments are very long with respect to the time delay between two consecutive echoes, and they are *a priori* unknown at the receiver.

On the contrary, there are no significant differences in terms of the spatial structure. If we know where the echoes are coming from, we can design spatial processing algorithms—for example beamformers—that use multiple copies of the same signal arriving from different directions.

Imagine first that we know the room geometry. Then, if we localize the source, we can predict where its echoes will come from using simple geometric rules [7], [8]. Localizing the direct signal in a reverberant environment is a well-understood problem [9]. What is more, we do not need to know the room shape in detail—locations of the most important reflectors (ceiling, floor, walls) suffice to localize the major echoes. In many cases this knowledge is readily available from the floor plans or measurements. In ad-hoc deployments, the room geometry may be difficult to obtain. If that is the case, we can first perform a calibration step to learn it. An appealing method to infer the room geometry is by using sound, as was demonstrated recently [10]–[13].

We may still be able to take advantage of the echoes without estimating the room geometry. Note that we are not after the room geometry itself; rather, we only need to know where the early echoes are coming from. Echoes can be seen as signals emitted by *image sources*—mirror images of the true source across reflecting walls [7]. Knowing where the echoes are coming from is equivalent to knowing where the image sources are.

Image source localization can be solved, for example, by *echo sorting* as described in [13]. Alternatively, O'Donovan, Duraiswami and Zotkin [5] propose to use an *audio camera* with a large number of microphones to find the images. Once the image sources are localized (in a calibration phase or otherwise), we can predict their movement using geometrical rules, as discussed in Section V. Thus, the acoustic raking is

Manuscript received July 15, 2014; revised January 15, 2015; accepted March 06, 2015. Date of publication March 23, 2015; date of current version nulldate. This work was supported by the ERC Advanced Investigators Grant: Sparse Sampling: Theory, Algorithms, and Applications SPARSAM no. 247006, and a Google Doctoral Fellowship. The guest editor coordinating the review of this manuscript and approving it for publication was XXXXX XXXXXX.

The authors are with the LCAV-EPFL, 1015 Lausanne, Switzerland (e-mail: ivan.dokmanic@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/JSTSP.2015.2415761

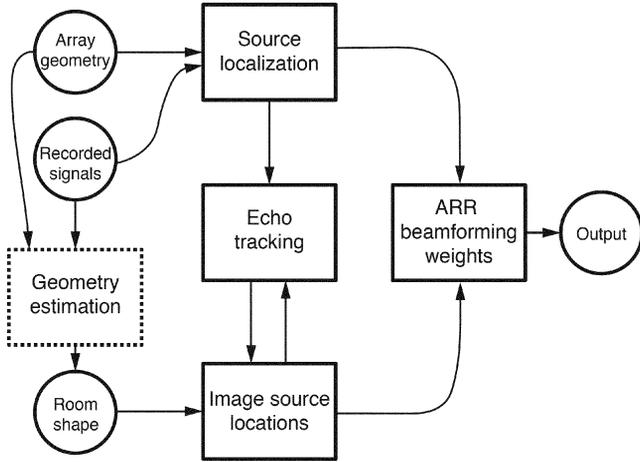


Fig. 1. A block diagram for acoustic rake receivers. In this paper, we focus on ARR beamforming weight computation, and we briefly discuss echo tracking and image source localization. The geometry estimation block is optional (room geometry could be known in advance), hence the dashed box.

a multi-stage process comprising image source localization, image source tracking, and beamforming weight computation. The complete block diagram of an acoustic rake receiver is shown in Fig. 1.

It is interesting to note the analogy between the ARR and the human auditory perception. It is well established that the early echoes improve speech intelligibility [14], [15]. In fact, adding energy in the form of early echoes (approximately within the first 50 ms of the room impulse response (RIR)) is equivalent to adding the same energy to the direct sound [14]. This observation suggests new designs for indoor beamformers, with different choices of performance measures and reference signals. A related discussion of this topic is given by Habets and co-authors [16], who examine the tradeoff between dereverberation and denoising in beamforming. In addition to the standard SNR, we propose to use the useful-to-detrimental ratio (UDR), first defined by Lochner and Burger [15], and used by Bradley, Sato and Picard [14]. We generalize UDR to a scenario with interferers, defining it as the ratio of the direct and early reflection energy to the energy of the noise and interference.

ARRs focus on the early part of the RIR, trying to concentrate the energy contained in the early echoes. In that regard, there are similarities between ARR and channel shortening [17], [18]. Channel shortening produces filters that are much better behaved than complete inversion, e.g., by the multiple-input-output-theorem (MINT) [19], [20]. Nevertheless, it is still tacitly assumed that we know the acoustic impulse responses between the sources and the microphones. In contrast to channel shortening, as well as other methods assuming this knowledge [19], [21], we never attempt the difficult task of estimating the impulse responses. Our task is simpler: we only need to detect the early echoes, and lift them to 3D space as image sources.

A. Main Contributions and Limitations

We introduce the acoustic rake receiver (ARR) as the echo-aware microphone beamformer. We present several formulations with different properties, and analyze their behavior theoretically and numerically. The analysis shows that ARR

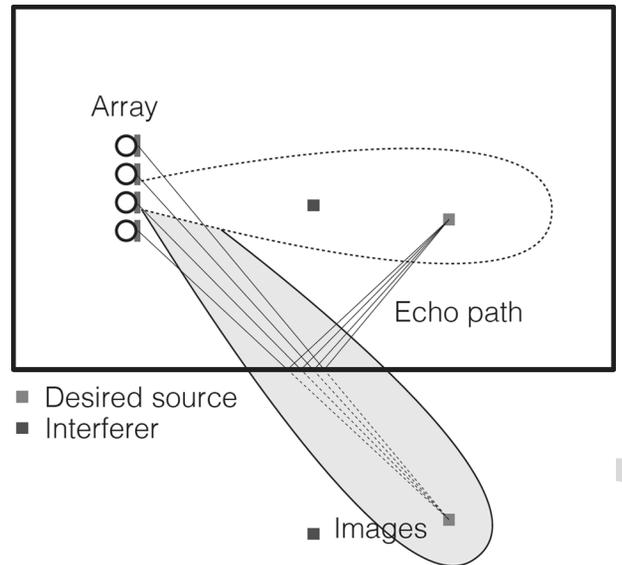


Fig. 2. Listening behind an interferer by listening to echoes (illustration). A beam directed towards the desired, green source will necessarily pick up the red interferer's signal. Acoustic rake receiver avoids it by beamforming towards the echoes of the desired source.

lead to significantly improved SNR and interference cancellation when compared with standard beamformers that only extract the direct path. ARR can suppress interference in cases when conventional beamforming is bound to fail, for example when an interferer is occluding the desired source (an illustration is given in Fig. 2; for a sneak-peak of real beam patterns, fast-forward to Fig. 7). We present optimal formulations that outperform the earlier delay-and-sum (DS) approaches [6], especially when interferers are present. Significant gains are observed not only in terms of signal-to-interference-and-noise ratio (SINR) and UDR, but also in terms of perceptual evaluation of speech quality (PESQ) [22].

The raking microphone beamformers are particularly well-suited to extracting the desired speech signal in the presence of interfering sounds, in part because they can focus on echoes of the desired sound and cancel the echoes of the interfering signals. The analogous human capacity to focus on a particular acoustic stimulus while not perceiving other, unwanted sounds is called the *cocktail party effect* [23]. This paper's title was inspired by that analogy.

We design and apply the ARR in the frequency domain. Frequency domain formulation is simple and concise; it allows us to focus on objective gains from acoustic raking. Time-domain designs [24] offer better control over the impulse responses of the beamforming filters, but they are out of the scope of this paper. For a recent time-domain approach to ARR, see [25].

Let us also mention some limitations of our results. For clarity, the numerical experiments are presented in a 2D "room", and as such are directly applicable to planar (e.g., linear or circular) arrays. Extension to 3D arrays is straightforward. We do not discuss robust formulations that address uncertainties in the array calibration. Microphones are assumed to be ideally omni-directional with a flat frequency response. Except for Section V, we assume that the locations of the image sources are known. We explain how to find the image sources

when the room geometry is either known or unknown, but we do not provide a deep overview of the geometry estimation techniques. To this end, we suggest a number of references for the interested reader. We consider the walls to be flat-fading; in reality, they are frequency selective. We do not discuss the estimation of various covariance matrices [26].

The results in this paper are reproducible. Python (NumPy) [27] code for the beamforming routines, for the STFT processing engine, and to generate the figures and the sound samples is available online at <http://lcav.github.io/AcousticRakeReceiver/>.

B. Paper Outline

In Section II we explain the notation and the signal model used in the paper. A brief overview of the relevant beamforming techniques and performance analysis is given in Section III. We formulate the acoustic rake receiver in Section IV, and we present a theoretical and numerical analysis of the corresponding beamformers. Section V explains how to locate the image sources, and comments on localizing the direct source. Numerical experiments are presented in Section VI.

II. NOTATION AND SIGNAL MODEL

We denote all matrices by bold uppercase letters, for example \mathbf{A} , and all vectors by bold lowercase letters, for example \mathbf{x} . The Hermitian transpose of a matrix or a vector is denoted by $(\cdot)^H$, as in \mathbf{A}^H , and the Euclidean norm of a vector by $\|\cdot\|$, that is, $\|\mathbf{x}\| \stackrel{\text{def}}{=} (\mathbf{x}^H \mathbf{x})^{1/2}$.

Suppose that the desired source of sound is at the location \mathbf{s}_0 in a room. Sound from this source arrives at the microphones located at $\{\mathbf{r}_m\}_{m=1}^M$ via the direct path, but also via the echoes from the walls. The echoes can be replaced by the image sources—mirror images of the true sources across the corresponding walls—according to the image source model [7], [8]. An important consequence is that instead of modeling the source of the desired or the interfering signal as a single point in a room, we can model it as a collection of points in free space. A more detailed discussion of the image source model is given in Section V.

Denote the signal emitted by the source $\tilde{x}[n]$ (e.g., the speech signal). Then all the image sources emit $\tilde{x}[n]$ as well, and the signals from the image sources reach the microphones with the appropriate delays. In our application, the essential fact is that the echoes correspond to image sources. We denote the image source positions by \mathbf{s}_k , $1 \leq k \leq K$, where K denotes the largest number of image sources considered. Note that we do not care about the sequence of walls that generates \mathbf{s}_k , nor do we care about how many walls are in the sequence. For us, all \mathbf{s}_k are simply additional sources of the desired signal. The described setup is illustrated Fig. 3.

Suppose further that there is an interferer at the location \mathbf{q}_0 (for simplicity, we consider only a single interferer). The interferer emits the signal $\tilde{z}[n]$, and its image sources emit $\tilde{z}[n]$ as well. Similarly as for the desired source, we denote by \mathbf{q}_k , $1 \leq k \leq K'$ the positions of the interfering image sources, with K' being the largest number of interfering image source considered. The model mismatch (e.g., the image sources of high orders and the late reverberation) and the noise are absorbed in the term $\tilde{n}_m[n]$.

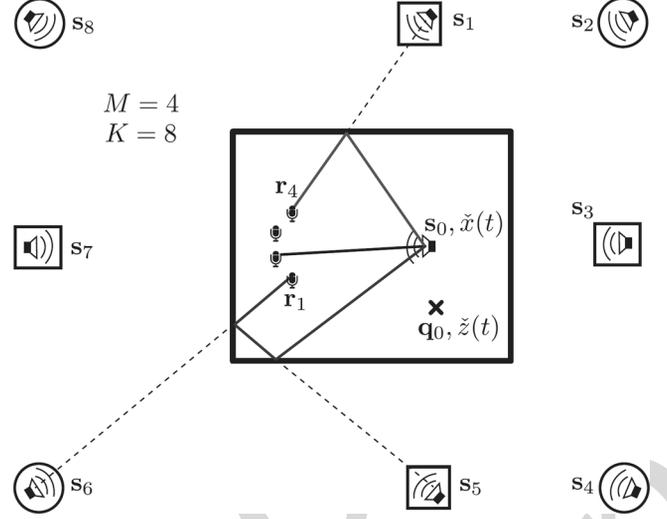


Fig. 3. Illustration of the notation and concepts. Echoes of the desired signal emitted at \mathbf{s}_0 can be modeled as a direct sound coming from the image sources of \mathbf{s}_0 . Two generations of image sources are illustrated: first ($\mathbf{s}_1, \mathbf{s}_3, \mathbf{s}_5, \mathbf{s}_7$) and second ($\mathbf{s}_2, \mathbf{s}_4, \mathbf{s}_6, \mathbf{s}_8$), as well as the corresponding sound rays for \mathbf{s}_5 and \mathbf{s}_6 . The interferer is located at \mathbf{q}_0 (its image sources are not shown), and the microphones are located at $\mathbf{r}_1, \dots, \mathbf{r}_4$.

The signal received by the m th microphone is then a sum of convolutions

$$\tilde{y}_m[n] = \sum_{k=0}^K (\tilde{a}_m(\mathbf{s}_k) * \tilde{x})[n] + \sum_{k=0}^{K'} (\tilde{a}_m(\mathbf{q}_k) * \tilde{z})[n] + \tilde{n}_m[n], \quad (1)$$

where $\tilde{a}_m(\mathbf{s}_k)$ denotes the impulse response of the channel between the source located at \mathbf{s}_k and the m th microphone—in this case a delay and a scaling factor.

We design and analyze all the beamformers in the frequency domain. That is, we will be working with the DTFT of the discrete-time signal \tilde{x} ,

$$x(e^{j\omega}) \stackrel{\text{def}}{=} \sum_{n \in \mathbb{Z}} \tilde{x}[n] e^{-j\omega n}. \quad (2)$$

In practical implementations, we use the discrete-time short-time Fourier transform (STFT). More implementation details are given in Section VI.

Using these notations, we can write the signal picked up by the m th microphone as

$$y_m(e^{j\omega}) = \sum_{k=0}^K a_m(\mathbf{s}_k, \Omega) x(e^{j\omega}) + \sum_{k=0}^{K'} a_m(\mathbf{q}_k, \Omega) z(e^{j\omega}) + n_m(e^{j\omega}), \quad (3)$$

where $n_m(e^{j\omega})$ models the noise and other errors, and $a_m(\mathbf{s}_k, \Omega)$ denotes the m th component of the steering vector for the source \mathbf{s}_k . The steering vector is the Fourier transform of the continuous version of the impulse response $\tilde{a}(\mathbf{s}_k)$, evaluated at the frequency Ω . The discrete-time frequency ω and the continuous-time frequency Ω are related as $\omega = \Omega T_s$, where

T_s is the sampling period. The steering vector is then simply $\mathbf{a}(\mathbf{s}_k, \Omega) = [a_m(\mathbf{s}_k, \Omega)]_{m=0}^{M-1}$.

We can write out the entries of the steering vectors explicitly for a point source in free space. They are given as the appropriately scaled free-space Green's functions for the Helmholtz equation [28],

$$a_m(\mathbf{s}_k, \Omega) = \frac{\alpha_k}{4\pi\|\mathbf{r}_m - \mathbf{s}_k\|} e^{-j\kappa\|\mathbf{r}_m - \mathbf{s}_k\|}, \quad (4)$$

where we define the wavenumber as $\kappa \stackrel{\text{def}}{=} \Omega/c$, and α_k is the attenuation corresponding to \mathbf{s}_k .

Using vector notation, the microphone signals can be written concisely as

$$\mathbf{y}(e^{j\omega}) = \mathbf{A}_s(e^{j\omega})\mathbf{1}x(e^{j\omega}) + \mathbf{A}_q(e^{j\omega})\mathbf{1}z(e^{j\omega}) + \mathbf{n}(e^{j\omega}), \quad (5)$$

where $\mathbf{A}_s(e^{j\omega}) \stackrel{\text{def}}{=} [\mathbf{a}(\mathbf{s}_1, \Omega), \dots, \mathbf{a}(\mathbf{s}_K, \Omega)]$, $\mathbf{A}_q(e^{j\omega}) \stackrel{\text{def}}{=} [\mathbf{a}(\mathbf{q}_1, \Omega), \dots, \mathbf{a}(\mathbf{q}_{K'}, \Omega)]$, and $\mathbf{1}$ is the all-ones vector. From here onward, we suppress the frequency dependency of the steering vectors and the beamforming weights to reduce the notational clutter.

III. BEAMFORMING PRELIMINARIES

Microphone beamformers combine the outputs of multiple microphones in order to achieve spatial selectivity, thereby suppressing noise and interference [29]. In the frequency domain, a beamformer forms a linear combination of the microphone outputs to yield the output u . That is,

$$u = \mathbf{w}^H \mathbf{y} = \mathbf{w}^H \mathbf{A}_s \mathbf{1}x + \mathbf{w}^H \mathbf{A}_q \mathbf{1}z + \mathbf{w}^H \mathbf{n}, \quad (6)$$

where the vector $\mathbf{w} \in \mathbb{C}^M$ contains the beamforming weights.

The weights \mathbf{w} are often selected so that they optimize some design criterion. Common examples of beamformers are the delay-and-sum (DS) beamformer, minimum-variance-distortionless-response (MVDR) beamformer, maximum-signal-to-interference-and-noise (Max-SINR) beamformer, and minimum-mean-squared-error (MMSE) beamformer. In this paper we discuss the rake formulation of the DS and the Max-SINR beamformers; for completeness, we first describe the non-raking variants.

A. Delay-and-Sum Beamformer

DS is the simplest and often quite effective beamformer [29]. Assume that we want to listen to a source at \mathbf{s} . Then we form the DS beamformer by compensating the propagation delays from the source \mathbf{s} to the microphones \mathbf{r}_m ,

$$u_{\text{DS}} = \frac{1}{M} \sum_{m=0}^{M-1} y_m e^{j\kappa\|\mathbf{r}_m - \mathbf{s}\|} \quad (7)$$

$$\approx \frac{x}{4\pi\|\bar{\mathbf{r}} - \mathbf{s}\|} + \frac{1}{M} \sum_{m=0}^{M-1} n_m, \quad (8)$$

where $\bar{\mathbf{r}} = 1/M \sum_{m=0}^{M-1} \mathbf{r}_m$ denotes the center of the array. The beamforming weights can be read out from (7) as

$$\mathbf{w}_{\text{DS}} = \frac{\mathbf{a}(\mathbf{s})}{\|\mathbf{a}(\mathbf{s})\|}, \quad (9)$$

where we used the definition of y_m (5) and the definition of the steering vector (4). We can see from (8) that if $\mathbf{n} \sim \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I}_M)$, then the output noise is distributed according to $\mathcal{N}(0, \sigma^2/M)$, that is, we obtain an M -fold decrease in the noise variance at the output with respect to any reference microphone.

B. Maximum Signal-to-Interference-and-Noise Ratio Beamformer

The SINR is an important figure of merit used to assess the performance of ARR, and to compare it with the standard non-raking beamformers. It is computed as the ratio of the power of the desired output signal to the power of the undesired output signal. The desired output signal is the output signal due to the desired source, while the undesired signal is the output signal due to the interferers and noise.

For a desired source at \mathbf{s} and an interfering source at \mathbf{q} we can write

$$\text{SINR} \stackrel{\text{def}}{=} \frac{\mathbb{E}|\mathbf{w}^H \mathbf{a}(\mathbf{s})x|^2}{\mathbb{E}|\mathbf{w}^H (\mathbf{a}(\mathbf{q})z + \mathbf{n})|^2} = \sigma_x^2 \frac{\mathbf{w}^H \mathbf{a}(\mathbf{s}) \mathbf{a}(\mathbf{s})^H \mathbf{w}}{\mathbf{w}^H \mathbf{K}_{nq} \mathbf{w}}, \quad (10)$$

where \mathbf{K}_{nq} is the covariance matrix of the noise and the interference.

It is compelling to pick \mathbf{w} that maximizes the SINR (10) [29]. The maximization can be solved by noting that the rescaling of the beamformer weights leaves the SINR unchanged. This means that we can minimize the denominator subject to numerator being an arbitrary constant. The solution is given as

$$\mathbf{w}_{\text{SINR}} = \frac{\mathbf{K}_{nq}^{-1} \mathbf{a}_s}{\mathbf{a}_s^H \mathbf{K}_{nq}^{-1} \mathbf{a}_s}. \quad (11)$$

Using the definition (10), we can derive the SINR for the Max-SINR beamformer as

$$\text{SINR} = \sigma_x^2 \mathbf{a}_s^H \mathbf{K}_{nq}^{-1} \mathbf{a}_s. \quad (12)$$

Because \mathbf{K}_{nq}^{-1} is a Hermitian symmetric positive definite matrix, it has an eigenvalue decomposition as $\mathbf{K}_{nq}^{-1} = \mathbf{U}^H \mathbf{\Lambda} \mathbf{U}$, where \mathbf{U} is unitary, and $\mathbf{\Lambda}$ is diagonal with positive entries. We can write $\mathbf{a}_s^H \mathbf{K}_{nq}^{-1} \mathbf{a}_s = (\mathbf{U} \mathbf{a}_s)^H \mathbf{\Lambda} (\mathbf{U} \mathbf{a}_s)$. Because $\|\mathbf{U} \mathbf{a}_s\|^2 = \|\mathbf{a}_s\|^2$, and because $\mathbf{\Lambda}$ is positive, increasing $\|\mathbf{a}_s\|^2$ typically leads to an increased SINR, although we can construct counterexamples. This will be important when we discuss the SINR gain of the Rake-Max-SINR beamformer.

IV. ACOUSTIC RAKE RECEIVER

In this section, we present several formulations of the ARR. The Rake-DS beamformer is a straightforward generalization of the conventional DS beamformer. The one-forcing beamformer implements the idea of steering a fixed beam power towards every image source, while trying to minimize the interference and noise. The Rake-Max-SINR and Rake-Max-UDR beamformers optimize the corresponding performance measures; we show in Section VI that the Rake-Max-SINR beamforming performs best in terms of all evaluation criteria except, as expected, in terms of UDR.

A. Delay-and-Sum Raking

If we had access to every echo separately (*i.e.*, not summed with all the other echoes), we could align them all to maximize the performance gain. Unfortunately, this is not the case: each microphone picks up the convolution of speech with the impulse response, which is effectively a sum of overlapping echoes of the speech signal. If we only wanted to extract the direct path, we would use the standard DS beamformer (9). To build the Rake-DS receiver, we create a DS beamformer for every image source, and average the outputs,

$$\frac{1}{K+1} \sum_{k=0}^K \frac{\alpha'_k}{M} \sum_{m=0}^{M-1} y_m e^{j\kappa \| \mathbf{r}_m - \mathbf{s}_k \|}, \quad (13)$$

where $\alpha'_k \stackrel{\text{def}}{=} \alpha_k / (4\pi \| \mathbf{r}_m - \mathbf{s}_k \|)$. We read out the beamforming weights from (13) as

$$\mathbf{w}_{\text{R-DS}} = \frac{1}{\| \sum_k \mathbf{a}(\mathbf{s}_k) \|} \sum_{k=0}^K \mathbf{a}(\mathbf{s}_k) = \frac{\mathbf{A}_s \mathbf{1}}{\| \mathbf{A}_s \mathbf{1} \|}, \quad (14)$$

where we chose the scaling in analogy with (9) (scaling of the weights does not alter the output SINR). It can be seen that this is just a scaled sum of the steering vectors for each image source.

B. One-Forcing Raking

A different approach, based on intuition, is to design a beamformer that listens to all K image sources with the same power, and at the same time minimizes the noise and interference energy:

$$\begin{aligned} & \underset{\mathbf{w} \in \mathbb{C}^M}{\text{minimize}} \mathbb{E} \left| \sum_{k=0}^{K'} \mathbf{w}^H \mathbf{a}(\mathbf{q}_k) z + \mathbf{w}^H \mathbf{n} \right|^2 \\ & \text{subject to } \mathbf{w}^H \mathbf{a}(\mathbf{s}_k) = 1, \forall 0 \leq k \leq K. \end{aligned} \quad (15)$$

Alternatively, we may choose to null the interfering source and its image sources. Both cases are an instance of the standard linearly-constrained-minimum-variance (LCMV) beamformer [30]. Collecting all the steering vectors in a matrix, we can write the constraint as $\mathbf{w}^H \mathbf{A}_s = \mathbf{1}^T$. The solution can be found in closed form as

$$\mathbf{w}_{\text{OF}} = \mathbf{K}_{nq}^{-1} \mathbf{A}_s (\mathbf{A}_s^H \mathbf{K}_{nq}^{-1} \mathbf{A}_s)^{-1} \mathbf{1}_M. \quad (16)$$

A few remarks are in order. First, with M microphones, it does not make sense to increase K beyond M , as this results in more constraints than degrees of freedom. Second, using this beamformer is a bad idea if there is an interferer along the ray through the microphone array and any of image sources.

As with all LCMV beamformers, adding linear constraints uses up degrees of freedom that could otherwise be used for noise and interference suppression. It is better to let the “beamformer decide” or “the beamforming procedure decide” on how to maximize a well-chosen cost function; one such procedure is described in the next subsection.

C. Max-SINR Raking

The main workhorse of the paper is the Rake-Max-SINR. We compute the weights so as to maximize the SINR, taking into account the echoes of the desired signal, and the echoes of the interfering signal,

$$\underset{\mathbf{w} \in \mathbb{C}^M}{\text{maximize}} \frac{\mathbb{E} \left| \sum_{k=0}^K \mathbf{w}^H \mathbf{a}(\mathbf{s}_k) x \right|^2}{\mathbb{E} \left| \sum_{k=0}^{K'} \mathbf{w}^H \mathbf{a}(\mathbf{q}_k) z + \mathbf{w}^H \mathbf{n} \right|^2}. \quad (17)$$

The logic behind this expression can be summarized as follows: we present the beamforming procedure with a set of good sources, whose influence we aim to maximize at the output, and with a set of bad sources, whose power we try to minimize at the output. Interestingly, this leads to the standard Max-SINR beamformer with a structured steering vector and covariance matrix. Define the combined noise and interference covariance matrix as

$$\mathbf{K}_{nq} \stackrel{\text{def}}{=} \mathbf{K}_n + \sigma_z^2 \left(\sum_{k=0}^{K'} \mathbf{a}(\mathbf{q}_k) \right) \left(\sum_{k=0}^{K'} \mathbf{a}(\mathbf{q}_k) \right)^H, \quad (18)$$

where \mathbf{K}_n is the covariance matrix of the noise term, and σ_z^2 is the power of the interferer at a particular frequency.

Then the solution to (17) is given as

$$\mathbf{w}_{\text{R-SINR}} = \frac{\mathbf{K}_{nq}^{-1} \mathbf{A}_s \mathbf{1}}{\mathbf{1}^H \mathbf{A}_s^H \mathbf{K}_{nq}^{-1} \mathbf{A}_s \mathbf{1}}. \quad (19)$$

It is interesting to note that when $\mathbf{K}_{nq} = \sigma^2 \mathbf{I}_M$ (*e.g.*, no interferers and iid noise), the Rake-Max-SINR beamformer reduces to $\mathbf{A}_s \mathbf{1} / \| \mathbf{A}_s \mathbf{1} \|$, which is exactly the Rake-DS beamformer. This is analogous to the non-raking DS beamformer (9).

D. Max-UDR Raking

Finally, it is interesting to investigate what happens if we choose the weights that optimize the perceptually motivated UDR [14], [15]. The UDR expresses the fact that adding early reflections (up to 50 ms in the RIR) is as good as adding the energy to the direct sound in terms of speech intelligibility. The useful signal is a *coherent* sum of the direct and early reflected speech energy, so that

$$\text{UDR} = \frac{\mathbb{E} \sum_{k=0}^K | \mathbf{w}^H \mathbf{a}(\mathbf{s}_k) x |^2}{\mathbb{E} \left| \sum_{k=0}^{K'} \mathbf{w}^H \mathbf{a}(\mathbf{q}_k) z + \mathbf{w}^H \mathbf{n} \right|^2}. \quad (20)$$

In applications K is rarely large enough to cover all the reflections occurring within 50 ms, simply because it is too optimistic to assume we know all the corresponding image sources. Therefore, (20) typically underestimates the UDR.

Alas, because (20) is specified in the frequency domain, it is challenging to control whether the reflections in the numerator arrive before or after the direct sound. Nevertheless, it is interesting to analyze it as it provides a basis for future work on

time-domain raking formulations, and a meaningful evaluation of the raking algorithms presented in this paper.

To compute the Rake-Max-UDR weights, we solve the following program

$$\underset{\mathbf{w} \in \mathbb{C}^M}{\text{maximize}} \frac{\mathbb{E} \sum_{k=0}^K \left| \mathbf{w}^H \mathbf{a}(s_k) x \right|^2}{\left| \mathbb{E} \mathbf{w}^H \sum_{k=0}^{K'} \mathbf{a}(q_k) z + \mathbf{w}^H \mathbf{n} \right|^2}. \quad (21)$$

This amounts to maximizing a particular generalized Rayleigh quotient,

$$\frac{\mathbf{w}^H \mathbf{A}_s \mathbf{A}_s^H \mathbf{w}}{\mathbf{w}^H \mathbf{K}_{nq} \mathbf{w}}. \quad (22)$$

Assuming that \mathbf{K}_{nq} has a Cholesky decomposition as $\mathbf{K}_{nq} = \mathbf{C}^H \mathbf{C}$ we can rewrite the quotient (22) as

$$\frac{\tilde{\mathbf{w}}^H (\mathbf{C}^{-1})^H \mathbf{A}_s \mathbf{A}_s^H \mathbf{C}^{-1} \tilde{\mathbf{w}}}{\tilde{\mathbf{w}}^H \tilde{\mathbf{w}}}, \quad (23)$$

where $\tilde{\mathbf{w}} \stackrel{\text{def}}{=} \mathbf{C} \mathbf{w}$. The maximum of this expression is

$$\lambda_{\max}((\mathbf{C}^{-1})^H \mathbf{A}_s \mathbf{A}_s^H \mathbf{C}^{-1}), \quad (24)$$

where $\lambda_{\max}(\cdot)$ denotes the largest eigenvalue of the matrix in the argument. The maximum is achieved by the corresponding eigenvector $\tilde{\mathbf{w}}_{\max}$. Then the optimal weights are given as

$$\mathbf{w}_{\text{R-UDR}} = \mathbf{C}^{-1} \tilde{\mathbf{w}}_{\max}. \quad (25)$$

E. SINR Gain From Raking

Intuitively, if we have multiple sources of the desired signal scattered in space, and we account for it in the design, we should do at least as well as when we ignore the image sources. Let us see how large the gain can be for the Rake-Max-SINR beamformer. We have that

$$\text{SINR} = \sigma_x^2 (\mathbf{A}_s \mathbf{1})^H \mathbf{K}_{nq}^{-1} (\mathbf{A}_s \mathbf{1}). \quad (26)$$

Intuitively, the larger the norm of $\mathbf{A}_s \mathbf{1}$, the better the SINR (as \mathbf{K}_{nq} is positive). To explicitly see if there is any gain in using the acoustic rake receiver, we should compare the standard Max-SINR beamformer with the Rake-Max-SINR, *e.g.*, we should evaluate

$$\frac{(\sum_k \mathbf{a}(s_k))^H \mathbf{K}_{nq}^{-1} (\sum_k \mathbf{a}(s_k))}{\mathbf{a}(s_0)^H \mathbf{K}_{nq}^{-1} \mathbf{a}(s_0)}. \quad (27)$$

One possible interpretation of (27) is that we ask whether the steering vectors $\mathbf{a}(s_k)$ sum coherently or they cancel out.

To answer this, assume that \mathbf{s}_k , $0 \leq k \leq K$ are the desired sources (true and image), and let $\beta \stackrel{\text{def}}{=} \sum_{k=1}^K (\alpha_k / \alpha_0)^2$, where α_k is the strength of the source \mathbf{s}_k received by the array. Then

$$\mathbb{E} \left(\left\| \sum_{k=0}^K \mathbf{a}(s_k) \right\|^2 \right) \approx (1 + \beta) \mathbb{E}(\|\mathbf{a}(s_0)\|^2), \quad (28)$$

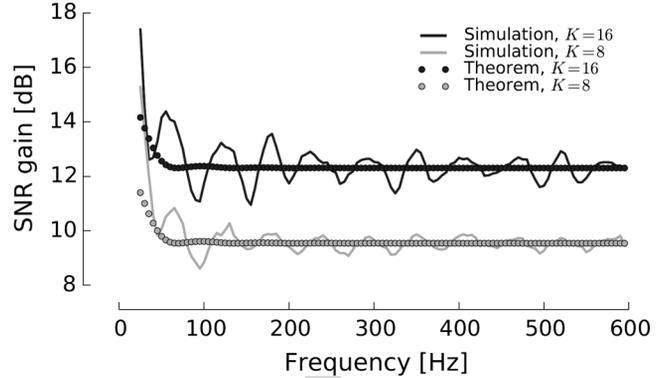


Fig. 4. Comparison of the simulated SNR gains and the theoretical prediction from Theorem 1 for $K = 8$, and $K = 16$. The theoretical prediction of the gain is $10 \log_{10}(8 + 1) \approx 9.54$ for $K = 8$, and $10 \log_{10}(16 + 1) \approx 12.30$ for $K = 16$.

that is, we can expect an increase in the output SINR approximately by a factor of $(1 + \beta)$ when using the Rake-Max-SINR beamformer. This statement is made precise in Theorem 1 in the Appendix. It holds when \mathbf{K}_{nq} has eigenvalues of similar magnitude, which is typically not the case in the presence of interferers. However, we show in Section VI that with interferers present, the gains actually increase.

A couple of remarks are in order:

- 1) This result is in expectation; it says that on average, the SINR will increase by a factor of $(1 + \beta)$. In the worst case, the steering vectors $\mathbf{a}(s_k)$ can even cancel out so that the SINR decreases. But the numerical experiments suggest that this is very rare in practice, and we can on the other hand observe large gains.
- 2) We see that summing the phasors in $a_m(s_k)$ behaves as a two-dimensional random walk. It is known that the root-mean-square distance of a 2D random walk from the origin after n steps is \sqrt{n} [31].
- 3) Due to the far-field assumption in Theorem 1, the attenuations α_k are assumed to be independent of the microphones; in reality they do depend on the source locations. However, they also depend on a number of additional factors, for example wall attenuations and radiation patterns of the sources. Therefore, for simplicity, we consider them to be independent. One can verify that this assumption does not change the described trend.

It is reassuring to observe the behavior suggested by (28) in practice. Fig. 4 shows the comparison of the prediction by Theorem 1 with the SNR gains observed in simulated rooms. In this case, we are comparing the pure SNR gain for white noise, without interferers. To generate Fig. 4, we randomized the location of the source inside the rectangular room. For simplicity we fixed the signal power as received by the microphones to the same value for all the image sources, so that the expected gain is $K + 1$ in the linear scale. The curves agree near-perfectly with the prediction of Theorem 1.

V. FINDING AND TRACKING THE ECHOES

Thus far we assumed that the locations of the image sources are known. In this section we briefly describe some methods to localize them when they are *a priori* unknown. We assume that

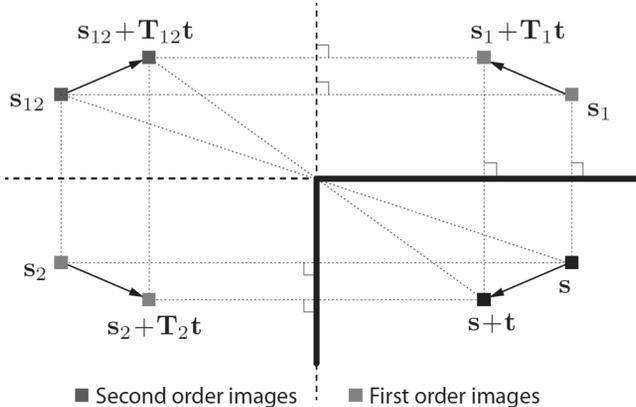


Fig. 5. Illustration of image source tracking in rectangular geometries.

we can localize the true source, or at least one image source. Combined with the knowledge of the room geometry, this suffices to find the locations of other image sources [32].

A. Known Room Geometry

In many cases, for example for fixed deployments, the room geometry is known. This knowledge could be obtained at the time of the deployment, or from blueprints. In most indoor environments, we encounter a large number of planar reflectors. These reflectors correspond to image sources. With reference to Fig. 5, we can easily compute the image source locations [7] (we note that the image source model is not limited to right angle geometries [8]).

Suppose that the real source is located at s . Then the image source with respect to wall i is computed as,

$$\mathbf{im}_i(s) = s + 2\langle \mathbf{p}_i - s, \mathbf{n}_i \rangle \mathbf{n}_i, \quad (29)$$

where i indexes the wall, \mathbf{n}_i is the outward normal associated with the i th wall, and \mathbf{p}_i is any point belonging to the i th wall. Analogously, we compute image sources corresponding to higher-order reflections,

$$\mathbf{im}_j(\mathbf{im}_i(s)) = \mathbf{im}_i(s) + 2\langle \mathbf{p}_j - \mathbf{im}_i(s), \mathbf{n}_j \rangle \mathbf{n}_j. \quad (30)$$

The above expressions are valid regardless of the dimensionality, concretely in 2D and 3D.

B. Acoustic Geometry Estimation

When the room geometry is not known, it is possible to estimate it using the same array that we use for beamforming. Recently a number of different methods appeared in the literature that propose to use sound to estimate the shape of a room. For example, in [10] the authors use a dictionary of wall impulse responses recorded with a particular array. In [11] the authors use tools from projective geometry together with the Hough transform to estimate the room geometry. In [13] the authors derive an *echo sorting* mechanism that finds the image sources, from which the room geometry is then derived.

C. Without Estimating the Room Geometry

To design an ARR, we do not really need to know how the room looks like; we only need to know where the major echoes are coming from. One possible approach is to locate the image

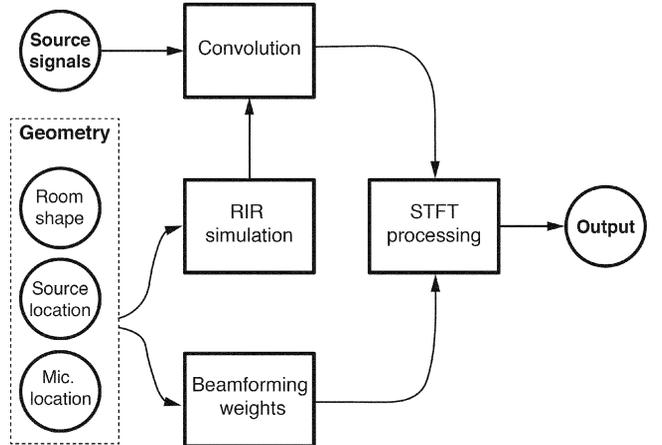


Fig. 6. Block diagram of the simulation setup used for numerical experiments.

sources in the initial calibration phase, and then track their movement by tracking the true source.

We propose a tracking rule that leverages the knowledge of the displacement of the true source. Again with reference to Fig. 5, we can state the following simple proposition.

Proposition 1: Suppose that the room has only right angles so that the walls are parallel with the coordinate axes. Let the source move from s to $s + t$. Then any image source s_k , moves to a point given by

$$s_k + \mathbf{T}t, \quad (31)$$

where $\mathbf{T} = \text{diag}(\pm 1, \mp 1)$ for odd generations, and $\mathbf{T} = \pm \mathbf{I}_2$ for even generations.

Proof: The proof follows directly from the figure. The displacement of the image source is the same as the displacement of the true source, passed through a series of reflections. Reflection matrices are diagonal matrices with ± 1 on the diagonal, and determinant equal to -1 , hence the result. \square

The usefulness of this proposition is that it gives us a tool to track the image sources even when we do not know the room geometry (as long as it has right angles). A possible use scenario is to start with a calibration procedure with a controlled source, and perform the echo sorting to find multiple image sources. Then if possible, we assign to each image source a generation (this is in fact a by-product of echo sorting), or we try different hypotheses using Proposition 1, and choose the one that maximizes the output SINR.

VI. NUMERICAL EXPERIMENTS

In this section, we validate the described theoretical results through numerical experiments. First, we analyze the beam-patterns produced by the ARR; second, we evaluate the SINR for various beamformers as a function of the number of image sources used in weight computation; and third, we evaluate the PESQ metric [22]. Finally, we show spectrograms that reveal visually the improved interferer and noise suppression achieved by the ARR.

A. Simulation Setup

We use a simple room acoustic framework written in Python, which relies on Numpy and Scipy for matrix computations [27].

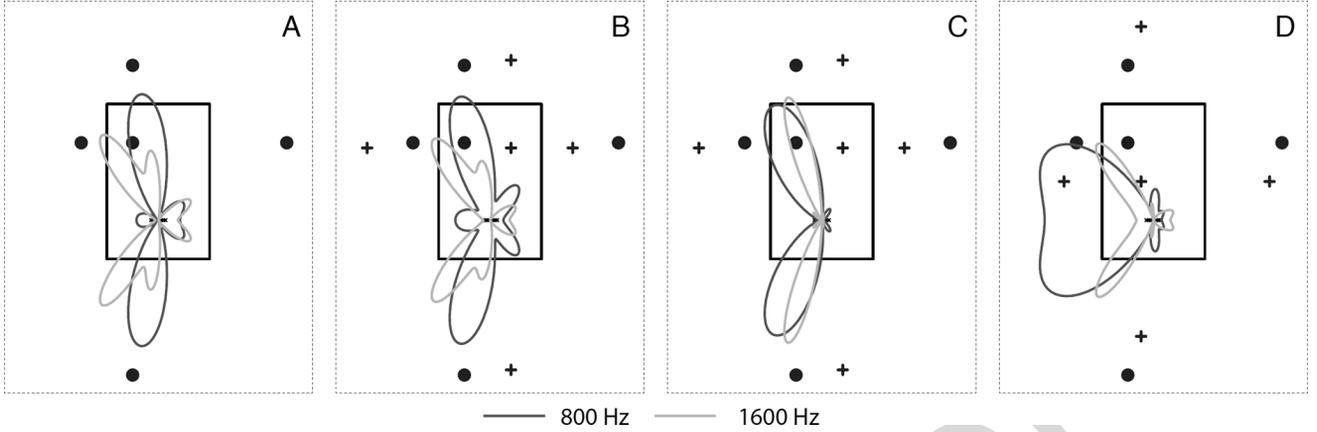


Fig. 7. Beam patterns in different scenarios. The rectangular room is 4 by 6 metres and contains a source of interest (●) and an interferer (+). The first-order image sources are also displayed. The weight computation of the beamformer includes the direct source and the first order image sources of both desired source and interferer (when applicable). (A) Rake-Max-SINR, no interferer, (B) Rake-Max-SINR, one interferer, (C) Rake-Max-UDR, one interferer, (D) Rake-Max-SINR, interferer is in direct path.

We limit ourselves in this paper to 2D geometry and rectangular rooms. In all experiments, the sampling frequency F_s was set to 8 kHz. An overview of the simulation setup is shown in Fig. 6.

Starting from the room geometry and the positions of the sources and microphones, we first compute the locations of all images sources up to ten generations. The reflectivity of the walls is fixed to 0.9. The RIR between the source \mathbf{s}_0 and the microphone \mathbf{r}_m is convolved with an ideal low-pass filter in the continuous domain and then sampled at the sampling frequency F_s ,

$$\tilde{a}_m(\mathbf{s}_0)[n] = \sum_{k=0}^K \frac{\alpha_k}{4\pi\|\mathbf{r}_m - \mathbf{s}_k\|} \text{sinc}\left(n - F_s \frac{\|\mathbf{r}_m - \mathbf{s}_k\|}{c}\right), \quad (32)$$

where K is the number of image sources considered. We choose the limits of n such that the cardinal sine in (32) decays sufficiently to avoid artifacts. The discrete signals from all sound sources are then convolved with their respective RIRs, and added together to obtain the m th microphone's signal.

We use the discrete-time STFT processing with a frame size of $L = 4096$ samples, 50% overlap and zero padding on both sides of the signal by $L/2$, and compute the beamforming in the frequency domain. A real fast Fourier transform of size $2L$ and a Hann window are used in the analysis. By exploiting the conjugate symmetry of the real FFT we only need to compute $L + 1$ beamforming weights, one for every positive frequency bin. The length L is dictated by the length of the beamforming filters in the time-domain and was set empirically so as to avoid any aliasing in the filter responses. We synthesize the output signal using the conventional overlap-add method [33].

B. Results

1) *Beampatterns*: We first inspect the beampatterns produced by the Rake-Max-SINR and Rake-Max-UDR beamformers for different source-interferer placements. We consider a 4 m × 6 m rectangular room with a source of interest at (1 m, 4.5 m) and a linear microphone array centered at (2 m, 1.5 m), parallel to the x -axis. Spacing between the microphones was set to 8 cm. In Fig. 7, we show the beampatterns for four different configurations of the source and the interferer. We consider a

scenario without an interferer, one with an interferer placed favorably at (2.8 m, 4.3 m), and finally one where the interferer is placed half-way between the desired source and the array, at (1.5 m, 3 m).

The last scenario is the least favorable. Interestingly, we can observe that the Rake-Max-SINR beampattern adjusts by completely ignoring the direct path, and steering the beam towards the echoes of the source of interest. This is validating the intuition that we can “hear behind an interferer by listening for the echoes”. Note that such a pattern cannot be achieved by a beamformer that only takes into account the direct path. We further note that, while the beampatterns only show the magnitude of the beamformer's response, the phase plays an important role with multiple sources present.

2) *SINR Gains From Raking*: In the experiments in this subsection, we set the power of the desired source and of the interferer to be equal, $\sigma_x^2 = \sigma_z^2 = 1$. The noise covariance matrix is set to $10^{-3} \cdot \mathbf{I}_M$. We use a circular array of $M = 12$ microphones with a diameter of 30 cm, and randomize the position of the desired source and the interferer inside the room. The resulting curves show median performance out of 20000 runs.

Fig. 8 shows output SINR for different beamformers. The one-forcing beamformer is left out because it performs poorly in terms of SINR, as predicted in the earlier discussion. Clearly, the Rake-Max-SINR beamformer outperforms all others. The output SINR for beamformers using only the direct path (Max-SINR and DS) remains approximately constant. The UDR is plotted against the number of image sources for various beamformers in Fig. 9. Even though the Rake-Max-UDR beamformer performs well in terms of the two measures, its output is perceptually displeasing due to audible pre-echoes; in informal listening tests, the Rake-Max-SINR beamformer did not produce such artifacts. It is interesting to note that the Rake-Max-SINR also performs well in terms of the UDR. Similar SINR gains to those in Fig. 8 are observed in Fig. 10 over a range of frequencies. It is therefore justified to extrapolate the results at one frequency in Fig. 8 to the wideband SINR.

3) *Evaluation of Speech Quality*: We complement the informal listening tests and the evaluation of SINR and UDR with extensive simulations to assess the improvement in speech

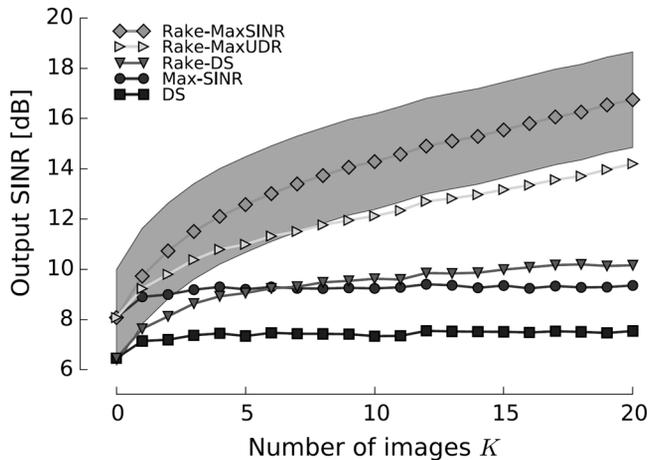


Fig. 8. Median output SINR plotted against the number of image sources used in the design for different beamformers, at a frequency $f = 1$ kHz. The shaded area contains the Rake-Max-SINR output SINRs for 50% of the 20000 Monte Carlo runs.

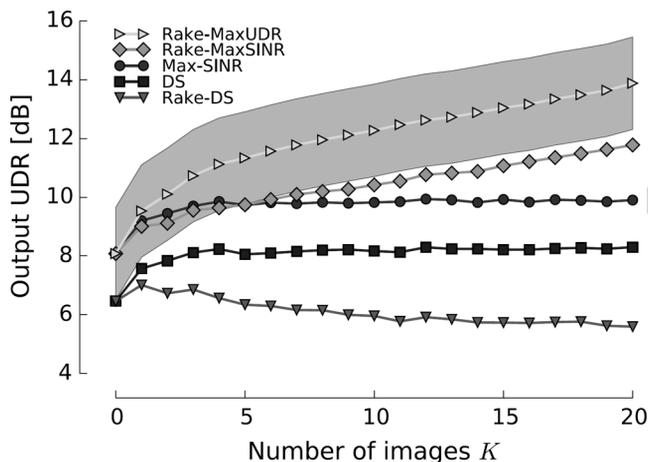


Fig. 9. Median output UDR plotted against the number of image sources used in the design for different beamformers, at a frequency $f = 1$ kHz. The shaded area contains the Rake-Max-UDR output UDRs for 50% of the 20000 Monte Carlo runs.

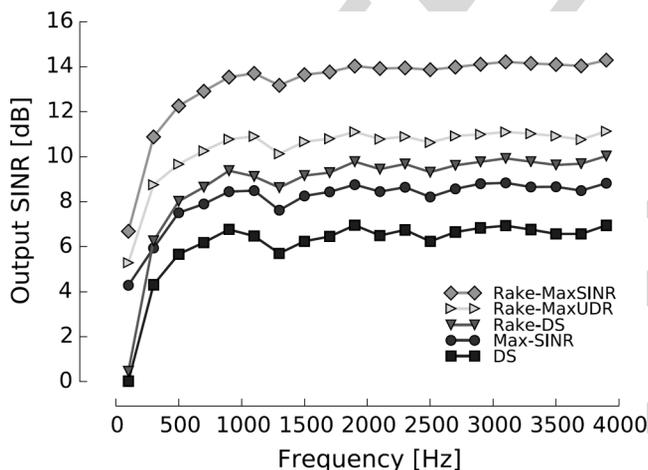


Fig. 10. Output SINR as a function of frequency for different beamformers and $K = K' = 10$. The curves show the average of 20000 runs, with averaging performed in the dB domain.

quality achieved by acoustic raking. We simulate a room with two sources—a desired source and an interferer—and compare the outputs of the Rake-DS, Rake-Max-SINR, and

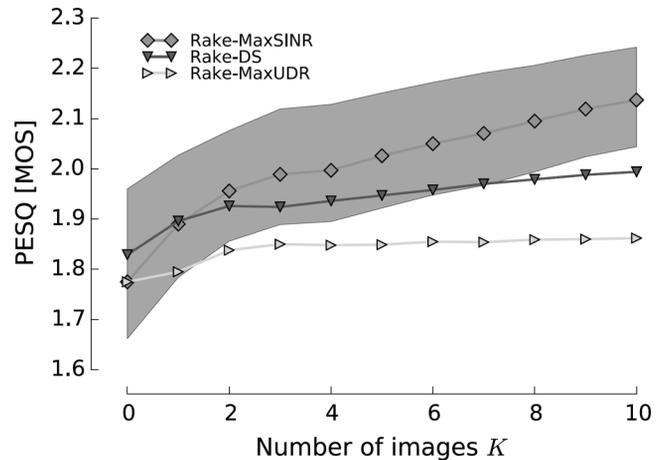


Fig. 11. Perceptual quality in MOS, evaluated using PESQ, as a function of the number of image sources used K . The lower limit of the ordinates is set to the median MOS of the degraded signal before processing, as measured at center of the array. The shaded area contains the Rake-Max-SINR output for 50% of the 10000 Monte Carlo runs.

Rake-Max-UDR as a function of the number of image sources used to design the beamformers.

The same number of image sources is used for the desired source and the interferer ($K = K'$). The performance metric used is the PESQ [22]. In particular, we use the reference implementation described by the ITU P.862 Amendment 2 [34]. PESQ compares the reference signal with the degraded signal and predicts the perceptual quality of the latter as it would be measured by the mean opinion score (MOS) value, on a scale from 1 to 4.5.

We consider the same room and microphone array setting as before (see Fig. 7(a)). The desired and the interfering sources are placed uniformly at random in a rectangular area with lower left corner at (1 m, 2.5 m) and upper right corner at (3 m, 5 m). To limit the experimental variation, the speech samples attributed to the sources are fixed throughout the simulation. The two sources start reproducing speech at the same time and approximately overlap for the total duration of the speech samples. The signals are normalized to have the same power at the source, and we add white Gaussian noise to the microphone signals, with power chosen so that the SNR of the direct sound for the desired source is 20 dB at the center of the microphone array. All signals are high-pass filtered with a cut-off frequency of 300 Hz. The reference for all PESQ results is the direct path of the desired source as measured at the center of the array (2 m, 1.5 m).

The median PESQ measure of 10000 Monte Carlo runs, given in raw MOS, is shown in Fig. 11. The median PESQ of the degraded signal measured at the center of the array before processing was found to be 1.6 raw MOS. When only the direct sound is used (*i.e.*, $K = 0$), all three beamformers yield the same improvement of about 0.2 raw MOS. We observe that Rake-DS is slightly outperforms the other beamformers. Using any number of echoes in addition to the direct sound results in larger MOS for all beamformers. When more than one image source is used, the Rake-Max-SINR beamformer always yields the largest MOS, with up to 0.5 MOS gain when using 10 images sources.

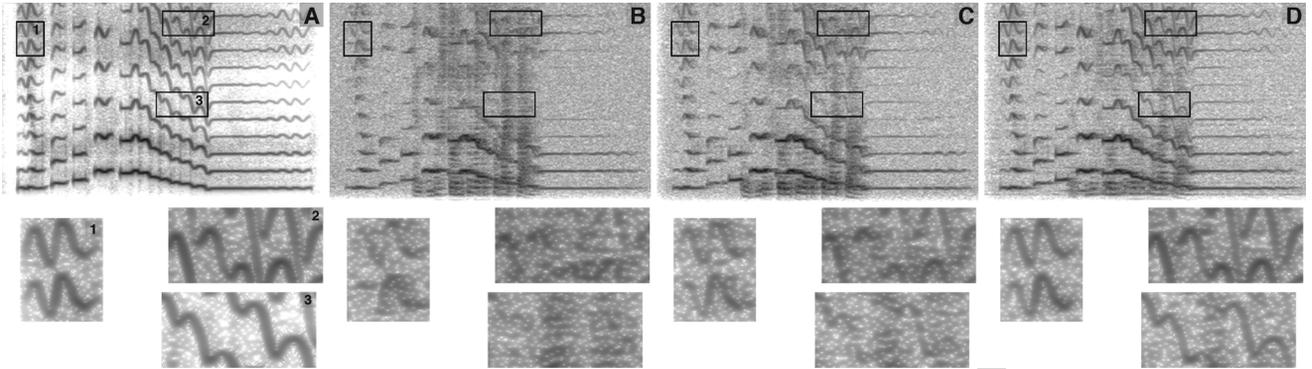


Fig. 12. Comparison of the conventional Max-SINR and Rake-Max-SINR beamformer on a real speech sample. Spectrograms of (A) clean signal of interest, (B) signal corrupted by an interferer and additive white Gaussian noise at the microphone input, outputs of (C) conventional Max-SINR and (D) Rake-Max-SINR beamformers. Time naturally goes from left to right, and frequency increases from zero at the bottom up to $F_s/2$. To highlight the improvement of Rake-Max-SINR over Max-SINR, we blow-up three parts of the spectrograms in the lower part of the figure. The boxes and the corresponding part of the original spectrogram are numbered in (A). The numbering is the same but omitted in the rest of the figure for clarity.

It is worth mentioning that in the beamformer design, we do not assume that we know the spectrum of the source or the interferer—we design as if it was flat. Thus the interferer acts as a strong source of colored, spatially correlated, non-stationary noise, spectrally mismatched with the designed beamformer. There is another source of model mismatch: while the RIRs were computed using hundreds of image sources, we use only up to ten to design the beamformers.

4) *Spectrograms and Sounds Samples*: Finally, we present the spectrograms for a scenario where we want to focus on a singer in the presence of interfering speech. We consider the same room, source, interferer, and microphone array geometry as in Fig. 7(b).

The source signal is a passage by a female opera singer (Fig. 12(a)), with strongly pronounced harmonics; the interfering signal is a male speech extract. The two signals are normalized to have unit maximum amplitude. We add white Gaussian noise to the microphone signals with power such that the SNR of the direct sound of the desired source is 20 dB at the center of the microphone array. All signals are high-pass filtered with a cut-off frequency of 300 Hz. The Rake-Max-SINR beamformer weights are computed using the direct source and three generations of image sources for both the desired sound source (singing) and the interferer (speech).

The output of the conventional Max-SINR beamformer (Fig. 12(c)) is compared to that of the Rake-Max-SINR (Fig. 12(d)). We can observe from the spectrogram that the Rake-Max-SINR reduces very effectively the power of the interfering signal at all frequencies, but particularly in the mid to high range. This is true even when the interferer overlaps significantly with the desired signal. Informal listening tests confirm that the Rake-Max-SINR maintains high quality of the desired signal while strongly reducing the interference. The Rake-Max-UDR beamformer provides good interference suppression, but it produces audible pre-echoes that render it unsuitable for speech processing applications. The sound clips can be found online along the code.

VII. CONCLUSION

We investigated the concept of acoustic rake receivers—beamformers that use echoes. Unlike earlier related

work, we presented optimal formulations that outperform the delay-and-sum style approaches by a large margin. This is especially true in the presence of interferers, hence the title “Raking the Cocktail Party”. We demonstrate theoretically that the ARR improves the SINR, and the numerical simulations agree well with these predictions.

Beyond theoretical and numerical evaluations of the performance measures, we demonstrated in informal listening tests the improved interference suppression by the ARR. Moreover, extensive simulation determined that the ARR improves the subjective quality, as predicted by PESQ, proportionally to the number of image sources used. A particularly illustrative example is when the interferer is occluding the desired source—the optimal ARR takes care of this simply by listening to the echoes.

Perhaps the most important aspect of ongoing work is the design of robust formulations of ARR. This may involve various heuristics, as well as combinatorial optimization due to the discrete nature of image sources. We expect that the raking beamformers described in this paper inherit the robustness properties of their classical counterparts. For example, the Rake-DS beamformer is likely to be more robust to array calibration errors than the Rake-Max-SINR beamformer. Furthermore, we expect that taking the image source perspective makes various ARR more robust to errors in source locations than the schemes that assume the knowledge of the RIR.

Another line of ongoing work investigates the time-domain formulations of the ARR, with some initial results already available [25]. Time-domain formulations offer better control over whether the echoes appear before or after the direct sound. This provides a more natural framework for optimizing perceptually motivated performance measures such as UDR.

APPENDIX THEOREM 1

We note that the theorem is stated for a linear array, but the described behavior is universal.

Theorem 1: Assume that there are $K + 1$ sources located at $\mathbf{s}_k = r_k[\cos \theta_k \ \sin \theta_k]^T$ where $\theta_k \sim \mathcal{U}(0, 2\pi)$ and $r_k \sim \mathcal{U}(a, b)$ are all independent, for some $0 < a < b$ such that the far-field assumption holds. Let \mathbf{A}_s collect the corresponding

TABLE II
SUMMARY OF BEAMFORMERS.

Acronym	Description	Beamforming Weights
DS	Align delayed copies of signal at the microphone	$\mathbf{w}_{\text{DS}} = \mathbf{a}(\mathbf{s})/\ \mathbf{a}(\mathbf{s})\ $
Max-SINR	$\max. \mathbf{w}^H \mathbf{a}_s \mathbf{a}_s^H \mathbf{w} / (\mathbf{w}^H \mathbf{K}_{nq} \mathbf{w})$	$\mathbf{w}_{\text{SINR}} = \mathbf{K}_{nq}^{-1} \mathbf{a}_s / (\mathbf{a}_s^H \mathbf{K}_{nq}^{-1} \mathbf{a}_s)$
Rake-DS	Weighted average of DS beamformers over image sources	$\mathbf{w}_{\text{R-DS}} = \mathbf{A}_s \mathbf{1} / \ \mathbf{A}_s \mathbf{1}\ $
Rake-OF	$\min. \mathbb{E} \left \sum_{k=0}^{K'} \mathbf{w}^H \mathbf{a}(\mathbf{q}_k) z + \mathbf{w}^H \mathbf{n} \right ^2, \text{ s.t. } \mathbf{w}^H \mathbf{A}_s = \mathbf{1}^T$	$\mathbf{w}_{\text{OF}} = \mathbf{K}_{nq}^{-1} \mathbf{A}_s (\mathbf{A}_s^H \mathbf{K}_{nq}^{-1} \mathbf{A}_s)^{-1} \mathbf{1}_M$
Rake-Max-SINR	$\max. \mathbb{E} \left \sum_{k=0}^K \mathbf{w}^H \mathbf{a}(\mathbf{s}_k) x \right ^2 / \mathbb{E} \left \sum_{k=0}^{K'} \mathbf{w}^H \mathbf{a}(\mathbf{q}_k) z + \mathbf{w}^H \mathbf{n} \right ^2$	$\mathbf{w}_{\text{R-SINR}} = \mathbf{K}_{nq}^{-1} \mathbf{A}_s \mathbf{1} / (\mathbf{1}^H \mathbf{A}_s^H \mathbf{K}_{nq}^{-1} \mathbf{A}_s \mathbf{1})$
Rake-Max-UDR	$\max. \mathbb{E} \sum_{k=0}^K \mathbf{w}^H \mathbf{a}(\mathbf{s}_k) x ^2 / \mathbb{E} \left \mathbf{w}^H \sum_{k=0}^{K'} \mathbf{a}(\mathbf{q}_k) z + \mathbf{w}^H \mathbf{n} \right ^2$	$\mathbf{w}_{\text{R-UDR}} = \mathbf{C}^{-1} \tilde{\mathbf{w}}_{\max} ((\mathbf{C}^{-1})^H \mathbf{A}_s \mathbf{A}_s^H \mathbf{C}^{-1})$

TABLE I
[AUTHOR: Please cite Tables 1 and 2 in the paper] SUMMARY OF NOTATION.

Symbol	Meaning
M	Number of microphones
\mathbf{r}_m	Location of the m th microphone
\mathbf{s}_0	Location of the desired source
\mathbf{s}_i	Location of the i th image of the desired source ($i \geq 1$)
\mathbf{q}_0	Location of the interfering source
\mathbf{q}_i	Location of the i th image of the interfering source ($i \geq 1$)
$x(e^{j\omega})$	Spectrum of the sound from the desired source
$z(e^{j\omega})$	Spectrum of the sound from the interfering source
$\mathbf{w}(e^{j\omega})$	Vector of beamformer weights
K	Number of considered desired image sources
K'	Number of considered interfering image sources
$a_m(\mathbf{s})$	m th component of the steering vector for a source at \mathbf{s}
y_m	Signal picked up by the m th microphone
$\ \cdot\ $	Euclidean norm, $\ \mathbf{x}\ = (\sum x_i ^2)^{1/2}$.

steering vectors for a uniform linear microphone array. Then $\mathbb{E} \|\mathbf{A}_s \mathbf{1}\|^2 \geq (1 + \beta) \mathbb{E} \|\mathbf{a}(\mathbf{s}_0)\|^2$, where $\beta = \sum_{k=1}^K (\alpha_k / \alpha_0)^2$, and α_k are attenuations of the steering vectors, assumed independent from the source locations. In fact, $\mathbb{E} (\|\mathbf{A}_s \mathbf{1}\|^2) = (1 + \beta) \mathbb{E} (\|\mathbf{a}(\mathbf{s}_0)\|^2) + O(1/\Omega^3)$.

Proof: Thanks to the far-field assumption, we can decompose the steering vector into a factor due to the array, and a phase factor due to different distances of different image sources. We have that

$$a_m = (\mathbf{A}_s \mathbf{1})_m = \sum_{k=0}^K \alpha_k e^{-j\kappa m d \sin \theta_k} e^{-j\Omega \delta_k / c}, \quad (33)$$

where d is the microphone spacing and $\kappa \stackrel{\text{def}}{=} \Omega/c$. Without loss of generality we assume that $\delta_k \sim \mathcal{U}(a, b)$. We can further write

$$\begin{aligned} \mathbb{E} |a_m|^2 &= \mathbb{E} \left[\left(\sum_{k=0}^K \alpha_k e^{-j\kappa m d \sin \theta_k} e^{-j\kappa \delta_k} \right) \right. \\ &\quad \times \left. \left(\sum_{\ell=0}^K \alpha_\ell e^{j\kappa m d \sin \theta_\ell} e^{j\kappa \delta_\ell} \right) \right] \\ &= \sum_{k=0}^K \alpha_k^2 + \sum_{k \neq \ell=0}^K \alpha_k \alpha_\ell \mathbb{E} \left[e^{j\kappa m d (\sin \theta_\ell - \sin \theta_k)} e^{j\kappa (\delta_\ell - \delta_k)} \right]. \end{aligned} \quad (34)$$

Invoking the independence for $k \neq \ell$, we compute the above expectation as

$$\begin{aligned} \mathbb{E} \left[e^{j\kappa m d (\sin \theta_\ell - \sin \theta_k)} e^{j\kappa (\delta_\ell - \delta_k)} \right] \\ = \frac{2J_0^2(m d \kappa) [1 - \cos(\Delta \kappa)]}{(\Delta \kappa)^2}, \quad (35) \end{aligned}$$

where J_0 denotes the Bessel function of the first kind and zeroth order and $\Delta \stackrel{\text{def}}{=} b - a$.

Plugging this back into (34), we obtain

$$\mathbb{E} |a_m|^2 = \sum_{k=0}^K \alpha_k^2 \left(1 + C \frac{2J_0^2(m d \kappa) [1 - \cos(\Delta \kappa)]}{(\Delta \kappa)^2} \right), \quad (36)$$

where $C = \sum_{k \neq \ell} \alpha_k \alpha_\ell / \sum_k \alpha_k^2$.

Because $|J_0(z)| \leq \sqrt{2/(\pi z)} + O(|z|^{-1})$ ([35], Eq. 9.2.1), we see that the expression in brackets is $1 + O(\Omega^{-3})$. Rewriting

$$\sum_{k=0}^K \alpha_k^2 = \frac{1}{M} \mathbb{E} \|\mathbf{a}(\mathbf{s}_0)\|^2 \left(1 + \sum_{k=1}^K \left(\frac{\alpha_k}{\alpha_0} \right)^2 \right) \quad (37)$$

concludes the proof. \square

ACKNOWLEDGMENT

We thank the anonymous reviewers for a number of constructive suggestions that have improved the manuscript.

REFERENCES

- [1] R. Price and P. E. Green, "A communication technique for multipath channels," *Proc. IRE. IEEE*, vol. 46, no. 3, pp. 555–570, Mar. 1958.
- [2] A. F. Naguib, "Space-time receivers for CDMA multipath signals," in *Proc. IEEE ICC. IEEE*, Montreal, QC, Canada, 1997, pp. 304–308.
- [3] B. H. Khalaj, A. Paulraj, and T. Kailath, "2D RAKE receivers for CDMA cellular systems," in *Proc. IEEE GLOBECOM*, 1994, pp. 400–404.
- [4] P. Annibale, F. Antonacci, P. Bestagini, A. Brutti, A. Canclini, L. Cristoforetti, J. Filos, E. Habets, W. Kellerman, K. Kowalczyk, A. Lombard, E. Mabande, D. Markovic, P. Naylor, and M. Omologo, "The SCENIC project: Space-time audio processing for environment-aware acoustic sensing and rendering," in *Proc. 131st Conv. Audio Eng. Soc.*, New York, NY, USA, 2011, Audio Engineering Society.
- [5] A. E. O'Donovan, R. Duraiswami, and D. N. Zotkin, "Automatic matched filter recovery via the audio camera," in *Proc. IEEE ICASSP. IEEE*, Dallas, TX, USA, 2010, pp. 2826–2829.
- [6] E.-E. Jan, P. Svaizer, and J. L. Flanagan, "Matched-filter processing of microphone array for spatial volume selectivity," in *Proc. IEEE ISCAS*, 1995, vol. 2, pp. 1460–1463.

- [7] J. B. Allen and D. A. Berkley, "Image method for efficiently simulating small-room acoustics," *J. Acoust. Soc. Amer.*, vol. 65, no. 4, pp. 943–950, 1979.
- [8] J. Borish, "Extension of the image model to arbitrary polyhedra," *J. Acoust. Soc. Amer.*, vol. 75, no. 6, pp. 1827–1836, 1984.
- [9] D. B. Ward, E. A. Lehmann, and R. C. S. Williamson, "Particle filtering algorithms for tracking an acoustic source in a reverberant environment," *IEEE Trans. Audio, Speech, Lang. Process.*, vol. 11, no. 6, pp. 826–836, Nov. 2003.
- [10] F. Ribeiro, D. A. Florencio, D. E. Ba, and C. Zhang, "Geometrically constrained room modeling with compact microphone arrays," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 20, no. 5, pp. 1449–1460, Jul. 2012.
- [11] F. Antonacci, J. Filos, M. R. P. Thomas, E. A. P. Habets, A. Sarti, P. A. Naylor, and S. Tubaro, "Inference of room geometry from acoustic impulse responses," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 20, no. 10, pp. 2683–2695, Oct. 2012.
- [12] I. Dokmanic, Y. M. Lu, and M. Vetterli, "Can one hear the shape of a room: The 2-D polygonal case," in *Proc. IEEE ICASSP*, Prague, Czech Republic, 2011, pp. 321–324.
- [13] **[AUTHOR: Please provide page range]** I. Dokmanic, R. Parhizkar, A. Walther, Y. M. Lu, and M. Vetterli, "Acoustic echoes reveal room shape," *Proc. Nat. Acad. Sci.*, vol. 110, no. 30, Jun. 2013.
- [14] J. S. Bradley, H. Sato, and M. Picard, "On the importance of early reflections for speech in rooms," *J. Acoust. Soc. Amer.*, vol. 113, no. 6, p. 3233, 2003.
- [15] J. Lochner and J. F. Burger, "The influence of reflections on auditorium acoustics," *J. Sound Vibr.*, vol. 1, no. 4, pp. 426–454, 1964.
- [16] E. Habets, J. Benesty, I. Cohen, S. Gannot, and J. Dmochowski, "New insights into the MVDR beamformer in room acoustics," *IEEE Trans. Audio, Speech, Lang. Process.*, vol. 18, no. 1, pp. 158–170, Jan. 2010.
- [17] M. R. P. Thomas, N. D. Gaubitch, and P. A. Naylor, "Application of channel shortening to acoustic channel equalization in the presence of noise and estimation error," in *Proc. IEEE WASPAA*, New Paltz, NY, USA, 2011, pp. 113–116.
- [18] W. Zhang, E. Habets, and P. A. Naylor, "On the use of channel shortening in multichannel acoustic system equalization," in *Proc. IWAENC*, Tel Aviv, Israel, 2010.
- [19] M. Miyoshi and Y. Kaneda, "Inverse filtering of room acoustics," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 36, no. 2, pp. 145–152, Feb. 1988.
- [20] K. Furuya, "Noise reduction and dereverberation using correlation matrix based on the multiple-input/output inverse-filtering theorem (MINT)," in *Proc. Int. Workshop on HSC*, Kyoto, Japan, 2001, pp. 59–62.
- [21] J. Benesty, J. Chen, Y. A. Huang, and J. Dmochowski, "On microphone-array beamforming from a MIMO acoustic signal processing perspective," *IEEE Trans., Audio, Speech, Lang. Process.*, vol. 15, no. 3, pp. 1053–1065, Mar. 2007.
- [22] A. W. Rix, J. G. Beerends, M. P. Hollier, and A. P. Hekstra, "Perceptual evaluation of speech quality (PESQ)—A new method for speech quality assessment of telephone networks and codecs," in *Proc. IEEE ICASSP*, Salt Lake City, UT, USA, 2001, pp. 749–752.
- [23] S. Haykin and Z. Chen, "The cocktail party problem," *Neural Comput.*, vol. 17, no. 9, pp. 1875–1902, 2005.
- [24] W. Herbordt and W. Kellermann, "Adaptive beamforming for audio signal acquisition," in *Adaptive Signal Processing*. Berlin/Heidelberg, Germany: Springer, 2003, pp. 155–194.
- [25] **[AUTHOR: Please update status]** R. Scheibler, I. Dokmanic R, and M. Vetterli, "Raking echoes in the time domain," in *Proc. IEEE ICASSP*, 2015, accepted for publication.
- [26] B. D. Carlson, "Covariance matrix estimation errors and diagonal loading in adaptive arrays," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 24, no. 4, pp. 397–401, Jul. 1988.
- [27] T. E. Oliphant, "Python for scientific computing," *Comput. Sci. Eng.*, vol. 9, no. 3, pp. 10–20, 2007.
- [28] D. Duffy, *Green's Functions With Applications*. Boca Raton, FL, USA: Chapman & Hall/CRC, 2001.
- [29] B. D. Van Veen and K. M. Buckley, "Beamforming: A versatile approach to spatial filtering," *IEEE ASSP Mag.*, vol. 5, no. 2, pp. 4–24, Apr. 1988.
- [30] O. L. I. Frost, "An algorithm for linearly constrained adaptive array processing," in *Proc. IEEE. IEEE*, 1972, pp. 926–935.
- [31] W. McCrea and F. Whipple, "Random paths in two and three dimensions," in *Proc. Roy. Soc., Edinburgh, U.K.*, 1940, vol. 60, pp. 281–298.
- [32] O. Öçal, I. Dokmanic, and M. Vetterli, "Source localization and tracking in non-convex rooms," in *Proc. IEEE ICASSP*, 2014, pp. 1429–1433.
- [33] J. J. Shynk, "Frequency-domain and multirate adaptive filtering," *IEEE Signal Process. Mag.*, vol. 9, no. 1, pp. 14–37, Jan. 1992.
- [34] "ITU-T P.862 Amendment 2, Reference Implementations and Conformance Testing for ITU-T Recs P.862, P.862.1 and P.862.2," 11, 2005, p. 862.
- [35] M. Abramowitz and I. A. Stegun, *Handbook of Mathematical Functions*. New York, NY, USA: Courier Dover, 1972.



ICASSP 2011. In 2014 he received a Google Ph.D. Fellowship.



acoustic beamforming.



and the IEEE Signal Processing Technical and Society Awards (2002, 2010). He is a Fellow of IEEE, ACM, and EURASIP. He is the coauthor of the books *Wavelets and Subband Coding* (1995), *Signal Processing for Communications* (2008), and *Foundations of Signal Processing* (2014). He is a highly cited researcher in engineering (Thomson ISI Web of Science, 2007, 2014).

Ivan Dokmanic is a Ph.D. candidate in Audiovisual Communications Laboratory (LCAV) at EPFL (expected graduation in May 2015). His interests include inverse problems, audio and acoustics, signal processing for sensor arrays/networks, and fundamental signal processing. Previously he was a Teaching Assistant at the University of Zagreb, a Codec Developer for MainConcept AG, and a Digital Audio Effects Designer for Little Endian Ltd. For his work on room shape reconstruction using sound he received the best student paper award at

Robin Scheibler (S'07) received the B.Sc. and M.Sc. in communications systems from Ecole Polytechnique Fédérale de Lausanne (EPFL), Switzerland, in 2006 and 2009, respectively. From 2009 to 2010, he was with IBM Research–Zürich. After this, he was with the NEC Media Information Processing group in Kawasaki, Japan, from 2011 to 2012. He then returned to EPFL where he is working towards the Ph.D. degree at the Laboratory for Audiovisual Communications. His research interests lie in algorithmic signal processing, speech processing, and

Martin Vetterli received his engineering degree from ETH Zurich, his M.S. from Stanford, and his Doctorate from EPF Lausanne. He was on the faculty of Columbia University and UC Berkeley before joining EPFL. He is currently President of the National Research Council of the Swiss NSF. He works on signal processing and communications, e.g., wavelet theory and applications, sparse sampling, joint source–channel coding and sensor networks. He received the best paper awards of IEEE Signal Processing Society (1991, 1996, 2006), and the IEEE Signal Processing Technical and Society Awards (2002, 2010). He is a Fellow of IEEE, ACM, and EURASIP. He is the coauthor of the books *Wavelets and Subband Coding* (1995), *Signal Processing for Communications* (2008), and *Foundations of Signal Processing* (2014). He is a highly cited researcher in engineering (Thomson ISI Web of Science, 2007, 2014).