Sampling and Exact Reconstruction of Pulses with Variable Width

Gilles Baechler, *Student Member, IEEE*, Adam Scholefield, *Member, IEEE*, Loïc Baboulaz,
and Martin Vetterli, *Fellow, IEEE*

Abstract—Recent sampling results enable the reconstruction of signals composed of streams of fixed-shaped pulses. These results have found applications in topics as varied as channel estimation, biomedical imaging and radio astronomy. However, in many real signals, the pulse shapes vary throughout the signal. In this paper, we show how to sample and perfectly reconstruct Lorentzian pulses with variable width. In the noiseless case, perfect recovery is guaranteed by a set of theorems. In addition, we verify that our algorithm is robust to model mismatch and noise. This allows us to apply the technique to two practical applications: electrocardiogram (ECG) compression and bidirectional reflectance distribution function (BRDF) sampling. ECG signals are one dimensional, but the BRDF is a higher dimensional signal, which is more naturally expressed in a spherical coordinate system; this motivated us to extend the theory to the 2D and spherical cases. Experiments on real data demonstrate the viability of the proposed model for ECG acquisition and compression, as well as the efficient representation and low-rate sampling of specular BRDFs.

Index Terms—Finite rate of innovation, variable width, signal sampling, signal denoising, pulse compression methods, spectral analysis, electrocardiogram, reflectance function, specularity.

I. INTRODUCTION

SAMPLING theorems provide a bridge between the analog signals of the real world and their discrete representations. In the ideal case, they allow perfect reconstruction of a continuous signal from only a limited number of discrete samples. The most notable embodiment of this concept is the well-known Nyquist-Shannon theorem [1]. This states that any signal that belongs to the shift-invariant subspace of band-limited functions can be perfectly reconstructed from a set of uniformly spaced discrete samples, provided that the sampling frequency is at least twice the bandwidth of the signal. While very powerful, this result falls short for many classes of signal.

Over the years, sampling results have emerged for signals belonging to more general classes, including signals that lie in general shift-invariant subspaces [2] and *some* other signals
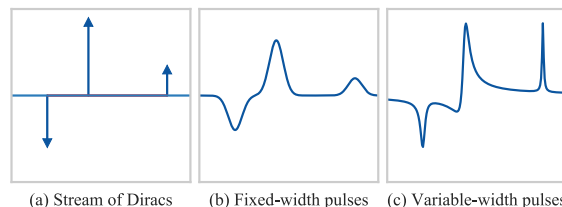
Fig. 1. Signals with finite rate of innovation: (a) stream of Diracs recoverable with the original FRI theory; (b) fixed-shaped pulses recoverable using Strang-Fix theory; (c) variable-width pulses, recoverable by the proposed technique.

with a finite rate of innovation (FRI) [3]—i.e., signals that have a finite number of degrees of freedom per unit time. This latter class extends the former to many other types of signals such as streams of Dirac deltas as depicted in Fig. 1(a), piecewise polynomials and piecewise sinusoidal signals [4].

All these schemes assume that we only have access to samples of a filtered version of the signal; that is, the original continuous-time signal $x(t)$ is filtered with a kernel $\varphi(t)$ before being uniformly sampled at time instants $t = nT$. Typically, $\varphi(t)$ is a lowpass filter that is either chosen by design, or, as is more common, is a characteristic of the acquisition system.

Like traditional Nyquist-Shannon sampling, the original FRI theory was developed for infinite support kernels, such as the sinc ideal low-pass filter; however, the results have been recently extended to physically realizable kernels of compact support that satisfy Strang-Fix conditions [5]. Another benefit of these more general kernels is that, when applied to streams of Diracs, they can be used to add more variety to the signals admitted by FRI. For example, it has enabled FRI theory to be utilised to retrieve streams of short pulses from ultrasound imaging [6] and action potentials of neurological data from calcium imaging [7]. Note that, although this allows more general pulses, as illustrated in Fig. 1(b), each pulse is still constrained to have the same shape (up to a multiplicative factor).

The question we address in this article is the following: Can we find a signal model that allows a different shape for each pulse? This is motivated by the fact that in nature, several signals can be parametrized as pulses with various shapes. In what follows, we show that it is indeed possible by considering signals that consist of sums of Lorentzian[1] pulses with variable width

---

[1]A Lorentzian function—also known as Cauchy distribution—is a bell-shaped curve that looks similar to a Gaussian but with heavier tails.

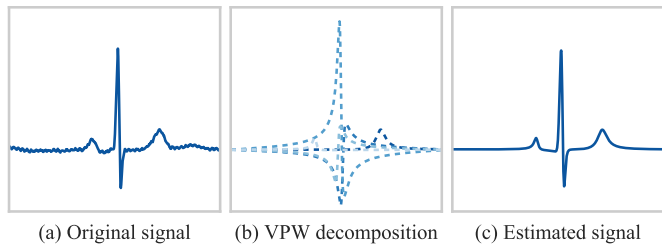(a) Original signal   (b) VPW decomposition   (c) Estimated signal

Fig. 2.   Example of decomposition of a signal into VPW pulses: (a) one period of a pulse from the MIT-BIH arrhythmia database [19]; (b) its decomposition into five asymmetric VPW pulses; (c) the VPW approximation (sum of the five pulses).

(a) Original   (b) VPW model

Fig. 3.   (a) Example of two dimensions of a measured BRDF and (b) its representation using a single VPW pulse.

and different degrees of asymmetry. While there are numerous parametric functions — Gaussians, B-splines, or wavelet basis functions to name a few — that are suitable to model signals with variable pulse shapes, the choice of Lorentzian functions is not innocuous; indeed, as we see in the next sections, the auto-regressive nature of their spectral coefficients enables us to estimate the parameters of the model using spectral estimation techniques such as *Prony's method*. Other algorithms estimating the parameters of sums of Lorentzians are described and analyzed in [8]–[10]; in this paper, we connect these results to sampling theory.

In addition to FRI techniques, recent extensions of compressive sensing also enable continuous parameter estimation but attack the problem from a different angle: while FRI theory relies on non-linear methods, these approaches enforce sparsity with the atomic norm and convex relaxations [11]–[14].

In this paper, we demonstrate, via a set of theorems, that sums of Lorentzians can be sampled and perfectly reconstructed, by adapting standard FRI machinery; whereas the core of the original FRI theory is composed of Diracs, the atoms of our signal model are called *VPW pulses*.[2] As highlighted in Fig. 1(c), the VPW model brings more versatility and flexibility to the set of FRI signals that can be acquired. We apply the VPW-FRI framework to two natural signals: electrocardiograms (ECGs) and bidirectional reflectance distribution functions (BRDFs).

ECG signals are composed of five different types of pulses, each of which corresponds to a particular event in a cardiac cycle [20]. The width of these pulses can vary greatly, making fixed-shaped pulse models less suitable; on the other hand, as illustrated in Fig. 2, a sum of VPW pulses closely approximates ECG pulse shapes. Other basis functions have been proposed to compress ECG signals, such as Gaussians [21] and Hermite polynomials [22]. In [21], nonlinear least square methods are used to estimate the parameters of the Gaussians, which does not always converge to a satisfactory solution. In [22], only the middle part of each heartbeat (the QRS complex) is compressed and, furthermore, it must be sampled at particular non-uniform locations. Hermite polynomials have also been used in other contexts, such as the detection of pathologies in ECG records [23]. During the completion of this manuscript, we also
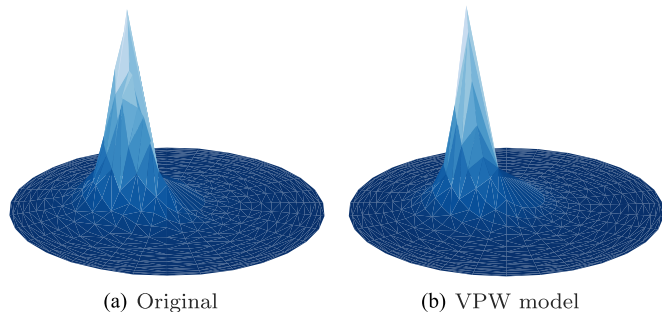
became aware of the work of Ouamri and Naït-Ali [24], in which a sum of Lorentzian functions is applied to ECG compression using an SVD-type algorithm. In addition to compression, we demonstrate accurate reconstruction of natural signals using a sampling rate that is of the same order as the rate of innovation of these signals, which is usually below the Nyquist rate.

The other catalyst for our research is the acquisition and representation of BRDFs, which are high-dimensional functions that characterize how light is reflected by surfaces. Usually, BRDFs contain high frequency components — called *specularities* — that are difficult to catch with classical sampling theory. Moreover, their shape varies through space, which makes existing FRI models inappropriate. The acquisition of BRDFs in the spherical domain requires dense sampling of the hemisphere, as can be seen in [25]. More elaborate techniques propose to perform the acquisition in a *transform* domain: for example in [26], the acquisition is done in the spherical harmonics domain, which substantially reduces the number of samples. The disadvantage is that it requires a specialized gantry, with a high number of light sources that can recreate the spherical harmonics functions. In this article, we leverage FRI theory to reduce the number of samples directly in the spherical domain. An example of a specularity approximated with a single spherical VPW pulse is shown in Fig. 3; here, the direction of outgoing light is kept fixed. By utilizing a simple parametric model, we are enforcing a strong prior that is particularly beneficial in certain scenarios. For example, in the BRDF problem it is typical to only have a few samples of sharp specularities; it is thus very challenging to obtain an accurate representation of the underlying reflectance function. Despite this, the proposed approach is able to recover a good approximation of these specularities, and the continuous representation enables us to later render the object from any illumination direction. This is in contrast to super-resolution approaches, where a denser sampling of the BRDF would only allow us to render the object under certain illumination directions or would force us to apply some other type of interpolation. In fact, graphics renderers deal with parametric forms of the reflectance function precisely for these reasons.

II. BACKGROUND

A. Signals with Finite Rate of Innovation

The *rate of innovation* ρ of a signal is defined as its number of degrees of freedom per unit time [3]. Ideally, we seek an

---

[2]Variable pulse width (VPW)-FRI was coined in [15], [16] in the context of electrocardiogram (ECG) representation. Other notable applications of VPW on multiple lead ECG signals are reported in [17], [18].
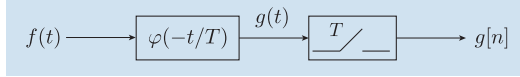
Fig. 4. Sampling operation. The continuous-time signal $f(t)$ is filtered with the kernel $\varphi(t)$ and then recorded at time instants $t = nT$ ($n \in \mathbb{Z}$) to obtain the discrete-time signal $g[n]$.

algorithm to reconstruct an FRI signal using *only* ρ parameters per unit time, in which case sampling at the rate of innovation allows perfect reconstruction. This is precisely what is achieved by the Nyquist-Shannon theorem, where the degrees of freedom are embodied by the expansion coefficients in the orthonormal basis formed by shifted sinc functions.

Another FRI signal, which is the archetype of the field, is a τ-periodic stream of K weighted Dirac pulses:

$$f(t) = \sum_{n \in \mathbb{Z}} \sum_{k=0}^{K-1} c_k \delta(t - t_k - n\tau), \qquad m \in \mathbb{Z}. \quad (1)$$

Since a Dirac can be completely described by its position and amplitude, one period of the signal $f(t)$ has precisely $2K$ free parameters (K positions and K amplitudes) and a rate of innovation of $\rho = 2K/\tau$.

In what follows, we will show how it is possible to reconstruct $f(t)$ from $N \geq 2K + 1$ samples obtained after an ideal low pass filter. To do this, we will first show how we can calculate some of the Fourier series coefficients of f from its samples and then show how these Fourier series coefficients can be used to calculate the $2K$ parameters of the signal.

B. From Samples to Fourier Series Coefficients

Although we do not have direct access to the Fourier series coefficients of $f(t)$, we can obtain a subset of them by taking the DFT of the samples obtained from the sampling scheme depicted in Fig. 4, with $\varphi(t)$ equal to an ideal low-pass filter. This result was first shown as part of [3, Theorem 1], and we reproduce it here in the following Lemma.

Lemma 1 (Vetterli, Blu and Marziliano 2002): Assume we take N samples of a τ-periodic function f, according to Fig. 4; that is, we have

$$g[n] = \langle f(t), \varphi(t - nT) \rangle, \quad n \in 0, 1, \ldots, N-1.$$

Furthermore, assume

$$\varphi(t) = B\mathrm{sinc}(Bt) = B\sin(\pi Bt)/(\pi Bt),$$

with $T = \tau/N$ and B satisfying $N = 2\lfloor B\tau/2 \rfloor + 1$. Define the DFT of $g[n]$ as

$$G[m] = \sum_{n=0}^{N-1} g[n]e^{-2\pi inm/N}$$

and the Fourier series coefficients of $f(t)$ as

$$F[m] = \frac{1}{\tau} \int_{-\tau/2}^{\tau/2} f(t)e^{-2\pi imt/\tau}\, dt.$$

Then, $F[m] = NG[m]$ for $|m| \leq \lfloor B\tau/2 \rfloor$.
Proof: See Appendix A. ∎

C. From Fourier Series Coefficients to Signal Parameters

When $f(t)$ is a τ-periodic stream of K weighted Dirac pulses, its Fourier series coefficients are given by

$$F[m] = \frac{1}{\tau} \sum_{k=0}^{K-1} c_k e^{-2\pi it_k m/\tau}, \quad (2)$$

where $m \in \mathbb{Z}$.

Thus, our problem is to calculate the parameters c_k and t_k from a finite number of Fourier series coefficients, $F[m]$. This is a common problem in spectral analysis, which can be solved using, for example, *Prony's method*. The idea is to find a filter $A[m]$ that, when convolved with the signal, produces a null signal; in other words, $A[m]$ annihilates the signal $F[m]$:

$$(F * A)[m] = 0, \qquad \forall m \in \mathbb{N}. \quad (3)$$

For this reason, *Prony's method* is also referred to as the *annihilating filter method*. Since $F[m]$ has the form given in Eq. (2), it is annihilated by the filter whose z-transform is

$$A(z) = \sum_{k=0}^{K} A[k]z^{-k} = \prod_{k=1}^{K}(1 - u_k z^{-1}), \quad (4)$$

where $u_k = e^{-2\pi it_k/\tau}$.

Therefore, given the roots of the annihilating filter, we can compute the locations of the Diracs as $t_k = -\frac{\tau \angle u_k}{2\pi}$, where $\angle u_k$ is the phase of u_k. It remains to show how to find the annihilating filter coefficients, and thus the roots. To that end, Eq. (3) can be rewritten in matrix form as

$$\boldsymbol{Sh} = \boldsymbol{0}, \quad (5)$$

where $\boldsymbol{S}$ is a rank deficient Toeplitz matrix (of rank K) formed by consecutive values of $F[m]$ and $\boldsymbol{h}$ is a vector containing the $K + 1$ annihilating filter coefficients $A[k]$. Observe that we need at least $2K + 1$ consecutive coefficients to solve the system of equations given in (5) and fully recover the annihilating filter; the set of coefficients we choose does not necessarily need to be located around the DC component, but the most efficient strategy is to pick $F[m]$ for $m = -K - 1, \ldots, K$, as it minimizes the frequency at which the signal is sampled. In this case, the bandwidth B of the low-pass filter $\varphi(t)$ should be larger than $(2K + 1)/\rho$, which represents one more sample per unit of time than the rate of innovation of the signal.

Finally, we need to show how to retrieve the coefficients $\{c_k\}_{k=0}^{K-1}$, which, after calculating $\{t_k\}_{k=0}^{K-1}$, is a simpler linear problem. Taking K consecutive Fourier coefficients (or more)
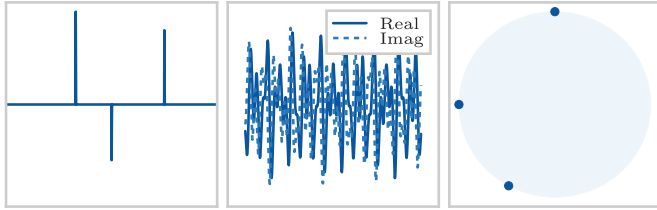
(a) Time domain   (b) Frequency domain   (c) Roots of the filter

Fig. 5. Classical FRI for streams of Diracs: (a) time-domain signal; (b) real and imaginary parts of the spectrum, which is made of a sum of complex exponentials; (c) complex roots of the z-transform of the filter annihilating the spectrum.

and letting

$$
V = \begin{bmatrix}
1 & 1 & \cdots & 1 \\
u_0 & u_1 & \cdots & u_{K-1} \\
\vdots & \vdots & \ddots & \vdots \\
u_0^{K-1} & u_1^{K-1} & \cdots & u_{K-1}^{K-1}
\end{bmatrix},
$$

$$
f = \begin{bmatrix} F[0] \; F[1] \; \cdots \; F[K-1] \end{bmatrix}^T,
$$

$$
\text{and} \quad c = \begin{bmatrix} c_0 \; c_1 \; \cdots \; c_{K-1} \end{bmatrix}^T,
$$

Eq. (2) can be written as

$$
\frac{1}{\tau} V c = f. \tag{6}
$$

The Vandermonde matrix V is non-singular since $u_k \neq u_\ell$ for any $k \neq \ell$. Therefore, the amplitudes c can be recovered by $c = \tau V^{-1} f$.

After this brisk review of FRI sampling and reconstruction, we are ready to extend the model to more general pulses.

III. PULSES WITH VARIABLE WIDTH - THE 1D CASE

This section demonstrates how to generalize the Dirac pulse model to account for signals with variable width and asymmetry. More precisely, we show how to sample and reconstruct variable-width pulses of the form

$$
f_k(t) = f_k^s(t) + f_k^a(t), \tag{7}
$$

where

$$
f_k^s(t) = c_k \sum_{n \in \mathbb{Z}} \frac{r_k}{\pi \left(r_k^2 + (t - t_k - n\tau)^2 \right)}, \tag{8}
$$

and

$$
f_k^a(t) = d_k \sum_{n \in \mathbb{Z}} \frac{t - t_k - n\tau}{\pi \left(r_k^2 + (t - t_k - n\tau)^2 \right)}. \tag{9}
$$

Here, $f_k^s(t)$ and $f_k^a(t)$ are symmetric and anti-symmetric components of the pulse, respectively.

These expressions follow naturally from a simple observation about the roots of the annihilating filter. In particular, in the classical FRI theory for streams of Diracs, the roots of the z-transform of the annihilating filter all have a magnitude of 1 and thus lie on the unit circle, as shown in Fig. 5. Since these roots are complex numbers, they can potentially carry more
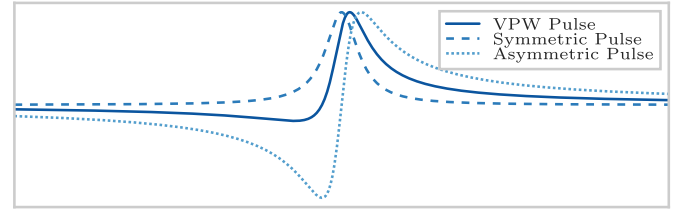

Fig. 6. Continuous line: VPW pulse; dashed line: symmetric pulse; dotted line: asymmetric pulse. The parameters used are $t_0 = 0.5, r_0 = 0.03$ and $c_0 = d_0 = 0.1$. Only one period is shown.

information than just the location of the pulses; indeed, in the classical case, only the phase is extracted and the information regarding the magnitude is simply disregarded.

A. Variable Width

To leverage the information about the magnitude of the root and control the width of the pulses, we introduce a new parameter r_k called the *width* and we redefine the roots as

$$
u_k = e^{-2\pi(r_k + it_k)/\tau}. \tag{10}
$$

With this adjustment and for a single root, we can build the following Fourier series coefficients:

$$
F_k^s[m] = \frac{c_k}{\tau} e^{-2\pi(r_k|m| + it_k m)/\tau}, \qquad m \in \mathbb{Z}. \tag{11}
$$

The absolute value of m is used to ensure conjugate symmetry of the spectrum. The corresponding τ-periodic time-domain signal is precisely the Lorentzian function described in (8).

B. Asymmetry

Pulses such as (8) are symmetric; in many practical applications however, asymmetric pulses are desirable. The symmetric formulation can be generalized by considering the following Fourier series coefficients:

$$
F_k^a[m] = -\frac{id_k}{\tau} \operatorname{sgn}(m) e^{-2\pi(r_k|m| + it_k m)/\tau}, \quad m \in \mathbb{Z}, \tag{12}
$$

where $d_k \in \mathbb{R}$. The signal F_k^a is essentially the Hilbert transform of F_k^s and its (τ-periodic) time-domain representation is given in (9). The combination of the symmetric and asymmetric pulses (7) is the main building block of our signal model and we call it *variable pulse width* (VPW) pulse. An example of a VPW pulse and its decomposition into symmetric and asymmetric components is shown in Fig. 6. An alternative formula for $f_k(t)$ that avoids the infinite sum is given by[3]

$$
f_k(t) = \frac{c_k}{\tau} \frac{1 - |z_k(t)|^2}{(1 - z_k(t))(1 - z_k^*(t))}
$$
$$
+ \frac{d_k}{\tau} \frac{2 \operatorname{Im}\{z_k(t)\}}{(1 - z_k(t))(1 - z_k^*(t))}, \tag{13}
$$

where $z_k(t) = e^{2\pi(-r_k + i(t-t_k))/\tau}$.

VPW pulses are in essence a generalization of Diracs; by setting $d_k = 0$ and taking the limit of $f_k(t)$ when r_k goes to zero, we obtain a Dirac delta with amplitude c_k located at time t_k.

---

[3]The details of the derivations are provided in Appendix B.

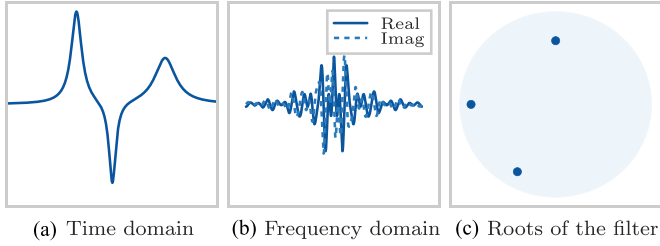(a) Time domain   (b) Frequency domain   (c) Roots of the filter

Fig. 7. VPW pulses: (a) time-domain signal; (b) real and imaginary parts of the spectrum, which is made of a damped sum of complex exponentials; (c) complex roots of the z-transform of the filter annihilating the spectrum.

As shown in Fig. 7, several VPW pulses can be combined to form a *VPW signal*:

$$f(t) = \sum_{k=0}^{K-1} f_k(t) \tag{14}$$

$$F[m] = \sum_{k=0}^{K-1} F_k[m], \qquad m \in \mathbb{Z}. \tag{15}$$

C. Recovery of VPW Pulse Parameters

We now show how to sample and reconstruct a sum of Lorentzian pulses.

Theorem 1: Let f be the sum of K continuous τ-periodic asymmetric Lorentzians as described in (14). Assume we take N samples of f, according to Fig. 4; i.e., we obtain

$$g[n] = \langle f(t), \varphi(t - nT) \rangle, \quad n \in 0, 1, \ldots, N-1.$$

Furthermore, assume $\varphi(t)$ is such that the Fourier series coefficients of f, denoted $F[m]$, can be obtained from the samples, for $0 \le m \le \lfloor (N-1)/2 \rfloor$. For example, if $\varphi(t) = B\mathrm{sinc}(Bt)$, with $T = \tau/N$ and B satisfying $N = 2\lfloor B\tau/2 \rfloor + 1$, this is achieved by Lemma 1.

Then, if $N \ge 4K + 1$, the samples, $\{g[n] : n \in 0, 1, \ldots, N-1\}$, are sufficient to recover f.

Proof: To prove the theorem, we first show how to construct the unique annihilating filter for the Fourier series coefficients of f. We then show how the $4K$ parameters of f can be retrieved from this filter.

Let

$$A(z) = \sum_{k=0}^{K} A[k] z^{-k} = \prod_{k=0}^{K-1} (1 - u_k z^{-1}), \tag{16}$$

and $F[m]$ be the Fourier series coefficients of f. Then,

$$
\begin{aligned}
(A * F)[m] &= \sum_{\ell=0}^{K} A[\ell] F[m-\ell] \\
&= \sum_{\ell=0}^{K} \sum_{k=0}^{K-1} (c_k - id_k) A[\ell] u_k^{m-\ell} \\
&= \sum_{k=0}^{K-1} (c_k - id_k) \left(\sum_{\ell=0}^{K} A[\ell] u_k^{-\ell} \right) u_k^m = 0.
\end{aligned}
$$

Algorithm 1: VPW-FRI Parameters Estimation.

Input: $N \ge 4K + 1$ discrete samples $x[n]$ ($n = 0, 1, \ldots N - 1$) of a VPW-FRI continuous-time signal $x(t)$ and K its number of pulses.

Output: The locations, widths and (complex) amplitudes $\{t_k, r_k, c_k\}_{k=0}^{K-1}$ of $x(t)$.

1: Compute $X[m]$ the FFT of $x[n]$
2: Build the Toeplitz matrix $\boldsymbol{S}$ from $X[m]$ for the positive part of the spectrum
3: Optionally denoise $\boldsymbol{S}$ or $x[n]$ according to one of the strategies of Section III-E
4: Find the annihilating filter coefficients by solving $\boldsymbol{Sh} = \boldsymbol{0}$
5: Compute u_k, the roots of $\boldsymbol{h}$
6: $t_k = -\frac{\tau \angle u_k}{2\pi}$
7: $r_k = \frac{\tau \log |u_k|}{2\pi}$
8: Build the Vandermonde matrix $\boldsymbol{V}$ from the roots u_k
9: Solve $\frac{1}{\tau} \boldsymbol{V} \boldsymbol{c} = \boldsymbol{x}$, where $\boldsymbol{x}$ contains the FFT coefficients $X[m]$ and $\boldsymbol{c} \in \mathbb{C}$ the amplitudes c_k
10: **return** $\{t_k, r_k, c_k\}_{k=0}^{K-1}$

Note that we have restricted ourselves to positive indices ($m \ge 0$). This is because, due to its decaying nature, the VPW spectrum has a *cusp* at $m = 0$, which prevents us from performing the annihilation on both positive and negative values of the spectrum.

Finding the unique filter A requires at least $2K + 1$ Fourier series coefficients corresponding to non-negative frequencies. By Lemma 1, this can be achieved with $N \ge 4K + 1$ samples of f.

As in the case of streams of Diracs, the locations $\{t_k\}_{k=0}^{K-1}$ can be retrieved from the K roots of $A(z)$: $u_k = e^{-2\pi i t_k / \tau}$. Moreover, the widths $\{r_k\}_{k=0}^{K-1}$ are given by $r_k = \tau \log |u_k| / (2\pi)$. Finally, the parameters $\{c_k\}_{k=0}^{K-1}$ and $\{d_k\}_{k=0}^{K-1}$ are retrieved by solving (6) with

$$\boldsymbol{c} = \begin{bmatrix} c_0 + id_0 & c_1 + id_1 & \cdots & c_{K-1} + id_{K-1} \end{bmatrix}^T. \qquad \blacksquare$$

The VPW-FRI estimation procedure is summarized in Algorithm 1. Since the algorithm requires $4K + 1$ samples to retrieve $4K$ parameters, it is only one sample away from the rate of innovation.

D. Sampling Kernels

Although in this paper we concentrate on the ideal low-pass filter or Dirichlet sampling kernel, VPW signals are fully compatible with the more general kernels presented in the FRI literature. In particular, as outlined in Theorem 1, we can use any kernel that allows us to retrieve enough Fourier series coefficients of the continuous-time signal. In [5], it is shown that this can be achieved with kernels that satisfy the generalized Strang-Fix conditions, such as *cardinal exponential splines* [27] or sum of sinc kernels [6].

Moreover, even if the sampling kernel does not meet these conditions, we can relax the requirement of perfect reconstruc-
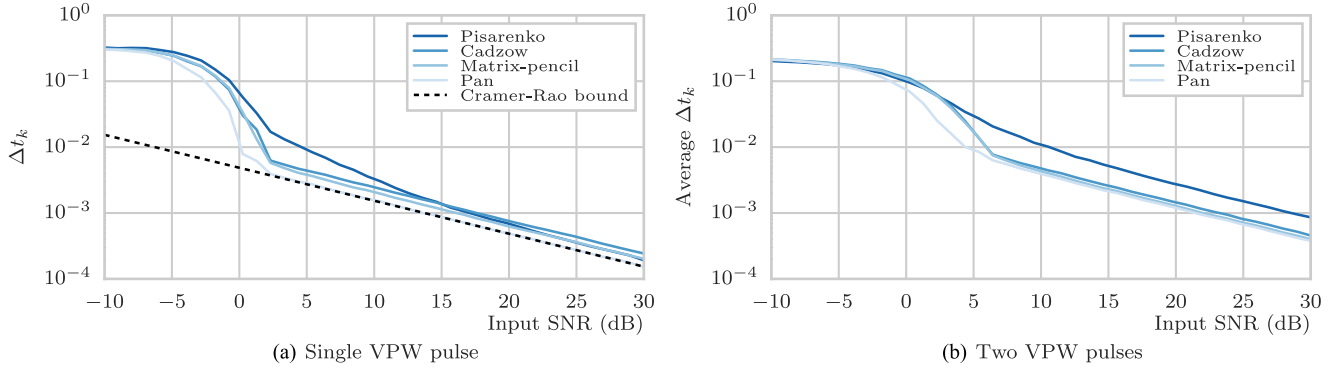
Fig. 8.   Comparison of the effect of different denoising approaches on the estimation of the location of (a) one and (b) two VPW pulses: 10 Fourier coefficients are used in the estimation procedure and the experiment has been repeated 5,000 times and averaged. For the single pulse case, the Cramér-Rao bound for the estimation of the location t_0 has been computed (dashed line).

tion and achieve accurate approximations using either the correction filter proposed in [28] or the framework set out in [29].

E. Denoising Techniques

In practice, it is unlikely that a signal perfectly matches the VPW-FRI model. Even if this is the case, the acquisition device might introduce some noise in the process. This justifies the need for a denoising block before the estimation of the parameters. Here, we review a few methods borrowed from spectral estimation theory [30] and evaluate their performance on our estimation procedure.

We model the noise ϵ_n as i.i.d. additive Gaussian with zero mean and covariance matrix $\sigma^2 I$. Furthermore, we assume that it is introduced after sampling; that is, the noisy measurements are

$$\tilde{f}[n] = f[n] + \epsilon_n \qquad n = 0, 1, \ldots N - 1. \qquad (17)$$

We denote the noisy Toeplitz matrix built from the samples $\tilde{f}[n]$ as $\tilde{S}$. Several denoising strategies have been introduced in the classical FRI framework and they are directly applicable to the VPW-FRI model. We summarize and compare the main techniques below.

1) Cadzow: Cadzow's algorithm [31] is an iterative method that alternatively enforces rank K and a Toeplitz structure for $\tilde{S}$. The low rank approximation is performed by computing the SVD of $\tilde{S}$ and keeping only its K largest singular values. To ensure that $\tilde{S}$ is Toeplitz, we simply average across diagonals. In practice, the method seems to converge to a solution that satisfies both properties, but it should be noted that this has not been proven theoretically. Furthermore, there is no guarantee that the final solution is close, in any way, to S.

2) The Matrix Pencil: Another widely used technique is the matrix pencil method [32] — also known as ESPRIT [33] — which exploits the rotational invariance property of the underlying signal subspace. Unlike Cadzow, it is a non-iterative algorithm.

3) Pisarenko: Pisarenko's method [34] estimates the annihilating filter by simply extracting the last column of the matrix V of the SVD of $\tilde{S} = U\Lambda V^*$. Indeed, the last columns of V provide an orthogonal basis for the nullspace of $\tilde{S}$ and in our case,

the nullspace of S is one-dimensional. Note that Pisarenko's method or the matrix pencil can actually be used sequentially after Cadzow.

4) Pan: The last method we survey is inspired by IQML [35] and has been recently introduced to the FRI framework by Pan in [36]. The annihilating filter h is found iteratively and its computation is formulated as

$$\min_{\hat{f}, h} \quad \|\tilde{f} - B\hat{f}\|_2^2$$

$$\text{subject to} \quad h * \hat{f} = 0 \qquad (18)$$

where B is a linear transformation that maps the signal that can be annihilated to the measurements. Unlike the other techniques, this approach is not applied to $\tilde{S}$ but directly to the measured signal: this can result in significant performance improvements since the transformation operation from the signal measurements to the signal to be annihilated sometimes dramatically alters the structure of the noise.

The main ingredient of the algorithm is the matrix B, which is different for every FRI problem. In the case of VPW-FRI, B is the transform that maps the denoised DFT coefficients $\hat{f}$ (corresponding to positive frequencies) to the discrete-time (noisy) measured signal $\tilde{f}$.

In Fig. 8, we compare the performance of the different denoising methods under various levels of noise for one and two VPW pulses. We observe that the matrix-pencil and Cadzow have a similar performance. Furthermore, we see that Pan's method outperforms all other methods by a fairly large margin.

For one pulse, we also derive and display the Cramér-Rao bound (CRB), which provides a theoretical lower bound on the best possible performance achievable by any unbiased estimator. The details of the computations of the CRB are provided in Appendix C. We see that, for a pulse with width $r_0 = 0.02$, the VPW-FRI estimation procedure is fully efficient down to SNRs of about 0 dB. Interestingly, the quality of estimation depends on the width parameter: the wider the pulse, the better the performance. This effect is showcased in the supplementary material provided with the paper.
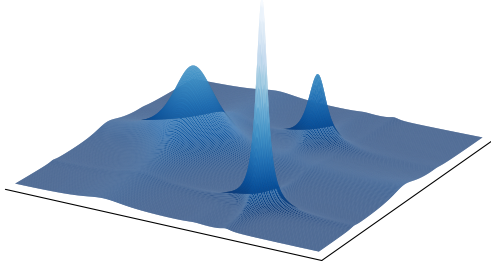
Fig. 9. Example of a 2D VPW signal composed of three pulses. The parameters used are $(x_k, y_k) = \{(0.6, 0.7), (0.7, 0.2), (0.3, 0.5)\}$, $r_k = \{0.035, 0.02, 0.085\}$, $c_k = \{0.75, 0.7, 5\}$ and $\tau = 1$.

IV. PULSES WITH VARIABLE WIDTH - THE 2D CASE

Before discussing the spherical case, we describe how the VPW model generalizes to 2D, by extending results on 2D streams of Diracs [37]. The added dimension introduces more flexibility in the model, allowing us to reproduce several different shapes. In what follows, we discuss four different pulse shapes, each offering different properties such as asymmetry or radiality. We also demonstrate how to estimate the parameters using the FRI algorithm.

We use the pair of variables (x, y) to index 2D spatial measurements and (m, n) to index 2D Fourier series coefficients. We also assume that the 2D sampling kernel is a separable kernel given by the tensor product of two 1D functions: $\varphi(x, y) = \varphi_x(x)\varphi_y(y)$. Therefore, the natural 2D extension of our model is given by the tensor product of two 1D VPW pulses, say $f_{k,x}(x)$ and $f_{k,y}(y)$. The Fourier series coefficients are given by

$$F_k[m, n] = \frac{1}{\tau}F_{k,x}[m]F_{k,y}[n]$$
$$= \frac{1}{\tau}e^{-\frac{2\pi}{\tau}i(x_k m + y_k n)}e^{-\frac{2\pi}{\tau}r_k(|m|+|n|)}, \quad (19)$$

Similarly, we define a VPW signal as a linear combination of K pulses; that is,

$$F[m, n] = \sum_{k=0}^{K-1} F_k[m, n], \qquad m, n \in \mathbb{Z}. \quad (20)$$

An example of a 2D VPW signal composed of three pulses is shown in Fig. 9.

A. Estimation of the Parameters

Observe that if we fix one of the frequency indices, say m, the expression (20) reduces to

$$F[n] = \frac{1}{\tau^2}\sum_{k=0}^{K-1} \tilde{c}_k e^{-\frac{2\pi}{\tau}iy_k n}e^{-\frac{2\pi}{\tau}r_k|n|}, \quad (21)$$

where $\tilde{c}_k = c_k e^{-i\frac{2\pi}{\tau}x_k m}e^{-\frac{2\pi}{\tau}r_k|m|}$. Now that we have reduced the problem to the 1D case, we can invoke Lemma 1 and Theorem 1 to connect the DFT coefficients with the Fourier series coefficients and guarantee the uniqueness of reconstruction of the continuous-space signal. As in the 1D case, we cannot

exploit the full spectrum for annihilation; due to the singularity at the DC component, we restrict ourselves to the positive indices m and n. We refer to the sets of positive n for a fixed m as *half-columns* and to the sets of positive m for a fixed n as *half-rows*. The coefficients used are part of what we call the *annihilation region*. The spectrum and the annihilation region of the signal described in Eq. (20) is shown in Fig. 10(a).

Following Eq. (21), we see that each half-column of the spectrum can be annihilated independently by the same filter. More precisely, we can build Toeplitz matrices with each of these half-columns and stack them on top of each other to obtain the matrix $\boldsymbol{S}^c$:

$$\boldsymbol{S}^c = \begin{bmatrix} \boldsymbol{S}_0^c \\ \boldsymbol{S}_1^c \\ \vdots \\ \boldsymbol{S}_{\lfloor M/2 \rfloor - 1}^c \end{bmatrix}, \quad (22)$$

where $\boldsymbol{S}_m^c$ is the Toeplitz matrix corresponding to the mth half-column of the spectrum. Similarly, by fixing n, the half-rows of the spectrum can be annihilated and we can create the matrix $\boldsymbol{S}^r$. Both $\boldsymbol{S}^c$ and $\boldsymbol{S}^r$ are block Toeplitz and can be used independently to recover the locations along the x and y axes, respectively, following the procedure presented in Section III-C.

To solve the annihilation equation, each half-column and half-row of the spectrum must have at least $K + 1$ elements. Therefore, we need at least $(2K + 1)^2$ samples to recover $4K$ parameters; unlike in the 1D case, we need to sample the signal at a significantly higher rate than the rate of innovation to achieve perfect recovery.

Another detail needs to be taken care of in 2D. Since we reformulate the estimation of the parameters as two 1D problems, we obtain two sets of locations $\{x_k\}_{k=0}^{K-1}$ and $\{y_k\}_{k=0}^{K-1}$ but we do not know the *corresponding pairs* of these recovered locations. If K is small, we can solve this issue by combinatorial search over all the possible pairs; for each combination of locations, we compute the amplitude of the corresponding pulses and compare the reproduced signal with the original. We then select the pair whose reproduction corresponds (or, in the noisy case, is closest) to the measured signal.

Unfortunately, this brute-force approach quickly becomes impractical as K gets larger. To cope with this issue, we use the following heuristic. Instead of searching through all pairs of K pulses, we suppose that all combinations are valid, that is, we have K^2 potential pulses. We then compute the amplitude of each of these pulses by solving Eq. (6). Finally, the correct pairs of locations are revealed by the K largest pulse amplitudes. This simple approach has no proven correctness, but experiments have shown that it performs well in practice.

B. Radial Formulation

Observe that the 2D Lorentzian pulse described in Eq. (20) has heavy tails along the two main axes. Perhaps a more suitable
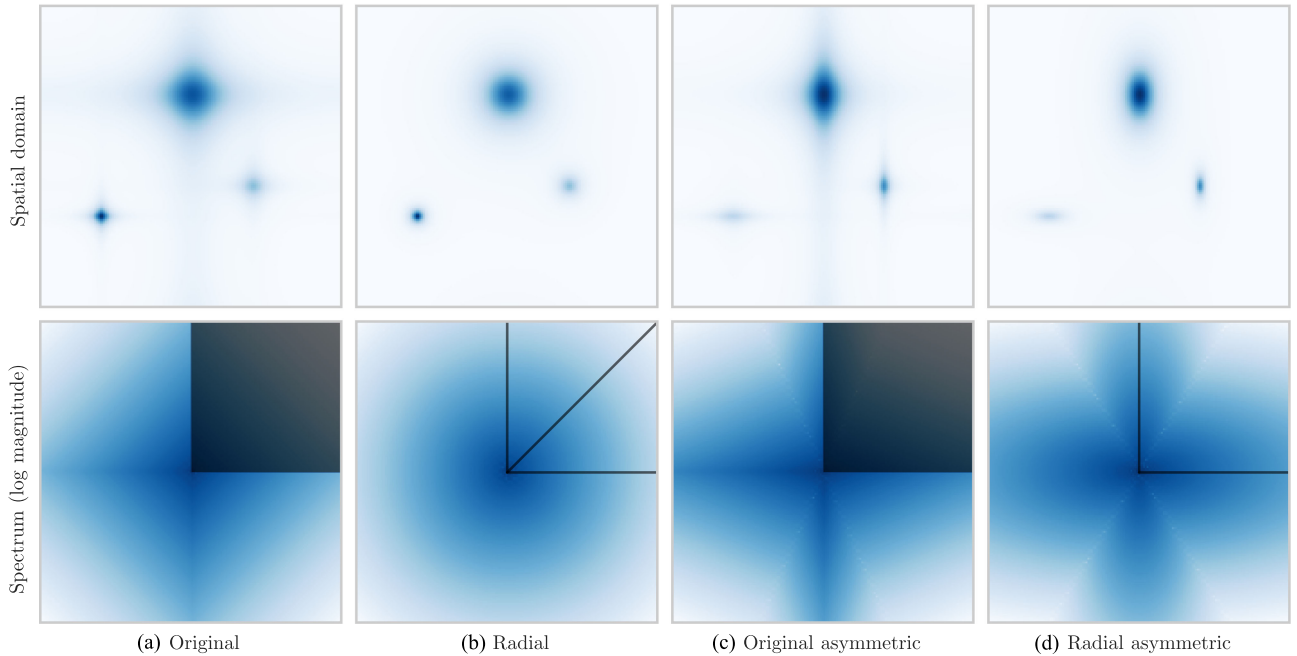
Fig. 10.   Family of 2D VPW pulses and their spectrum: (a) original formulation, (b) radial formulation, (c) original anisotropic formulation, and (d) radial anisotropic formulation. The area of annihilation is represented by a black overlay.

model is given by

$$F[m,n] = \frac{1}{\tau^2} \sum_{k=0}^{K-1} c_k e^{-\frac{2\pi}{\tau} i (x_k m + y_k n)} e^{-\frac{2\pi}{\tau} r_k \sqrt{m^2 + n^2}}, \quad (23)$$

for $m, n \in \mathbb{Z}$. The difference with Eq. (20) resides in the damping factor: in (20), the exponent is the ℓ^1-norm of the Fourier indices, while in (23), the ℓ^2-norm is used. As a result, the spectrum of (23) does not exhibit the sharp transitions of the original representation; this corresponds to radial pulses in the primal domain.

The disadvantage of this representation is that its region of annihilation is not as large as in the previous model. Such a signal can be annihilated along the DC components, that is, when we set either m or n equal to zero, but also along the diagonal line defined as $m = n$. More generally, the spectrum can be annihilated along the line defined as $m = an$ or $n = bm$ for $a, b \in \mathbb{N}$, as long as we have a sufficient number (at least $K + 1$) of coefficients on that line. The annihilation region of this signal can be seen in Fig. 10(b).

C. Anisotropic Pulses

We can further generalize the 2D pulse shape and introduce anisotropy in the model by using a different width parameter for each axis. For instance, the radial formulation can be rewritten as

$$F[m,n] = \frac{1}{\tau^2} \sum_{k=0}^{K-1} c_k e^{-\frac{2\pi}{\tau} i (x_k m + y_k n)} e^{-\frac{2\pi}{\tau} \sqrt{r_{m,k}^2 m^2 + r_{n,k}^2 n^2}},$$

where $r_{m,k}$ and $r_{n,k}$ are the respective widths along the main axes. Examples of anisotropic pulses for both the original and radial formulations are illustrated in Fig. 10(c) and 10(d).

V. PULSES WITH VARIABLE WIDTH - ON THE SPHERE

Recently, FRI theory has been developed for spherical signals by Deslauriers-Gauthier and Marziliano [38], and the spectrum usage has been further improved by Dokmanić and Lu [39]. Building on their work, we show that the VPW framework can also be applied to the spherical domain.

A. Preliminaries

We start by reviewing a few definitions on spherical signal processing. We refer the interested reader to [39], [40] for a more thorough introduction.

Functions on the sphere $\mathbf{S}^2 = \{\boldsymbol{x} \in \mathbb{R}^3 : \|\boldsymbol{x}\| = 1\}$ are usually characterized by elevation and azimuth angles, denoted θ and ϕ. The spherical Dirac delta is defined as

$$\delta(\theta, \phi; \theta_0, \phi_0) = \delta(\cos\theta - \cos\theta_0)\delta(\phi - \phi_0), \quad (24)$$

where (θ_0, ϕ_0) denotes its location. A sum of weighted Diracs is then defined as

$$f(\theta, \phi) = \sum_{k=0}^{K-1} c_k f_k(\theta, \phi) = \sum_{k=0}^{K-1} c_k \delta(\theta, \phi; \theta_k, \phi_k). \quad (25)$$

On the sphere, the spherical harmonic (SH) decomposition provides the analog to the Fourier series expansion. As shown in [39], for the signal $f(\theta, \phi)$, it is given by[4]

$$\hat{f}_{\ell,m} = N_{\ell,m} \sum_{k=0}^{K-1} c_k P_{\ell,|m|}(\cos\theta_k) e^{-im\phi_k}, \quad (26)$$

---

[4]In some of the spherical harmonics related literature, indices are expressed with a superscript/subscript notation, e.g. f_ℓ^m. We choose to include both indices in the subscript, to avoid confusion with powers.
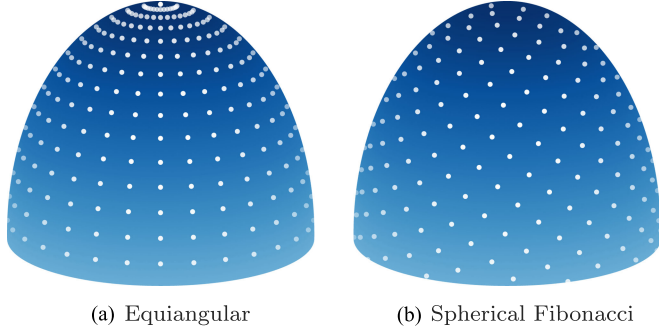
(a) Equiangular   (b) Spherical Fibonacci

Fig. 11. Two possible sampling schemes on the hemisphere: (a) equiangular and (b) spherical Fibonacci sampling. Each of these hemispheres contains 300 samples.

with $\ell = 0, 1, \ldots$ and $|m| < \ell$. Here, $N_{\ell,m}$ is a normalization constant, and $P_{\ell,m}$ denotes the *associated Legendre polynomial*, defined as

$$P_{\ell,m}(\cos\theta_k) = (-1)^m (\sin\theta_k)^{|m|} \frac{d^m}{d(\cos\theta_k)^m} P_\ell(\cos\theta_k),$$

where $P_\ell(\cdot)$ is the Legendre polynomial. Since P_ℓ is of degree ℓ, the degree of P_ℓ^m is at most $\ell - m$.

B. Sampling the Hemisphere

Several strategies can be employed to sample a hemisphere; the simplest one is *equiangular sampling*, i.e. uniform sampling along the azimuth and elevation angles. The drawback of this approach is that it results in non-uniform sampling of the hemisphere, as illustrated in Fig. 11(a): the area around the pole has a higher density than the area near the equator.

Another simple technique that leads to almost uniform sampling is *spherical Fibonacci sampling* and can be seen in Fig. 11(b). In the remainder of the paper, we nevertheless use the former strategy since in practice it allows for a faster and more natural scanning of the hemisphere.

C. From Samples to Spherical Harmonics Coefficients

We define the *bandwidth* L of a spherical signal as the smallest L such that $\hat{f}_{\ell,m} = 0$ for all $\ell \geq L$. In [41, Theorem 3], it is proved that the SH coefficients of a spherical band-limited function can be exactly computed given a minimum number of equiangular samples. For completeness, we restate this result below.

Theorem 2 (Driscoll and Healy 1994): Let g be the result of convolving the spherical signal f with an ideal lowpass filter φ with bandwidth L. Assume we further sample g at points $(\theta_n, \phi_{n'})$, where $\theta_n = \frac{\pi n}{2L}$ for $n = 0, 1, \ldots, 2L - 1$ and $\phi_{n'} = \frac{\pi n'}{L}$ for $n' = 0, 1, \ldots, 2L - 1$. Then the SH coefficients of g are given by

$$\hat{g}_{\ell,m} = \frac{\sqrt{2\pi}}{2L} \sum_{n=0}^{2L-1} \sum_{n'=0}^{2L-1} \alpha_n g(\theta_n, \phi_{n'}) Y_{\ell,m}^*(\theta_n, \phi_{n'}), \quad (27)$$

where $Y_{\ell,m}^*(\theta_n, \phi_{n'})$ are the SH basis functions evaluated at $(\theta_n, \phi_{n'})$ and $\{\alpha_n\}_{n=0}^{2L-1}$ is the unique solution to

$$\sum_{n=0}^{2L-1} \alpha_n P_\ell(\cos\theta_n) = \sqrt{2}\delta_{\ell 0},$$

for $\ell = 0, 1, \ldots, 2L - 1$.

Proof: See [41]. ∎

Other sampling theorems exists for different sampling strategies: for example, if we choose N samples uniformly at random on the sphere, we can retrieve the SH coefficients of g when $N \geq L^2$, with probability one [42]. Finally, remark that if we have fewer samples, the spectrum often decays fast enough for (27) to yield a good approximation.

D. Retrieving Diracs from Sectoral SH Coefficients

In [38], the locations and amplitudes of the Diracs are retrieved from the *sectoral* spherical harmonics coefficients, which are the coefficients where $m = \ell$:

$$\hat{f}_{\ell,\ell} = \tilde{N}_{\ell,\ell} \sum_{k=0}^{K-1} c_k (\sin\theta_k)^\ell e^{-i\ell\phi_k}. \quad (28)$$

Here, $\tilde{N}_{\ell,\ell}$ is a constant that encompasses $N_{\ell,\ell}$ as well as the appropriate associated Legendre coefficient. Since $\tilde{N}_{\ell,\ell}$ is known and nonzero, we can divide the expression in (28) by it to obtain

$$g_\ell = \frac{\hat{f}_{\ell,\ell}}{\tilde{N}_{\ell,\ell}} = \sum_{k=0}^{K-1} c_k u_k^\ell, \quad (29)$$

where $u_k = \sin\theta_k e^{-i\phi_k}$. We recognize in (29) the familiar expression of a sum of complex exponentials, which can be annihilated using Prony's method to recover the roots $\{u_k\}_{k=0}^{K-1}$. To find the annihilating filter, we need the bandwidth L of the signal to be at least $2K$.

E. Retrieving VPW Pulses from Sectoral SH Coefficients

We now extend the result of [38], adding width to the Diracs. Like the Cartesian case, we can introduce a damping factor $0 \leq r_k \leq 1$, which reduces the magnitude of the signal at higher frequencies. The spherical VPW pulses are thus defined as

$$\hat{f}_{\ell,m} = N_{\ell,m} \sum_{k=0}^{K-1} c_k P_{\ell,|m|}(\cos\theta_k) e^{-im\phi_k} e^{-r_k(\ell+|m|)}. \quad (30)$$

An example of a spherical VPW signal composed of two pulses is shown in Fig. 12. The following theorem shows how to retrieve the parameters of this new signal model.

Theorem 3: Let f be the sum of K continuous spherical VPW pulses as in (30). The spherical harmonics coefficients $\hat{f}_\ell^m$ with $m = 0, 1, \ldots, 2K - 1$ and $|\ell| < m$ are sufficient to recover f.

Proof: As in [38], we can compute

$$g_\ell = \frac{\hat{f}_{\ell,\ell}}{\tilde{N}_{\ell,\ell}} = \sum_{k=0}^{K-1} c_k u_k^\ell, \quad (31)$$
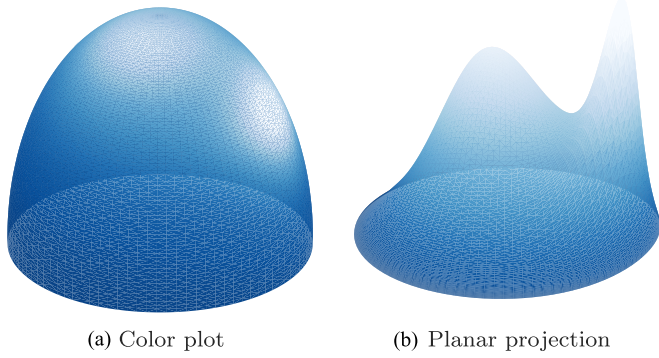
(a) Color plot   (b) Planar projection

Fig. 12. Example a spherical VPW signal composed of two pulses. There are different ways to display (hemi)spherical signals: (a) a plot of a unit hemisphere where the brightness of colors represents the intensity, and (b) a projection of the same signal onto the xy plane with the intensity represented on the z axis.

where we now have

$$\hat{f}_{\ell,\ell} = \tilde{N}_{\ell,\ell} \sum_{k=0}^{K-1} c_k (\sin\theta_k)^\ell e^{-i\ell\phi_k} e^{-2r_k \ell}. \qquad (32)$$

The roots, which can be found via the annihilating filter method, become $u_k = \sin\theta_k e^{-i\phi_k} e^{-2r_k}$. We can extract the azimuth ϕ_k from the phase of u_k, but the magnitude of u_k depends on both θ_k and r_k and at this point we cannot resolve the ambiguity. What we can do however is find the amplitude c_k by solving the linear system in Eq. (31).

Furthermore, for $m = \ell - 1$ and $\ell \geq 1$, we have

$$\hat{f}_{\ell,\ell-1} = \tilde{N}_{\ell,\ell-1} \sum_{k=0}^{K-1} c_k e^{-r_k} \cos\theta_k u_k^{\ell-1}, \qquad (33)$$

$$g_{\ell-1} = \frac{\hat{f}_{\ell,\ell-1}}{\tilde{N}_{\ell,\ell-1}} = \sum_{k=0}^{K-1} c_k b_k u_k^{\ell-1}, \qquad (34)$$

where $b_k = e^{-r_k}\cos\theta_k$. Note that, since we know $\hat{f}_\ell^{\ell-1}$ and $u_k^{\ell-1}$, we can compute b_k. Moreover, the knowledge of the azimuth ϕ_k enables us to calculate

$$a_k = \frac{u_k}{e^{-i\phi_k}} = \sin\theta_k e^{-2r_k}. \qquad (35)$$

Finally, combining a_k and b_k and assuming a_k is nonzero, we compute the width

$$r_k = \frac{1}{2} \log\left(-b_k^2 + \frac{1}{2a_k^2}\sqrt{4a_k^2 + b_k^4} \right) \qquad (36)$$

by keeping only the positive solution to the square root. Observe that b_k is raised to the 4th power and a_k is squared. Furthermore, in the noiseless case,[5] both terms are real, so the expression under the square root is always positive. Using this result, we calculate the elevation as $\theta_k = \arccos\left(\frac{b_k}{e^{r_k}}\right)$. ∎

Like [38], our method suffers from a relatively poor spectrum usage. Indeed, only a small number of the spherical harmonics coefficients are used for the parameters' estimation.[6] In

[5]In the noisy case, it is possible that b_k and a_k are complex. To prevent negative terms in the square root, we project them onto the real axis.

[6]A better spectrum usage was proposed in [39], but we have been unable to generalize their method to variable width pulses.
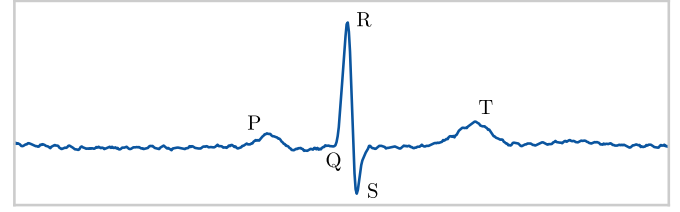

Fig. 13. An ECG recording from the MIT-BIH arrhythmia database and the P, Q, R, S, and T pulses annotated.

Section VII, we address this shortcoming by adding a nonlinear optimization routine that uses the entire spectrum to refine the VPW estimation.

VI. APPLICATION I: ECG COMPRESSION

Several signals can be approximated using VPW pulses; in the next two sections, we show two practical cases where VPW pulses can be used. In the first example, we consider the problem of sampling and representing electrocardiogram (ECG) signals. ECG signals are composed of five types of pulses, labeled P, Q, R, S, T in the medical literature (see also Fig. 13). Each of these pulses represents a different stage in the activity of the heart and each has a different width and amplitude; while the P and T pulses are fairly wide, the Q,R,S pulses are narrow and close to each other. Throughout this section, we use ECG records from the MIT-BIH database [19].

A. Practical Considerations

There are a number of details to pay attention to in order to bridge the gap between theory and practice and implement the VPW algorithm on real ECG records.

First, even though real ECG signals are not periodic, they exhibit a cyclic behavior since the heart goes through the same states during each heart beat. We propose therefore to estimate each heart beat independently; this requires the segmentation of ECG records via the detection of the QRS complex [43]: the segmentation is performed between the T and P wave, where the activity of the heart is minimal. Each heart beat is then processed by Algorithm 1. Since VPW pulses are periodic, we only extract one period of each before concatenating them. To allow for a smoother transition between two consecutive reconstructed beats, we make the extracted beats slightly overlap and connect their reconstructed versions together using pairs of cosine-shaped windows.

Second, we need to choose a value for K, the number of VPW pulses. Since a heart beat is naturally composed of 5 waves, it is tempting to represent them with $K = 5$ pulses. Practical experiments however have shown that 6 or 7 pulses give more flexibility to the model. More pulses can be added, at the expense of a lower resulting compression ratio.

Finally, to counteract noise and model mismatch, we make use of one or more of the denoising techniques discussed in Section III-E.

B. Acquisition at Low Sampling Rates

In many applications, power consumption is critical and it is thus beneficial to reduce the sampling frequency at the point of acquisition. While our framework guarantees perfect reconstruction when sampling at or above the rate of innovation *in the noiseless case*, noise and model mismatch induce inaccuracies in the estimation of the parameters. Furthermore, as high frequency information is lost during the sampling procedure, a strong prior is needed to recover the original signal from low sampling rates. To demonstrate the validity of our approach, we now compare the VPW-FRI algorithm to sinc interpolation.

We use the *signal-to-residual ratio* (SRR) to quantify the quality of the reconstruction:

$$\mathrm{SRR} = 20 \log_{10} \left(\frac{\|\boldsymbol{f} - \mu_f\|_2}{\|\boldsymbol{f} - \hat{\boldsymbol{f}}\|_2} \right),$$

where μ_f is the average value of $\boldsymbol{f}$.[7]

We run our algorithm on the 48 different records from the MIT-BIH arrhythmia database. These measurements are originally sampled at 360 Hz. To simulate different acquisitions at lower sampling rates, we lowpass filter and downsample the original signal. From the subsampled signals, we try to reconstruct the original signal with both VPW-FRI (using 5 pulses per heart beat) and sinc interpolation. Since a 5 pulse model requires at least 20 samples per beat and the heart rate of the ECG records varies between 60 and 80 beats per minute, the lowest sampling rate we can theoretically achieve is around 25 Hz.

We observed two distinct regimes: on ECG records composed of broader pulses and containing most of the energy in the low frequency part of the spectrum, sinc interpolation surpasses VPW-FRI. On the other hand, on records having sharper pulses with a spectrum not only concentrated around the low frequencies, VPW-FRI outperforms sinc interpolation *at low sampling rates*. To stress this behavior, we split the database into two sets: the first one — called the *high frequency set* and containing 12 records with higher frequency components — and the second — denoted the *low frequency set* and containing all other records. The results on these two sets are illustrated in Fig. 14. For the sake of completeness, we also show results on the set of all records. As expected from a nonlinear parametric estimation technique, we attain very good performance, provided that the data is well fit by the model.

As a final remark, notice that the error rate for VPW-FRI is achieved using only a limited number of coefficients (about 25 per second), while sinc interpolation uses up to 60 coefficients per second. This suggests that the model could be particularly suited to compression.

C. Compression

In compression, the goal is to minimize both the reconstruction error as well as the number of parameters needed per second.[8] The difference with the previous experiment is that we

---

[7]The mean is subtracted from $\boldsymbol{f}$ to avoid artificially large SRR values.
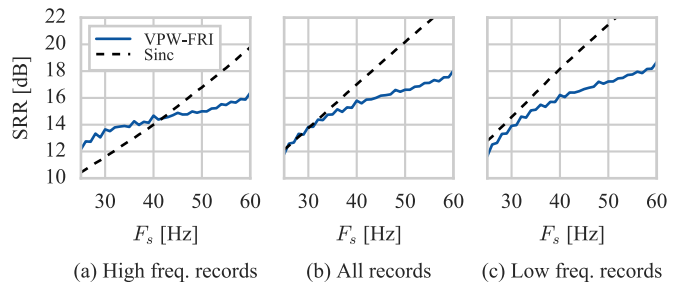[8]We do not discuss here quantization of the parameters.

Fig. 14. SRR for acquisition of ECG signals at low sampling rates on the first 20 seconds of the records of the MIB-BIH database: (a) average score on records containing higher frequency components; (b) average score on all the records; (c) average score on records containing less high frequencies.
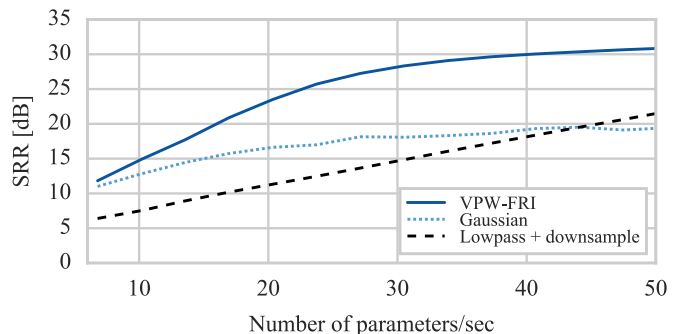

Fig. 15. SRR for ECG compression on the first 20 seconds of all records of the MIB-BIH database for three compression algorithms: VPW-FRI, Gaussian pulses and downsampling operation.

have access to the original signals, which means better parameter estimation. We compare the VPW-FRI pulse model with a Gaussian model, as is for example described in [21]. Each Gaussian pulse is expressed with 3 parameters: location, scale and amplitude. We fit them to the original signal using nonlinear optimization techniques. Additionally, we compare the results with ideal lowpass filtering followed by downsampling. Fig. 15 shows that VPW-FRI is superior in terms of robustness and precision and Fig. 16 shows an example of a compressed reconstruction. On the single record of Fig. 16, VPW-FRI achieved an SRR of 22.3 dB compared to 14.0 dB for the Gaussian pulse model. We also observe that the Gaussian model sometimes fails to identify narrow pulses. Another disadvantage of this technique is that it is difficult to initialize the nonlinear optimization. In fact, in the results presented here, we initialized it with the VPW-FRI estimated values, since this led to faster convergence and a significant improvement in the estimation.

Finally, we compare our results to [44], where a wavelet-based compressed sensing algorithm is applied to a selection of four records from the MIT-BIH arrhythmia database. For example, to achieve a compression ratio of 10 on the first 10 minutes on record 107, an SRR marginally below 22 dB is reported in [44], while our approach achieves an SRR of 25.4 dB.

VII. APPLICATION II: BRDF ACQUISITION

We demonstrate in this section how the VPW pulse model can be exploited to acquire and estimate the bidirectional
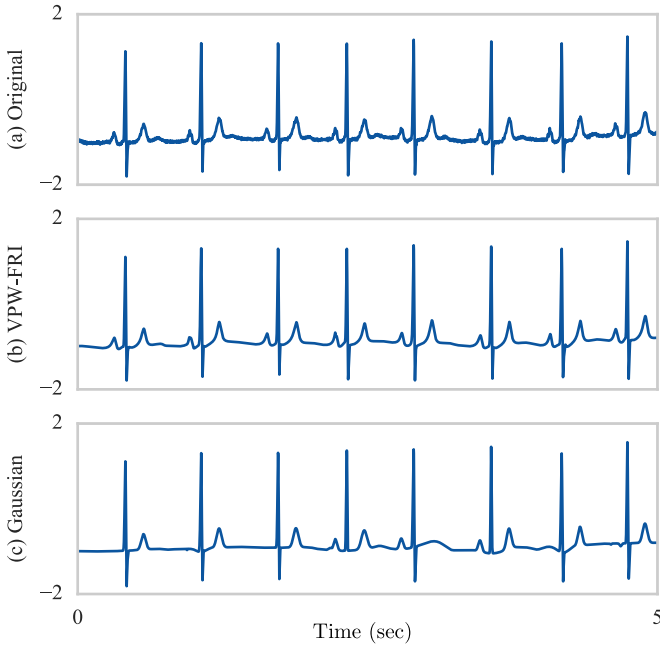
Fig. 16. ECG compression on the first 10 seconds of record 123 from the MIT-BIH database: (a) original ECG signal; (b) VPW estimation with 5 pulses for each heart beat (16 parameters/sec); (c) estimation with 7 Gaussian functions (16.8 parameters/sec).

Fig. 17. Vectors involved in the BRDF function: $\boldsymbol{\omega}_i$ is the incoming light direction, $\boldsymbol{\omega}_o$ is the outgoing light direction, and $\boldsymbol{n}$ is the surface normal. The vectors ℓ_x and ℓ_y represent the projection of $\boldsymbol{\omega}_i$ onto the xy plane.

reflectance distribution function (BRDF). In particular, we focus on the specularity, which is a narrow peak that can be challenging to sample. We start with a brief review of the BRDF definitions and existing models and then describe three experiments to justify the use of VPW-FRI; each of them focuses on a different domain (1D, 2D and spherical) and makes use of several data sources: synthetic data, real data from a publicly-available dataset, as well our own acquired data using a custom-made gantry.

A. Background

The illumination has a strong influence on the appearance of objects: for example, a shiny material might look brighter if light is directed at it from a particular angle and darker otherwise. The role of the BRDF [45] is precisely to describe the reflective characteristics of materials; it enables us to decouple the reflectance properties of objects from the illumination. The full knowledge of the BRDF also allows the rendering of realistic scenes under new illuminations and views.

More formally, the BRDF $\rho(\boldsymbol{w}_i, \boldsymbol{w}_o)$ is a four-dimensional function that measures the amount of light that is reflected in the viewing direction $\boldsymbol{w}_o = (\theta_o, \phi_o)$ given a light source with incoming direction $\boldsymbol{w}_i = (\theta_i, \phi_i)$ (see Fig. 17). The BRDF is specific to the wavelength, and in practice we often use a separate reflectance function for each spectral band; for the sake of clarity, we focus here on a single channel.

Acquisition of BRDFs is performed using a *goniophotometer* — a device that takes several photographs of a material under different illuminations; for instance in [25], almost 3 million BRDF samples have been acquired for 100 different materials that now compose the publicly available MERL BRDF database.
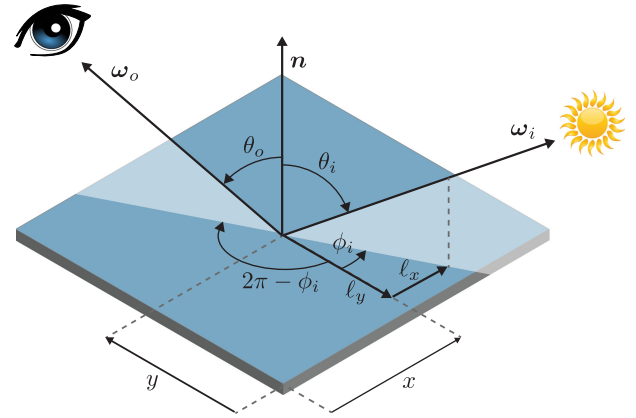
The high dimensionality of BRDFs often results in huge datasets, and it motivates the need for lower-dimensional approximations and reparametrizations. To this end, it is common to represent the incoming and outgoing light directions in the half-angle coordinate system [46]. With this parametrization, the BRDF is expresed in terms of the half-angle $\boldsymbol{h}$ and the difference $\boldsymbol{d}$: $\boldsymbol{h} = (\theta_h, \phi_h) = \frac{\boldsymbol{\omega}_i + \boldsymbol{\omega}_o}{\|\boldsymbol{\omega}_i + \boldsymbol{\omega}_o\|}$ is the angle that is half way between $\boldsymbol{\omega}_i$ and $\boldsymbol{\omega}_o$, while $\boldsymbol{d} = (\theta_d, \phi_d)$ is the angle between $\boldsymbol{h}$ and $\boldsymbol{\omega}_i$. The advantage of this representation is that it is independent of ϕ_d for *isotropic materials*. This is the approach taken in [25] to reduce the dataset size to three dimensions. Other simplifications consist of fixing the viewing angle or projecting the vector $\boldsymbol{\omega}_i$ onto the xy plane, as in [47]: in Fig. 17, this projection is illustrated by the two components ℓ_x and ℓ_y. For each of these representations, we demonstrate usages of the VPW model.

B. Modeling Glossy Materials

First, we show that a VPW pulse is adequate to model glossy materials. In its elementary form, a BRDF is decomposed into two parts: the *diffuse* (or *Lambertian*) part – a constant term that represents the amount of light that is equally reflected in all directions — and the *specular* component — a peak whose direction has the same angle as the incoming light with respect to the surface normal. In mathematical terms,

$$\rho(\boldsymbol{w}_i, \boldsymbol{w}_o) = \underbrace{\frac{k_d}{\pi}}_{\text{diffuse}} + \underbrace{k_s \rho_s(\boldsymbol{w}_i, \boldsymbol{w}_o)}_{\text{specular}}, \qquad (37)$$

where k_d is the diffuse albedo, k_s the specular albedo, and ρ_s is a function of a few arguments that represents the shape of the specular peak. A wide range of parametric models have been proposed to represent ρ_s [25]; one of the most flexible and accurate [48] being the Cook-Torrance model [49], also known as the *microfacet model*. It is given by

$$\rho_s(\boldsymbol{w}_i, \boldsymbol{w}_o) = \frac{F(\boldsymbol{\omega}_o, \boldsymbol{h}) G(\boldsymbol{\omega}_i, \boldsymbol{\omega}_o, \boldsymbol{h}) D(\boldsymbol{h})}{4 \langle \boldsymbol{\omega}_i, \boldsymbol{n} \rangle \langle \boldsymbol{\omega}_o, \boldsymbol{n} \rangle}, \qquad (38)$$

where the function F is called the *Fresnel factor*, G is the *geometry term*, and D is the *microfacet distribution*. While F
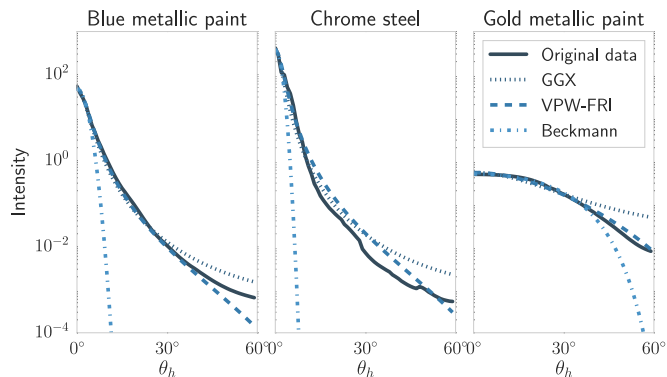
Fig. 18. Modeling of a few glossy materials from the MERL BRDF database using a GGX distribution, Beckmann distribution and a VPW pulse.

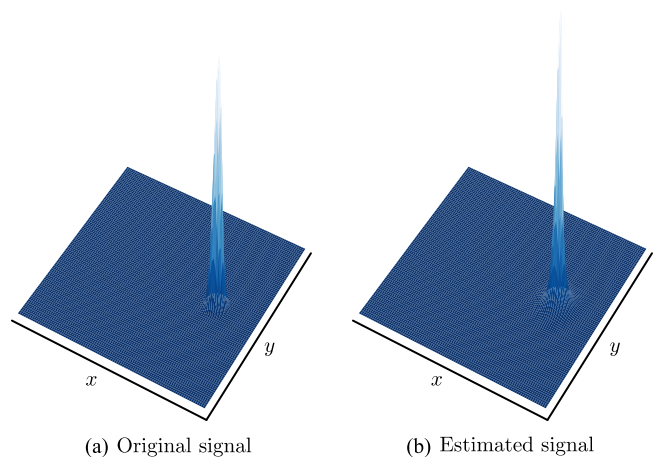(a) Original signal   (b) Estimated signal

Fig. 19. 2D fit of a specularity: (a) Beckmann distribution and (b) VPW pulse. For this experiment, we used all DFT coefficients corresponding to positive frequencies — in this case 50 in each dimension — in the estimation algorithm.

and G are fairly constant throughout the range of values of θ_h (except near grazing angle, where they exhibit a particular behavior), D is the factor we are most interested in since it describes the global shape of the specularity. Again, several shapes have been proposed for D; two of the most widely used are the Beckmann [50] and the GGX [51] distributions:

$$D_{Beckmann}(\boldsymbol{h}) = \frac{1}{\pi\alpha^2\langle\boldsymbol{n},\boldsymbol{h}\rangle^4}\exp\left(-\frac{\tan(\theta_h)^2}{\alpha^2}\right),$$

$$D_{GGX}(\boldsymbol{h}) = \frac{\alpha^2}{\pi\langle\boldsymbol{n},\boldsymbol{h}\rangle^4(\alpha^2+\tan(\theta_h)^2)^2},$$

where θ_h is the angle between $\boldsymbol{h}$ and $\boldsymbol{n}$, and α is a *roughness parameter*, analogous to r_0 in the VPW model. Observe that the microfacet distribution is 1-dimensional in these models: it depends only on $\tan(\theta_h)$.

We show that a single VPW pulse is appropriate to model the specularity ρ_s; indeed, we observe that if we formulate it with respect to $\tan(\theta_h)$, its shape is similar to the GGX distribution, up to a normalization factor and a square root term.[9] We compare in Fig. 18 the fitting of these models to a few measured BRDFs from the MERL BRDF database [25]. In this experiment, and as we neglect the Fresnel and geometry terms, we focus on θ_h in the range $[0°, 60°]$. Within this range, we see that the VPW model compares favorably to the Beckmann and GGX distributions and presents a good trade-off between the heavy tails of the GGX distribution and the relatively narrow Beckmann distribution.

C. Specularity Localization on the Plane

Next, we demonstrate on synthetic data how our algorithm and model can be used to estimate the width and location of a specular peak projected onto the xy plane. For this experiment, we fix the viewing angle and only study the variation with respect to the incoming light angle. The pulse in Fig. 19(a) is a specularity generated according to a Beckmann distribution with $\alpha = 0.003$ and $k_s = 0.005$; we fit a VPW pulse to it and estimate its parameters using the FRI procedure coupled with a 2D version of Cadzow's algorithm. As we can see, a single

Fig. 20. (a) Photograph of our light dome for BRDF acquisition; (b) closeup of an artwork from which we captured the BRDF.

VPW pulse — see Fig. 19(b) — provides a good approximation of its shape.

D. Specularity Localization on the Sphere

Finally, we propose to perform a practical acquisition with a custom-made gantry: a light dome that consists of an extensible motorized arm with a light source at its extremity [52]. The arm can be rotated around the center, where a camera is installed (see Fig. 20). This enables us to position the incoming light at almost any location on the hemisphere and densely sample the two dimensions of the BRDF representing the incoming light (the viewing angle is fixed). Depending on the number of angular samples, the acquisition process can be lengthy and tedious. Moreover, most light domes are composed of a limited number of lights at fixed positions; this motivates the need for a method that can faithfully estimate the BRDF from as few samples as possible.

Using a Nikon D810 full-frame DSLR camera, we acquired samples of the reflectance field of the artwork shown in Fig. 20 using $N = 18$ different light positions; we discretized the azimuth into 6 uniform bins and the elevation into 3 uniform bins.[10] In practice the light source is not perfectly pointwise; this fact, combined with the acquisition system, acts as a low-pass filter that eliminates the high frequency components. Using

---

[9]To account for this difference, we actually fit the *square root* of the measurements to the VPW pulse.

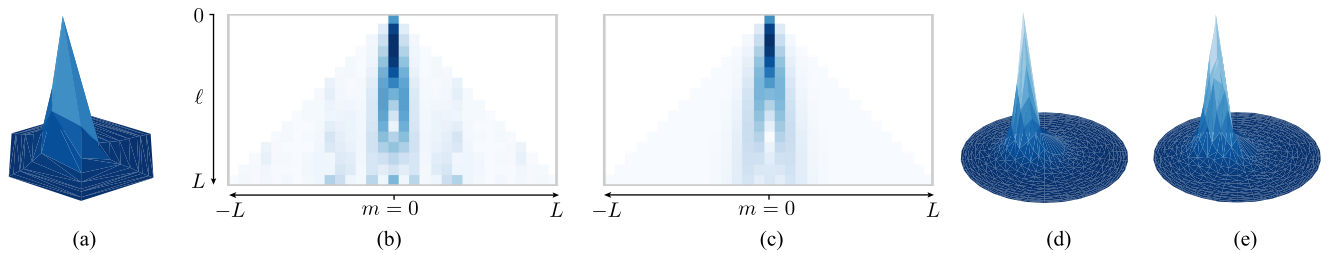[10]Our dataset is available at http://rr.epfl.ch/paper/BSB2016

Fig. 21. Specularity estimation using spherical VPW-FRI algorithm: (a) 18 samples of the acquired BRDF; (b) the spectrum computed from the 18 samples; (c) the spectrum estimated using VPW-FRI with $K = 1$; (d) the corresponding reconstructed spherical signal; (e) the same BRDF acquired at a higher sampling rate (450 samples) for comparison.

Theorem 2, we can approximate the SH coefficients from the spatial samples.

We then use spherical VPW-FRI to estimate the parameters of the pulses from these coefficients; for most BRDFs, one or two pulses suffice to represent the specularity. As a consequence, we can obtain a good approximation of the specular peak even with a relatively low number of samples on the sphere. This idea is illustrated in Fig. 21: from only 18 spherical samples, we can approximate the spectrum of the reflectance field — see Fig. 21(b) — using Eq. (27). Moreover, from these SH coefficients, the FRI algorithm allows us to estimate the spectrum, as illustrated in Fig. 21(c). Since VPW-FRI is a parametric model, its spectrum can have infinite support. This enables a continuous representation of the reflectance function in the spherical domain; Fig. 21(d) shows the estimated VPW pulse sampled at 450 locations on the hemisphere. For the sake of comparison, we also acquired a denser version of the reflectance field — shown in Fig. 21(e); we observe that the VPW-FRI approximation from 18 samples is extremely close to the ground truth.

To summarize, the algorithm allows us to accurately estimate the specularity from a few samples; this reduces the acquisition time as well as storage requirements. Furthermore, given this continuous VPW representation, we can render objects under virtually any new illumination.

E. Practical Considerations

We already mentioned in Section V that the method does not exploit the entire spectrum but only uses its two outer diagonals, i.e. the SH coefficients for $m = \ell$ and $m = \ell - 1$. Furthermore, as depicted in Fig. 21(b), the energy of a typical BRDF signal is concentrated along the central column of the spectrum, and these coefficients are unfortunately not used by the currently proposed algorithm. To achieve better spectrum usage, we add a nonlinear optimization routine that minimizes the squared error between the model and the measurements over the full spectrum. As the objective function is non-convex, this method relies on the VPW-FRI parameters to provide a good initialization. Moreover, to make the VPW-FRI estimation more robust, we first denoise the SH coefficients with Cadzow's algorithm.

Another limitation is model mismatch. While our experiments have demonstrated that VPW pulses are adequate to represent specular peaks, there are other components in the BRDF that are
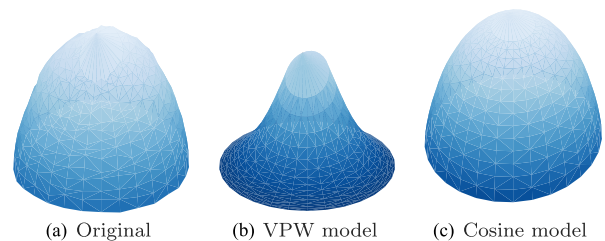

Fig. 22. Example of model mismatch: (a) a measured BRDF composed only of a Lambertian part, and (b) its spherical VPW approximation. Simpler models such as (c) a cosine can be used instead.

not easily approximated by VPW pulses: the Lambertian component and hard cast shadows are such examples. The mismatch between the Lambertian component and its VPW approximation is illustrated in Fig 22. Fortunately, the Lambertian part is much less challenging to detect and represent than the specular part, since it can easily be approximated using a constant value modulated with the cosine of the azimuth of the incoming light angle θ_i. Hard shadows are more difficult to cope with, as they also contain high frequency components, and thus provide an interesting topic of investigation for future work.

VIII. CONCLUSION

We have presented VPW pulses, a new model that fits within the FRI framework and allows for a wider range of pulse shapes. We also proposed a method for estimating the parameters and quantified its performance on noisy signals.

Moreover, we showed that VPW can be extended to 2D and spherical signals. In 2 dimensions, we proposed several pulse shapes and discussed their characteristics. Unlike in 1D, the spectrum usage is not optimal in both 2D and on the sphere; we need to sample at a rate higher than the rate of innovation to ensure perfect recovery. Another limitation is that the spherical case allows only radial pulses.

We demonstrated that a small number of VPW pulses are adequate to acquire, represent and compress ECG signals. Finally, we applied our signal model to the estimation of the BRDF, with emphasis on the detection of specularities; we showed that with the VPW-FRI framework we can accurately retrieve and represent the shape of specularities from just a handful of measurements. In both applications, we stressed the importance of the data matching the underlying parametric model: when this condition was satisfied, results were very convincing.

APPENDIX

A. Proof of Lemma 1

Proof: Since $f(t)$ is periodic, with period τ, it has the following Fourier series expansion:

$$f(t) = \sum_{m \in \mathbb{Z}} F[m] e^{2\pi i m t / \tau}.$$

From the sampling scheme in Fig. 4,

$$
\begin{aligned}
g[n] &= \langle f(t), \varphi(t - nT) \rangle \\
&= \sum_{m \in \mathbb{Z}} F[m] \left\langle e^{2\pi i m t / \tau}, B \operatorname{sinc}(B(t - nT)) \right\rangle \\
&= \sum_{m \in \mathbb{Z}} F[m] e^{2\pi i m n T / \tau} \int_{-\infty}^{\infty} \operatorname{sinc}(t) e^{2\pi i m t / (B \tau)} dt \\
&= \sum_{m = -\lfloor B\tau/2 \rfloor}^{\lfloor B\tau/2 \rfloor} F[m] e^{2\pi i m n / N}.
\end{aligned}
$$

This is a system of N equations in $2 \lfloor B\tau/2 \rfloor + 1$ unknown Fourier series coefficients $F[m]$, which can be inverted when $N \geq 2 \lfloor B\tau/2 \rfloor + 1$. When $N = 2 \lfloor B\tau/2 \rfloor + 1$, the system is, up to a $1/N$ factor, the $N \times N$ inverse-DFT of $F[m]$. Therefore,

$$F[m] = N G[m],$$

for $|m| \leq \lfloor B\tau/2 \rfloor$. ∎

B. Time-domain Formulae for VPW Pulse

We show here the details of the derivation of the time-domain formulae for the VPW pulse. The time-domain representation of F_k^s — cf. Eq. (11) — is computed as

$$
\begin{aligned}
f_k^s(t) &= \sum_{m=-\infty}^{\infty} F_k^s[m] e^{i 2\pi m t / \tau} \\
&= \frac{1}{\tau} \left(\sum_{m=0}^{\infty} z_k^m(t) + \sum_{m=0}^{\infty} (z_k^*(t))^m - 1 \right) \\
&\stackrel{(a)}{=} \frac{1}{\tau} \left(\frac{1}{1 - z_k(t)} + \frac{1}{1 - z_k^*(t)} - 1 \right) \\
&= \frac{1}{\tau} \frac{1 - |z_k(t)|^2}{(1 - z_k(t))(1 - z_k^*(t))},
\end{aligned}
$$

where $z_k(t) = e^{2\pi/\tau(-r_k + i(t - t_k))}$. Note that (a) assumes that $r_k > 0$, which only holds when the annihilating filter is *stable* (i.e., when its roots are inside the unit circle).

We can do the same reasoning for the asymmetric component F_k^a of the signal, whose time-domain representation is

$$
\begin{aligned}
f_k^a(t) &= \sum_{m=-\infty}^{\infty} F_k^a[m] e^{i 2\pi m t / \tau} \\
&= \frac{i}{\tau} \left(\sum_{m=0}^{\infty} z_k^m(t) - \sum_{m=0}^{\infty} (z_k^*(t))^m + 1 \right)
\end{aligned}
$$

$$
\begin{aligned}
&\stackrel{(d)}{=} \frac{i}{\tau} \left(\frac{1}{1 - z_k(t)} - \frac{1}{1 - z_k^*(t)} \right) \\
&= \frac{1}{\tau} \frac{2 \operatorname{Im}\{z_k(t)\}}{(1 - z_k(t))(1 - z_k^*(t))}.
\end{aligned}
$$

Again, (d) only holds for $r_k > 0$.

C. Computation of the Cramer-Rao Bound for VPW Signals

Let $\boldsymbol{\Theta} = [t_0, r_0, c_0]^T$ be the vector of unknown parameters. The CRB relies on the Fisher information matrix, which is given by (cf. [53] for more details)

$$
\begin{aligned}
\boldsymbol{I}(\boldsymbol{\theta}) &= \mathrm{E} \left[\left(\frac{\partial \log \tilde{f}[n])}{\partial \boldsymbol{\theta}} \right) \left(\frac{\partial \log \tilde{f}[n])}{\partial \boldsymbol{\theta}} \right)^T \right] \\
&= \frac{1}{\sigma_\epsilon^2} \sum_{n=0}^{N-1} \nabla f[n] \nabla f[n]^T.
\end{aligned}
$$

Note that $f[n]$ and $\tilde{f}[n]$ also depend on $\boldsymbol{\theta}$, but we have omitted it to simplify the notation.

Given any unbiased estimate $\hat{\boldsymbol{\theta}}$ of the parameters, the CRB provides a lower bound on the covariance matrix of $\hat{\boldsymbol{\theta}}$:

$$\operatorname{cov}(\hat{\boldsymbol{\theta}}) \succeq \boldsymbol{I}^{-1}(\boldsymbol{\theta}).$$

Let our signal $\boldsymbol{f} = [f[0], f[1], \ldots, f[N-1]]^T$ and its DFT $\hat{\boldsymbol{f}} = [F[0], F[1], \ldots, F[N-1]]^T$. Since $\boldsymbol{f} = \boldsymbol{W}^\times \hat{\boldsymbol{f}}$ — where $\boldsymbol{W}$ is the DFT matrix of size N — and we have a closed-form formula for $F[m]$, we choose to first compute the derivatives for the gradient in the frequency domain and then perform the inverse DFT to get $\nabla f[n]$. If we restrict ourselves to positive m and a single VPW pulse, we have, given $F[m] = c_0 e^{-2\pi m (r_0 + i t_0)/\tau}$,

$$
\begin{aligned}
\frac{\partial F[m]}{\partial t_0} &= -c_0 \frac{2\pi i m}{\tau} e^{-2\pi m (r_0 + i t_0)/\tau} \\
\frac{\partial F[m]}{\partial r_0} &= -c_0 \frac{2\pi m}{\tau} e^{-2\pi m (r_0 + i t_0)/\tau} \\
\frac{\partial F[m]}{\partial c_0} &= e^{-2\pi m (r_0 + i t_0)/\tau}.
\end{aligned}
$$

ACKNOWLEDGMENT

REFERENCES

[1] C. E. Shannon, "A mathematical theory of communication," *Bell Syst. Tech. J.*, vol. 27, pp. 379–423, 1948.

[2] M. Unser, "Sampling - 50 years after shannon," *Proc. IEEE*, vol. 88, no. 4, pp. 569–587, 2000.

[3] M. Vetterli, P. Marziliano, and T. Blu, "Sampling signals with finite rate of innovation," *IEEE Trans. Signal Process.*, vol. 17, no. 6, pp. 1417–1428, 2002.

[4] J. Berent, P. L. Dragotti, and T. Blu, "Sampling piecewise sinusoidal signals with finite rate of innovation methods," *IEEE Trans. Signal Process.*, vol. 58, no. 2, pp. 613–625, 2010.

[5] P. L. Dragotti, M. Vetterli, and T. Blu, "Sampling moments and reconstructing signals of finite rate of innovation: Shannon meets Strang-Fix," *IEEE Trans. Signal Process.*, vol. 55, no. 5, pp. 1741–1757, 2007.

[6] R. Tur, Y. C. Eldar, and Z. Friedman, "Innovation rate sampling of pulse streams with application to ultrasound imaging," *IEEE Trans. Signal Process.*, vol. 59, no. 4, pp. 1827–1842, 2011.

[7] J. Oñativia, S. R. Schultz, and P. L. Dragotti, "A finite rate of innovation algorithm for fast and accurate spike detection from two-photon calcium imaging," *J. Neural Eng.*, vol. 10, no. 4, pp. 1–30, 2013.

[8] G. Beylkin and L. Monzón, "Nonlinear inversion of a band-limited Fourier transform," *Appl. Comput. Harmon. Anal.*, vol. 27, no. 3, pp. 351–366, 2009.

[9] F. Andersson, M. Carlsson, and M. V. de Hoop, "Sparse approximation of functions using sums of exponentials and AAK theory," *J. Approx. Theory*, vol. 163, no. 2, pp. 213–248, 2011.

[10] F. Andersson, M. Carlsson, and M. V. de Hoop, "Frequency extrapolation through sparse sums of Lorentzians," *J. Earth Sci.*, vol. 25, no. 1, pp. 117–125, 2014.

[11] V. Chandrasekaran, B. Recht, P. A. Parrilo, and A. S. Willsky, "The convex geometry of linear inverse problems," *Found. Comput. Math.*, vol. 12, no. 6, pp. 805–849, 2012.

[12] B. N. Bhaskar, G. Tang, and B. Recht, "Atomic norm denoising with applications to line spectral estimation," *IEEE Trans. Signal Process.*, vol. 61, no. 23, pp. 5987–5999, 2013.

[13] G. Tang, B. N. Bhaskar, P. Shah, and B. Recht, "Compressed sensing off the grid," *IEEE Trans. Inf. Theory*, vol. 59, no. 11, pp. 7465–7490, Nov. 2013.

[14] E. J. Candès and C. Fernandez-Granda, "Towards a mathematical theory of super-resolution," *Commun. Pure Appl. Math.*, vol. 67, no. 6, pp. 906–956, 2014.

[15] R. F. Quick, R. E. Crochiere, J. H. Hong, A. Hormati, and G. Baechler, "Extension of FRI for modeling of electrocardiogram signals," in *Proc. Annu. Int. Conf. IEEE Eng. Med. Biol. Soc*, 2012, pp. 2909–2912.

[16] G. Baechler, "Sensing ECG signals with variable pulse width finite rate of innovation," Master's thesis, EPFL, 2012.

[17] A. Nair, P. Marziliano, R. F. Quick, R. E. Crochiere, and G. Baechler, "Multichannel ECG analysis using VPW-FRI," *Proc. 10th Int. Conf. Sampling Theory Appl.*, no. 2, pp. 125–128, 2014.

[18] A. Nair and P. Marziliano, "Noisy channel detection using the common annihilator with an application to electrocardiograms," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process.*, 2015, pp. 5972–5976.

[19] A. L. Goldberger *et al.*, "PhysioBank, physio toolkit, and physio net: Components of a new research resource for complex physiologic signals," *Circulation*, vol. 101, no. 23, pp. e215–e220, 2000.

[20] A. Goldberger, *Clinical Electrocardiography: A Simplified Approach*. Philadelphia, PA, USA: Elsevier Health Sciences, 2012.

[21] M. S. Billah, T. B. Mahmud, F. S. Snigdha, and M. A. Arafat, "A novel method to model ECG beats using Gaussian functions," in *Proc. IEEE 4th Int. Conf. Biomed. Eng. Inf.*, vol. 2, 2011, pp. 612–616.

[22] A. Sandryhaila, S. Saba, M. Püschel, and J. Kovacević, "Efficient compression of QRS complexes using Hermite expansion," *IEEE Trans. Signal Process.*, vol. 60, no. 2, pp. 947–955, Feb. 2012.

[23] M. Lagerholm, C. Peterson, G. Braccini, L. Edenbrandt, and L. Sornmo, "Clustering ECG complexes using Hermite functions and self-organizing maps," *IEEE Trans. Biomed. Eng.*, vol. 47, no. 7, pp. 838–848, 2000.

[24] A. Ouamri and A. Naït-Ali, "ECG compression method using Lorentzian functions model," *Digit. Signal Process.*, vol. 17, no. 1, pp. 319–326, 2007.

[25] W. Matusik, H. Pfister, M. Brand, and L. McMillan, "A data-driven reflectance model," *ACM Trans. Graph.*, vol. 22, p. 759, 2003.

[26] B. Tunwattanapong, G. Fyffe, P. Graham, J. Busch, X. Yu, A. Ghosh, and P. Debevec, "Acquiring reflectance and shape from continuous spherical harmonic illumination," *ACM Trans. Graph.*, vol. 32, no. 4, Jul. 2013.

[27] M. Unser and T. Blu, "Cardinal exponential splines: Part I - Theory and filtering algorithms," *IEEE Trans. Signal Process.*, vol. 53, no. 4, pp. 1425–1438, 2005.

[28] T. Blu and M. Unser, "Quantitative fourier analysis of approximation techniques. i. Interpolators and projectors," *IEEE Trans. Signal Process.*, vol. 47, no. 10, pp. 2783–2795, Oct 1999.

[29] J. A. Uriguën, T. Blu, and P. L. Dragotti, "FRI sampling with arbitrary kernels," *IEEE Trans. Signal Process.*, vol. 61, no. 21, pp. 5310–5323, 2013.

[30] P. Stoica and R. Moses, *Spectral Analysis of Signals*. Upper Saddle River, NJ, USA: Pearson Prentice-Hall, 2005.

[31] J. A. Cadzow, "Signal enhancement—a composite property mapping algorithm," *IEEE Trans. Acoust., Speech Signal Process.*, vol. 36, no. I, 1988.

[32] Y. Hua and T. K. Sarkar, "Matrix pencil method for estimating parameters of exponentially damped/undamped sinusoids in noise," *IEEE Trans. Acoust., Speech Signal Process.*, vol. 38, no. 5, pp. 814–824, 1990.

[33] R. Roy and T. Kailath, "ESPRIT - estimation of signal parameters via rotational invariance techniques," *IEEE Trans. Acoust., Speech Signal Process.*, vol. 37, no. 7, pp. 984–995, 1989.

[34] V. F. Pisarenko, "The retrieval of harmonies from a covariance function," *Geophys. J. R. Astrophys. Soc.*, vol. 33, pp. 347–366, 1973.

[35] Y. Bresler and A. Macovski, "Exact maximum likelihood parameter estimation of superimposed exponential signals in noise," *IEEE Trans. Acoust., Speech Signal Process.*, vol. 34, no. 5, pp. 1081–1089, 1986.

[36] H. Pan, T. Blu, and M. Vetterli, "Towards generalized FRI sampling with an application to source resolution in radioastronomy," *IEEE Trans. Signal Process.*, vol. 65, no. 4, pp. 821–835, 2017.

[37] I. Maravić and M. Vetterli, "Exact sampling results for some classes of parametric nonbandlimited 2-D signals," *IEEE Trans. Signal Process.*, vol. 52, no. 1, pp. 175–189, 2004.

[38] S. Deslauriers-Gauthier and P. Marziliano, "Sampling signals with a finite rate of innovation on the sphere," *IEEE Trans. Signal Process.*, vol. 61, no. 18, pp. 4552–4561, 2013.

[39] I. Dokmanić and Y. Lu, "Sampling sparse signals on the sphere: Algorithms and applications," *IEEE Trans. Signal Process.*, vol. 64, no. 1, pp. 189–202, 2016.

[40] B. Rafaely, *Fundamentals of Spherical Array Processing*. Berlin, Germany: Springer-Verlag, 2015.

[41] J. R. Driscoll and D. M. Healy, "Computing Fourier transforms and convolutions on the 2-sphere," *Adv. Appl. Math.*, vol. 15, no. 2, pp. 202–250, 1994.

[42] R. F. Bass and K. Gröchenig, "Random sampling of multivariate trigonometric polynomials," *SIAM J. Math. Anal.*, vol. 36, no. 3, pp. 773–795, 2005.

[43] L. Sörnmo and P. Laguna, *Bioelectrical Signal Processing in Cardiac and Neurological Applications*. Elsevier Academic Press, 2005.

[44] L. F. Polania, R. E. Carrillo, M. Blanco-Velasco, and K. E. Barner, "Compressed sensing based method for ECG compression," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process.*, May 2011, pp. 761–764.

[45] F. E. Nicodemus, "Directional reflectance and emissivity of an opaque surface," *Appl. Opt.*, vol. 4, no. 7, pp. 767–775, Jul. 1965.

[46] S. M. Rusinkiewicz, "A new change of variables for efficient BRDF representation," *Proc. Eurographics Rendering Workshop Rendering Techniques*, pp. 11–22, 1998.

[47] T. Malzbender, D. Gelb, and H. Wolters, "Polynomial texture maps," *Proc. SIGGRAPH*, vol. 33, pp. 519–528, 2001.

[48] A. Ngan, F. Durand, and W. Matusik, "Experimental analysis of BRDF models," in *Proc. Eurographics Symp. Rendering*, pp. 117–126, 2005.

[49] R. L. Cook and K. E. Torrance, "A reflectance model for computer graphics," *ACM SIGGRAPH Comput. Graph.*, vol. 15, no. 3, pp. 307–316, 1981.

[50] G. J. Ward, "Measuring and modeling anisotropic reflection," *ACM SIGGRAPH Comput. Graph.*, vol. 26, no. Jul., pp. 265–272, 1992.

[51] B. Walter, S. Marschner, H. Li, and K. Torrance, "Microfacet models for refraction through rough surfaces," *Eurographics*, pp. 195–206, 2007.

[52] L. A. Baboulaz, G. G. Soudan, and M. Vetterli, "Device for determining a bidirectional reflectance distribution function of a subject," U.S. Patent US 14 559 726, 2014.

[53] P. L. Dragotti and F. Homann, "Sampling signals with finite rate of innovation in the presence of noise," in *IEEE Int. Conf. Acoust., Speech Signal Process. (ICASSP)*, 2009, pp. 2941–2944.

Gilles Baechler (S'12), photograph and biography not available at the time of publication.

Adam Scholefield (M'12), photograph and biography not available at the time of publication.

Loïc Baboulaz, photograph and biography not available at the time of publication.

Martin Vetterli (S'86–M'86–SM'90–F'95), photograph and biography not available at the time of publication.