

Performance Bounds and Design Criteria for Estimating Finite Rate of Innovation Signals

Zvika Ben-Haim, *Student Member, IEEE*, Tomer Michaeli, *Student Member, IEEE*,
and Yonina C. Eldar, *Senior Member, IEEE*

Abstract—In this paper, we consider the problem of estimating finite rate of innovation (FRI) signals from noisy measurements, and specifically analyze the interaction between FRI techniques and the underlying sampling methods. We first obtain a fundamental limit on the estimation accuracy attainable regardless of the sampling method. Next, we provide a bound on the performance achievable using any specific sampling approach. Essential differences between the noisy and noise-free cases arise from this analysis. In particular, we identify settings in which noise-free recovery techniques deteriorate substantially under slight noise levels, thus quantifying the numerical instability inherent in such methods. This instability, which is only present in some families of FRI signals, is shown to be related to a specific type of structure, which can be characterized by viewing the signal model as a union of subspaces. Finally, we develop a methodology for choosing the optimal sampling kernels based on a generalization of the Karhunen–Loève transform. The results are illustrated for several types of time-delay estimation problems.

Index Terms—Finite rate of innovation, Sampling, Cramér–Rao bound, Union of subspaces, Time-delay estimation

I. INTRODUCTION

The field of digital signal processing hinges on the availability of techniques for sampling analog signals, thus converting them to discrete measurements. The sampling mechanism aims to preserve the information present in the analog domain, ideally permitting flawless recovery of the original signal. For example, one may wish to recover a continuous-time signal $x(t)$ from a discrete set of samples. The archetypical manifestation of this concept is the Shannon sampling theorem, which states that a B -bandlimited function can be reconstructed from samples taken at the Nyquist rate $2B$ [1].

Recently, considerable attention has been devoted to the extension of sampling theory to functions having a finite rate of innovation (FRI). These are signals determined by a finite number ρ of parameters per time unit [2]. Such a definition encompasses a rich variety of signals, including splines, shift-invariant signals, multiband signals, and pulse streams. In many FRI settings, several existing algorithms are guaranteed to recover the signal $x(t)$ from samples taken at rate ρ [2]–[7]. In other words, signals which correspond to the FRI model can

be reconstructed from samples taken at the rate of innovation, which is potentially much lower than their Nyquist rate.

The set of signals described by an FRI model can often be viewed as a union of subspaces [5]–[8]. For example, consider a stream of pulses parameterized by pulse locations and amplitudes. The set of all pulses having a given location is a subspace of the space of continuous-time functions. Thus, the set of all pulses having arbitrary locations is a union of such subspaces. As we will see, this point of view yields a flexible and productive framework for understanding the types of constraints implied by the model.

Real-world signals are often contaminated by continuous-time noise and thus do not conform precisely to the FRI model. Furthermore, like any mathematical model, the FRI framework is an approximation which does not precisely hold in practical scenarios, an effect known as mismodeling error [8]. It is therefore of interest to quantify the effect of noise and mismodeling errors on FRI techniques [3], [6], [7], [9]. In the noisy case, it is no longer possible to perfectly recover the original signal from its samples. Nevertheless, one might hope for an appropriate finite-rate technique which achieves the best possible estimation accuracy, in the sense that increasing the sampling rate confers no further performance benefits. For example, to recover a B -bandlimited signal contaminated by continuous-time white noise, one can use an ideal low-pass filter with cutoff B prior to sampling at a rate of $2B$. This strategy removes all noise components with frequencies larger than B , while leaving all signal components intact. Consequently, any alternative method which does not zero out frequencies above B can be improved upon, whereas methods which zero out some of the signal frequencies can suffer from an arbitrarily large reconstruction error. Thus, sampling at a rate of $2B$ is indeed optimal in the case of a B -bandlimited signal, if the signal is corrupted by continuous-time noise prior to sampling. Sampling at a rate higher than $2B$ can be beneficial only when the sampling process itself introduces additional noise into the system, e.g., as a result of quantization.

By contrast, empirical observations indicate that, for some noisy FRI signals, substantial performance improvements are achievable when the sampling rate is increased beyond the rate of innovation [4], [7]. Thus, in some cases, there appears to be a fundamental difference between the noiseless and noise-corrupted settings, in terms of the required sampling rate. Our first goal in this paper will be to provide an analytical justification and quantification of these empirical findings. As we will see, the fact that oversampling improves performance is not

The authors are with the Department of Electrical Engineering, Technion—Israel Institute of Technology, Haifa 32000, Israel (e-mail: {zvika@tx, tomermic@tx, yonina@ee}.technion.ac.il). This work was supported in part by a Magnet grant from the Israel Ministry of Industry and Trade, by the Israel Science Foundation under Grant 1081/07, and by the European Commission’s FP7 Network of Excellence in Wireless COMMunications NEWCOM++ (grant agreement no. 216715).

merely indicative of flaws in existing algorithms; rather, it is a consequence of the inherent difficulty of reconstructing FRI signals under noise. Indeed, we will demonstrate that for some FRI signals, unless considerable oversampling is employed, performance will necessarily deteriorate by several orders of magnitude relative to the optimal achievable reconstruction capability. Such effects occur even when the noise level is exceedingly low. Our analysis will also enable us to identify and characterize the types of signals for which oversampling is necessary.

To demonstrate these results, we first derive the Cramér–Rao bound (CRB) for estimating a finite-duration segment of an FRI signal $x(t)$ directly from continuous-time measurements $y(t) = x(t) + w(t)$, where $w(t)$ is a Gaussian white noise process. This yields a lower bound on the accuracy whereby $x(t)$ can be recovered by any technique, regardless of its sampling rate. This setting is to be distinguished from previous bounds in the FRI literature [9], [10] in three respects. First and most importantly, the measurements are a continuous-time process $y(t)$ and the bound therefore applies regardless of the sampling method. Second, in our model, the noise is added prior to sampling. Thus, as will be shown below, even sampling at an arbitrarily high rate will not completely compensate for the noise. Third, we bound the mean-squared error (MSE) in estimating $x(t)$ and not the parameters defining it, since we seek to determine the accuracy with which $x(t)$ itself can be recovered. Such a bound does not depend on the specific parametrization of the signal, and consequently, possesses a simpler analytical expression.

In practice, rather than processing the continuous-time signal $y(t)$, it is typically desired to estimate $x(t)$ from a discrete set of samples $\{c_n\}$ of $y(t)$. In this scenario, in addition to the continuous-time noise $w(t)$, digital noise may arise from the sampling process itself, for example due to quantization. To quantify the extent to which sampling degrades the ability to recover the signal, we next derive the CRB for estimating $x(t)$ from the measurements $\{c_n\}$. This analysis depends on the relative power of the two noise factors. When only digital noise is present, oversampling can be used to completely overcome its effect. On the other hand, when there exists only continuous-time noise, the bound converges to the continuous-time CRB as the sampling rate increases. In some cases, these bounds coincide at a finite sampling rate, which implies that the sampling scheme has captured all of the information present in the continuous-time signal, and any further increase in the sampling rate is useless. Conversely, when the continuous-time and sampled CRBs differ, the gap between these bounds is indicative of the degree to which information is lost in the sampling process. Our technique can then be used to plot the best possible performance as a function of the sampling rate, and thus provide the practitioner with a tool for evaluating the benefits of oversampling.

When a certain sampling technique achieves the performance of continuous-time measurements, it can be identified using the method described above. However, in some cases no such technique exists, or the sampling rate it requires may be prohibitive. In these cases, it is desirable to determine the optimal sampling scheme having an allowed rate. Since

different signals are likely to perform successfully with different sampling kernels, a Bayesian or average-case analysis is well-suited for this problem. Specifically, we assume that the signal $x(t)$ has a known prior distribution over the class of signals, and determine the linear sampling and reconstruction technique which minimizes the MSE for recovering $x(t)$ from its measurements. While nonlinear reconstruction techniques are commonly used and typically outperform the best linear estimator, this approach provides a simple means for identifying an appropriate sampling method. The resulting method can then be used in conjunction with standard nonlinear FRI algorithms.

We demonstrate our results via the problem of estimating a finite-duration sequence of pulses having unknown positions and amplitudes [2], [4], [5], [7]. In this case, a simple sufficient condition is obtained for the existence of a sampling scheme whose performance bound coincides with the continuous-time CRB. This scheme is based on sampling the Fourier coefficients of the pulse shape, and is reminiscent of recent time-delay estimation algorithms [7]. However, while the sampling scheme is theoretically sufficient for optimal recovery of $x(t)$, we show that in some cases there is room for substantial improvement in the reconstruction stage of these algorithms. Finally, we demonstrate that the Fourier domain is also optimal (in the sense of minimizing the reconstruction MSE) when the sampling budget is limited. Specifically, given an allowed number of samples N , the reconstruction MSE is minimized by sampling the N highest-variance Fourier coefficients of the signal $x(t)$.

The rest of this paper is organized as follows. The problem setting is defined in Section II, and some examples of signals conforming to this model are presented in Section III. We then briefly summarize our main results in Section IV. In Section V, we provide a technical generalization of the CRB to general spaces. This result is used to obtain bounds on the achievable reconstruction error from continuous-time measurements (Section VI) and using a sampling mechanism (Section VII). Next, in Section VIII a Bayesian viewpoint is introduced and utilized to determine the optimal sampling kernels having a given rate budget. The results are demonstrated for the specific signal model of time-delay estimation in Section IX.

II. DEFINITIONS

A. Notation

The following notation is used throughout the paper. A boldface lowercase letter \mathbf{v} denotes a vector, while a boldface uppercase letter \mathbf{M} denotes a matrix. \mathbf{I}_N is the $N \times N$ identity matrix. For a vector \mathbf{v} , the notation $\|\mathbf{v}\|$ indicates the Euclidean norm. Given a complex number $z \in \mathbb{C}$, the symbols z^* and $\Re\{z\}$ denote the complex conjugate and the real part of z , respectively. For an operator P , the range space and null space are $\mathcal{R}(P)$ and $\mathcal{N}(P)$, respectively, while the trace and adjoint are denoted, respectively, by $\text{Tr}(P)$ and P^* . The Kronecker delta, denoted by $\delta_{m,n}$, equals 1 when $m = n$ and 0 otherwise. The expectation of a random variable v is written as $\mathbb{E}\{v\}$.

The Hilbert space of square-integrable complex-valued functions over $[0, T_0]$ is denoted $L_2[0, T_0]$ or simply L_2 . The corresponding inner product is

$$\langle f, g \rangle \triangleq \int_0^{T_0} f(t)g^*(t)dt \quad (1)$$

and the induced norm is $\|f\|_{L_2}^2 \triangleq \langle f, f \rangle$. For an ordered set of K functions g_1, \dots, g_K in L_2 , we define the associated *set transformation* $G : \mathbb{C}^K \rightarrow L_2$ as

$$(G\mathbf{v})(t) = \sum_{k=1}^K \mathbf{v}_k g_k(t). \quad (2)$$

By the definition of the adjoint, it follows that

$$G^*f = (\langle f, g_1 \rangle, \dots, \langle f, g_K \rangle)^T. \quad (3)$$

B. Setting

In this work, we are interested in the problem of estimating FRI signals from noisy measurements. To define FRI signals formally, let the T_0 -local number of degrees of freedom $N_{T_0}(t)$ of a signal $x(t)$ at time t be the number of parameters defining the segment $\{x(t) : t \in [t - T_0/2, t + T_0/2]\}$. The T_0 -local rate of innovation of $x(t)$ is then defined as [2]

$$\rho_{T_0} = \max_{t \in \mathbb{R}} \frac{N_{T_0}(t)}{T_0}. \quad (4)$$

We then say that $x(t)$ is an FRI signal if ρ_{T_0} is finite for all sufficiently large values of T_0 . In Section III, we will give several examples of FRI signals and compute their rates of innovation.

For concreteness, let us focus on the problem of estimating the finite-duration segment $\{x(t) : t \in [0, T_0]\}$, for some constant T_0 , and let $K \triangleq N_{T_0}(T_0/2)$ denote the number of parameters defining this segment. We then have

$$x \in \mathfrak{X} \triangleq \{h_\theta \in L_2[0, T_0] : \theta \in \Theta\} \quad (5)$$

where h_θ is a set of functions parameterized by the vector θ , and Θ is an open subset of \mathbb{R}^K .

We wish to examine the random process

$$y(t) = x(t) + w(t), \quad t \in [0, T_0] \quad (6)$$

where $w(t)$ is continuous-time white Gaussian noise. Recall that formally, it is not possible to define Gaussian white noise over a continuous-time probability space [11]. Instead, we interpret (6) as a simplified notation for the equivalent set of measurements

$$z(t) = \int_0^t x(\tau)d\tau + \sigma_c b(t), \quad t \in [0, T_0] \quad (7)$$

where $b(t)$ is a standard Wiener process (also called Brownian motion) [12]. It follows that $w(t)$ can be considered as a random process such that, for any $f, g \in L_2$, the inner products $a = \langle f, w \rangle$ and $b = \langle g, w \rangle$ are zero-mean jointly Gaussian random variables satisfying $\mathbb{E}\{ab^*\} = \sigma_c^2 \langle f, g \rangle$ [11]. The subscript c in σ_c is meant as a reminder of the fact that $w(t)$ is continuous-time noise. By contrast, when examining samples of the random process $y(t)$, we will also consider digital noise which is added during the sampling process.

In this paper, we consider estimators which are functions either of the entire continuous-time process (6) or of some subset of the information present in (6), such as a discrete set of samples of $y(t)$. To treat these two cases in a unified way, let (Ω, \mathcal{F}) be a measurable space and let $\{P_\theta : \theta \in \Theta\}$ be a family of probability measures over (Ω, \mathcal{F}) . Let $(\mathcal{Y}, \mathcal{U})$ be a measurable space, and let the random variable $y : \Omega \rightarrow \mathcal{Y}$ denote the measurements. This random variable can represent either $y(t)$ itself or samples of this quantity.

An estimator can be defined in this general setting as a measurable function $\hat{x} : \mathcal{Y} \rightarrow L_2$. The MSE of an estimator \hat{x} at x is defined as

$$\text{MSE}(\hat{x}, x) \triangleq \mathbb{E}\{\|\hat{x} - x\|_{L_2}^2\} = \mathbb{E}\left\{\int_0^{T_0} |\hat{x}(t) - x(t)|^2 dt\right\}. \quad (8)$$

An estimator \hat{x} is said to be unbiased if

$$\mathbb{E}\{\hat{x}(t)\} = x(t) \text{ for all } x \in \mathfrak{X} \text{ and almost all } t \in [0, T_0]. \quad (9)$$

In the next section, we demonstrate the applicability of our model by reviewing several scenarios which can be formulated using the FRI framework. Some of these settings will also be used in the sequel to exemplify our theoretical results.

III. TYPES OF FRI SIGNALS

Numerous FRI signal structures have been proposed and analyzed in the sampling literature. Whereas most of these can be treated within our framework, some FRI structures do not conform exactly to our problem setting. Thus, before delving into the derivation of the CRB, we first provide examples for scenarios that can be analyzed via our model and discuss some of its limitations.

A. Shift-Invariant Spaces

Consider the class of signals that can be expressed as

$$x(t) = \sum_{m \in \mathbb{Z}} a[m]g(t - mT) \quad (10)$$

with some arbitrary square-integrable sequence $\{a[m]\}_{m \in \mathbb{Z}}$, where $g(t)$ is a given pulse in $L_2(\mathbb{R})$ and $T > 0$ is a given scalar. This set of signals is a linear subspace of $L_2(\mathbb{R})$, which is often termed a *shift-invariant* (SI) space [13], [14]. The class of functions that can be represented in the form (10) is quite large. For example, choosing $g(t) = \text{sinc}(t/T)$ leads to the subspace of π/T -bandlimited signals. Other important examples include the space of spline functions (obtained by letting $g(t)$ be a B-spline function) and communication signals such as pulse-amplitude modulation (PAM) and quadrature amplitude modulation (QAM). Reconstruction in SI spaces from noiseless samples has been addressed in [15], [16] and extended to the noisy setting in [17]–[19].

Intuitively, every signal lying in a SI space with spacing T has one degree of freedom per T seconds (corresponding to one coefficient from the sequence $\{a[m]\}$). It is thus tempting to regard the rate of innovation of such signals as $1/T$. However, this is only true in an asymptotic sense and for compactly supported pulses $g(t)$. For any finite window size

T_0 , the T_0 -local rate of innovation ρ_{T_0} is generally larger. Specifically, suppose that the support of $g(t)$ is contained in $[t_a, t_b]$ and consider intervals of the form $[t, t + MT]$, where M is an integer. Then, due to the overlaps of the pulses, for any such interval we can only assure that there are no more than $M + \lceil (t_b - t_a)/T \rceil$ coefficients affecting the values of $x(t)$. Thus, the MT -local rate of innovation of signals of the form (10) is given by

$$\rho_{MT} = \frac{1}{T} \left(1 + \frac{\lceil \frac{t_b - t_a}{T} \rceil}{M} \right). \quad (11)$$

In particular, signals of the form (10) having a generator $g(t)$ which is not compactly supported have an infinite T_0 -local rate of innovation, for any finite T_0 . This is the case, for example, with bandlimited signals, which are therefore not FRI functions under our definition. As will be discussed in the sequel, this is not a flaw of the definition we use for the rate of innovation. Rather, it reflects the fact that it is impossible to recover any finite-duration segment $[T_1, T_2]$ of such signals from a finite number of measurements.

B. Nonlinearly-Distorted Shift-Invariant Spaces

In certain communication scenarios, nonlinearities are introduced in order to avoid amplitude clipping, an operation known as companding [20]. When the original signal lies in a SI space, the resulting transmission takes the form

$$x(t) = r \left(\sum_{m \in \mathbb{Z}} a[m]g(t - mT) \right), \quad (12)$$

where $r(\cdot)$ is a nonlinear, invertible function. Clearly, the MT -local rate of innovation ρ_{MT} of this type of signals is the same as that of the underlying SI function, and is thus given by (11). The recovery of nonlinearly distorted SI signals from noiseless samples was treated in [20]–[23]. We are not aware of research works treating the noisy case.

C. Union of Subspaces

Much of the FRI literature treats signal classes which are unions of subspaces [5], [6], [8], [24]. We now give examples of a few of these models.

1) *Finite Union of Subspaces*: There are various situations in which a continuous-time signal is known to belong to one of a finite set of spaces. One such signal model is described by

$$x(t) = \sum_{m \in \mathbb{Z}} \sum_{k=1}^K a_k[m]g_k(t - mT), \quad (13)$$

where $\{g_k(t)\}_{k=1}^K$ are a set of generators. In this model it is assumed that only $L < K$ out of the K sequences $\{a_1[m]\}_{m \in \mathbb{Z}}, \dots, \{a_K[m]\}_{m \in \mathbb{Z}}$ are not identically zero [25]. Therefore, the signal $x(t)$ is known to reside in one of $\binom{K}{L}$ spaces, each of which is spanned by an L -element subset of the set of generators $\{g_k(t)\}_{k=1}^K$. This class of functions can be used to describe multiband signals [24], [26]. However, the discrete nature of these models precludes analysis using the differential tools employed in the remainder of this paper.

Therefore, in this work we will focus on infinite unions of subspaces.

2) *Single-Burst Channel Sounding*: In certain medium identification and channel sounding scenarios, the echoes of a transmitted pulse $g(t)$ are analyzed to identify the positions and reflectance coefficients of scatterers in the medium [7], [27]. In these cases, the received signal has the form

$$x(t) = \sum_{\ell=1}^L a_\ell g(t - t_\ell), \quad (14)$$

where L is the number of scatterers and the amplitudes $\{a_\ell\}_{\ell=1}^L$ and time-delays $\{t_\ell\}_{\ell=1}^L$ correspond to the reflectance and location of the scatterers. Such signals can be thought of as belonging to a union of subspaces, where the parameters $\{t_\ell\}_{\ell=1}^L$ determine an L -dimensional subspace, and the coefficients $\{a_\ell\}_{\ell=1}^L$ describe the position within the subspace. In contrast with the previous example, however, in this setting we have a union of an infinite number of subspaces, since there are infinitely many possible values for the parameters t_1, \dots, t_L .

In this case, for any window of size $T_0 > \max_\ell \{t_\ell\} - \min_\ell \{t_\ell\}$, the T_0 -local rate of innovation is given by

$$\rho_{T_0} = \frac{2L}{T_0}. \quad (15)$$

3) *Periodic Channel Sounding*: Occasionally, channel sounding techniques consist of repeatedly probing the medium [28]. Assuming the medium does not change throughout the experiment, the result is a periodic signal

$$x(t) = \sum_{m \in \mathbb{Z}} \sum_{\ell=1}^L a_\ell g(t - t_\ell - mT). \quad (16)$$

As before, the set \mathfrak{X} of feasible signals is an infinite union of finite-dimensional subspaces in which $\{t_\ell\}_{\ell=1}^L$ determine the subspace and $\{a_\ell\}_{\ell=1}^L$ define the position within the subspace. The T_0 -local rate of innovation in this case coincides with (15).

4) *Semi-Periodic Channel Sounding*: There are situations in which a channel consists of L paths whose amplitudes change rapidly, but the time delays can be assumed constant throughout the duration of the experiment [5], [28], [29]. In these cases, the output of a channel sounding experiment will have the form

$$x(t) = \sum_{m \in \mathbb{Z}} \sum_{\ell=1}^L a_\ell[m]g(t - t_\ell - mT), \quad (17)$$

where $a_\ell[m]$ is the amplitude of the ℓ th path at the m th probing experiment. This is, once again, a union of subspaces, but here each subspace is infinite-dimensional, as it is determined by the infinite set of parameters $\{a_\ell[m]\}$. In this case, the MT -local rate of innovation can be shown to be

$$\rho_{MT} = \frac{L}{T} \left(1 + \frac{1 + \lceil \frac{t_b - t_a}{T} \rceil}{M} \right). \quad (18)$$

5) *Multiband Signals*: Multiuser communication channels are often characterized by a small number of utilized subbands interspersed by large unused frequency bands [26]. The resulting signal can be described as

$$x(t) = \sum_{n \in \mathbb{Z}} \sum_{\ell=1}^L a_\ell[n] g(t - nT) e^{j\omega_\ell t}, \quad (19)$$

where $\{a_\ell[n]\}_{n \in \mathbb{Z}}$ is the data transmitted by the ℓ th user, and ω_ℓ is the corresponding carrier frequency. In some cases the transmission frequencies are unknown [24], [26], resulting in an infinite union of infinite-dimensional subspaces. This setting is analogous in many respects to the semi-periodic channel sounding case; in particular, the MT -local rate of innovation can be shown to be the same as that given by (18).

IV. SUMMARY OF MAIN RESULTS

Before delving into the mathematical details, we provide in this section a high-level overview of our main contributions and summarize the resulting conclusions.

The overarching objective of this paper is to design and analyze sampling schemes for reconstructing FRI signals from noisy measurements. This goal is accomplished in three stages. First, we identify the best achievable MSE for estimating an FRI signal $x(t)$ from its continuous-time measurements $y(t) = x(t) + w(t)$, providing a fundamental lower bound which is independent of the sampling method. We then compare this continuous-time bound with the lowest possible MSE for a given sampling scheme, thus measuring the loss entailed in any particular technique. Finally, we provide a mechanism for choosing the optimal sampling kernels (in a specific Bayesian sense) utilizing a pre-specified sampling rate budget. Our results can be applied to specific families of FRI signals, but they also yield some general conclusions as to the relative difficulty of various classes of estimation problems. These general observations are summarized below.

A. Continuous-Time Bound

Our first goal in this paper is to derive the continuous-time CRB, which defines a fundamental limit on the accuracy with which an FRI signal can be estimated, regardless of the sampling technique. This bound turns out to have a particularly simple closed form expression which depends on the rate of innovation, but not on the class \mathfrak{X} of FRI signals being estimated. Specifically, under suitable regularity conditions, the MSE of any unbiased estimator \hat{x} satisfies

$$\frac{1}{T_0} \text{MSE}(\hat{x}, x) \geq \rho_{T_0} \sigma_c^2. \quad (20)$$

Thus, the rate of innovation can be given a new interpretation as the ratio between the best achievable MSE and the noise variance σ_c^2 . This is to be contrasted with the characterization of the rate of innovation in the noise-free case as the lowest sampling rate allowing for perfect recovery of the signal; indeed, when noise is present, perfect recovery is no longer possible.

B. Bound for Sampled Measurements

We next consider lower bounds for estimating $x(t)$ from *samples* of the signal $y(t)$. In this setting, the samples inherit the noise $w(t)$ embedded in the signal $y(t)$, and may suffer from additional discrete-time noise, for example, due to quantization. We derive the CRB for estimating $x(t)$ from sampled measurements in the presence of both types of noise. However, the combination of the two noise models complicates the mathematical analysis. Consequently, since the sampling noise model has been previously analyzed [9], [10], we focus in this paper on the assumption that the discrete-time noise is negligible.

In this setting, the sampled CRB can be designed so as to converge to the continuous-time bound as the sampling rate increases. Moreover, if the family \mathfrak{X} of FRI signals is contained in a finite-dimensional subspace \mathcal{M} of L_2 , then a sampling scheme achieving the continuous-time CRB can be constructed. Such a sampling scheme is obtained by choosing kernels which span the subspace \mathcal{M} , and yields samples which fully capture the information present in the signal $y(t)$. Contrariwise, if \mathfrak{X} is not contained in a finite-dimensional subspace, then no finite-rate sampling method achieves the continuous-time CRB. In this case, any increase in the sampling rate can improve performance, and the continuous-time bound is obtained only asymptotically.

It is interesting to examine this distinction from a union of subspaces viewpoint. Suppose that, as in the examples of Section III-C, the family \mathfrak{X} can be described as a union of an infinite number of subspaces $\{\mathcal{U}_\alpha\}$ indexed by the continuous parameter α , so that

$$\mathfrak{X} = \bigcup_{\alpha} \mathcal{U}_\alpha. \quad (21)$$

In this case, a finite sampling rate captures all of the information present in the signal if and only if

$$\dim \left(\sum_{\alpha} \mathcal{U}_\alpha \right) < \infty \quad (22)$$

where $\dim(\mathcal{M})$ is the dimension of the subspace \mathcal{M} . By contrast, in the noise-free case, it has been previously shown [30] that the number of samples required to recover $x(t)$ is given by

$$\max_{\alpha_1, \alpha_2} \dim(\mathcal{U}_{\alpha_1} + \mathcal{U}_{\alpha_2}), \quad (23)$$

i.e., the largest dimension among sums of *two* subspaces belonging to the union. In general, the dimension of (22) will be much higher than (23), illustrating the qualitative difference between the noisy and noise-free settings. For example, if the subspaces \mathcal{U}_α are finite-dimensional, then (23) is also necessarily finite, whereas (22) need not be.

Nevertheless, one may hope that the structure embodied in \mathfrak{X} will allow *nearly* optimal recovery using a sampling rate close to the rate of innovation. This is certainly the case in many noise-free FRI settings. For example, there exist techniques which recover the pulse stream (14) from samples taken at the rate of innovation, despite the fact that in this case \mathfrak{X} is typically not contained in a finite-dimensional subspace.

However, this situation often changes when noise is added, in which case standard techniques improve considerably under oversampling. This empirical observation can be quantified using the CRB: as we show, the CRB for samples taken at the rate of innovation is substantially higher in this case than the optimal, continuous-time bound. This demonstrates that the sensitivity to noise is a fundamental aspect of estimating signals of the form (14), rather than a limitation of existing algorithms. On the other hand, other FRI models, such as the semi-periodic pulse stream (17), exhibit considerable noise resilience, and indeed in these cases the CRB converges to the continuous-time value much more quickly.

As we discuss in Section IX-E, the different levels of robustness to noise can be explained when the signal models are examined in a union of subspaces context. In this case, the parameters θ defining $x(t)$ can be partitioned into parameters defining the subspace \mathcal{U}_α and parameters pinpointing the position within the subspace. Our analysis hints that estimation of the position within a subspace is often easier than estimation of the subspace itself. Thus, when most parameters are used to select an intra-subspace position, estimation at the rate of innovation is successful, as occurs in the semi-periodic case (17). By contrast, when a large portion of the parameters define the subspace in use, a sampling rate higher than the rate of innovation is necessary; this is the case in the non-periodic pulse stream (14), wherein θ is evenly divided among subspace-selecting parameters $\{t_\ell\}$ and intra-subspace parameters $\{a_\ell\}$. Thus we see that the CRB, together with the union of subspaces viewpoint, provide valuable insights into the relative degrees of success of various FRI estimation techniques.

C. Choosing the Sampling Kernels

In some cases, one may choose the sampling system according to design specifications, leading to the question: What sampling kernels should be chosen given an allotted number of samples? We tackle this problem by adopting a Bayesian framework, wherein the signal $x(t)$ is a random process distributed according to a known prior distribution. We further assume that both the sampling and reconstruction techniques are linear. While nonlinear reconstruction methods are often used for estimating FRI signals, this assumption is required for analytical tractability, and is used only for the purpose of identifying sampling kernels. Once these kernels are chosen, they can be used in conjunction with nonlinear reconstruction algorithms.

Under these assumptions, we identify the sampling kernels yielding the minimal MSE. An additional advantage of our assumption of linearity is that in this case, the optimal kernels depend only on the autocorrelation

$$R_X(t, \tau) = \mathbb{E}\{x(t)x^*(\tau)\} \quad (24)$$

of the signal $x(t)$, rather than on higher-order statistics. Indeed, given a budget of N samples, the optimal sampling kernels are given by the N eigenfunctions of R_X corresponding to the N largest eigenvalues. This is reminiscent of the Karhunen–Loève transform (KLT), which can be used to identify the

optimal sampling kernels in the noiseless setting. However, in our case, shrinkage is applied to the measurements prior to reconstruction, as is typically the case with Bayesian estimation of signals in additive noise.

A setting of particular interest occurs when the autocorrelation R_X is cyclic, in the sense that

$$R_X(t, \tau) = R_X((t - \tau) \bmod T) \quad (25)$$

for some T . This scenario occurs, for example, in the periodic pulse stream (16) and the semi-periodic pulse stream (17), assuming a reasonable prior distribution on the parameters θ . It is not difficult to show that the eigenfunctions of R_X are given, in this case, by the complex exponentials

$$\psi_n(t) = \frac{1}{\sqrt{T}} e^{j\frac{2\pi}{T}nt}, \quad n \in \mathbb{Z}. \quad (26)$$

Furthermore, in the case of the periodic and semi-periodic pulse streams, the magnitudes of the eigenvalues of R_X are directly proportional to the magnitudes of the respective Fourier coefficients of the pulse shape $g(t)$. It follows that the optimal sampling kernels are the exponentials (26) corresponding to the largest Fourier coefficients of $g(t)$. This result is encouraging in light of recently proposed FRI reconstruction techniques which utilize exponential sampling kernels [6], and demonstrates the suitability of the Bayesian approach for designing practical estimation kernels.

V. MATHEMATICAL PREREQUISITES: CRB FOR GENERAL PARAMETER SPACES

In statistics and signal processing textbooks, the CRB is typically derived for parameters belonging to a finite-dimensional Euclidean space [31]–[33]. However, this result is insufficient when it is required to estimate a parameter x belonging to other Hilbert spaces, such as the L_2 space defined above. When no knowledge about the structure of the parameter x is available, a bound for estimating $x(t)$ from measurements contaminated by *colored* noise was derived in [34]. However, this bound does not hold when the noise $w(t)$ is white. Indeed, in the white noise case, it can be shown that no finite-MSE unbiased estimators exist, unless further information about $x(t)$ is available. For example, the naive estimator $\hat{x}(t) = y(t)$ has an error $\hat{x}(t) - x(t)$ equal to $w(t)$, whose variance is infinite.

In our setting, we are given the additional information that x belongs to the constraint set \mathfrak{X} of (5). To the best of our knowledge, the CRB has not been previously defined for any type of constraint set $\mathfrak{X} \in L_2$, a task which will be accomplished in the present section. As we show below, a finite-valued CRB can be constructed by requiring unbiasedness only within the constraint set \mathfrak{X} , as per (9). As we will see, the CRB increases linearly with the dimension of the manifold \mathfrak{X} . Thus, in particular, the CRB is infinite when $\mathfrak{X} = L_2$. However, for FRI signals, the dimension of \mathfrak{X} is finite by definition, implying that a finite-valued CRB can be constructed. Although the development of this bound invokes some deep concepts from measure theory, it is a direct analog of the CRB for finite-dimensional parameters [32, Th. 2.5.15].

To derive the bound in the broadest setting possible, in this section we temporarily generalize the scenario of Section II, and consider estimation of a parameter x belonging to an arbitrary measurable and separable Hilbert space \mathcal{H} . The MSE of an estimator \hat{x} in this setting is defined as $\text{MSE}(\hat{x}, x) = \mathbb{E}\{\|\hat{x} - x\|_{\mathcal{H}}^2\}$. The concept of bias can similarly be extended if one defines the expectation $\mathbb{E}\{v\}$ of a random variable $v : \Omega \rightarrow \mathcal{H}$ as an element $k \in \mathcal{H}$ such that $\langle k, \varphi \rangle = \mathbb{E}\{\langle v, \varphi \rangle\}$ for any $\varphi \in \mathcal{H}$. If no such element exist, the expectation is said to be undefined.

The derivation of the CRB requires the existence of a ‘‘probability density’’ $p_{\theta}(y)$ (more precisely, a Radon–Nikodym derivative) which is differentiable with respect to θ , and such that its differentiation with respect to θ can be interchanged with integration with respect to y . The CRB also requires the mapping h_{θ} between θ and x to be non-redundant and differentiable. The formal statement of these regularity conditions is specified below. For the measurement setting (6), with reasonable mappings h_{θ} , these conditions are guaranteed to hold, as we will demonstrate in the sequel; in this section, however, we list these conditions in full so that a more general statement of the CRB will be possible.

- P1) There exists a value $\theta_0 \in \Theta$ such that the measure P_{θ_0} dominates $\{P_{\theta} : \theta \in \Theta\}$. In other words, there exists a Radon–Nikodym derivative $p_{\theta}(y) \triangleq dP_{\theta}/dP_{\theta_0}$ such that, for any event $A \in \mathcal{U}$,

$$P_{\theta}(A) = \int_A p_{\theta}(y) P_{\theta_0}(dy). \quad (27)$$

- P2) For all y such that $p_{\theta}(y) > 0$, the functions $p_{\theta}(y)$ and $\log p_{\theta}(y)$ are continuously differentiable with respect to θ . We denote by $\partial p_{\theta}(y)/\partial \theta$ and $\partial \log p_{\theta}(y)/\partial \theta$ the column vectors of the gradients of these two functions.
- P3) The support $\{y \in \mathcal{Y} : p_{\theta}(y) > 0\}$ of $p_{\theta}(y)$ is independent of θ .
- P4) There exists a measurable function $q : \mathcal{Y} \times \Theta \rightarrow \mathbb{R}$ such that for all sufficiently small $\Delta > 0$, for all $i = 1, \dots, K$, for all y , and for all θ ,

$$\frac{1}{\Delta} |p_{\theta+\Delta e_i}(y) - p_{\theta}(y)| \leq q(y, \theta) \quad (28)$$

and such that for all θ ,

$$\int q^2(y, \theta) P_{\theta_0}(dy) < \infty. \quad (29)$$

In (28), e_i represents the i th column of the $K \times K$ identity matrix.

- P5) For each θ , the $K \times K$ Fisher information matrix (FIM)

$$\mathbf{J}_{\theta} \triangleq \mathbb{E}\left\{\left(\frac{\partial \log p_{\theta}(y)}{\partial \theta}\right)\left(\frac{\partial \log p_{\theta}(y)}{\partial \theta}\right)^*\right\} \quad (30)$$

is finite and invertible.

- P6) h_{θ} is Fréchet differentiable with respect to θ , in the sense that for each θ , there exists a continuous linear operator $\partial h_{\theta}/\partial \theta : \mathbb{R}^K \rightarrow L_2$ such that, for any sufficiently small $\delta \in \mathbb{R}^K$,

$$\frac{h_{\theta+\delta} - h_{\theta}}{\|\delta\|} = \frac{\partial h_{\theta}}{\partial \theta} \delta + o(\|\delta\|) \quad \text{as } \|\delta\| \rightarrow 0. \quad (31)$$

- P7) The null space of the mapping $\partial h_{\theta}/\partial \theta$ contains only the zero vector. This assumption is required to ensure that the mapping from θ to x is non-redundant, in the sense that there does not exist a parametrization of \mathfrak{X} in which the number of degrees of freedom is smaller than K .

We are now ready to state the CRB for the estimation of a parameter $x \in L_2[0, T]$ parameterized by a finite-dimensional vector θ . The proof of this theorem is given in Appendix A.

Theorem 1. *Let $\theta \in \Theta$ be a deterministic parameter, where Θ is an open set in \mathbb{R}^K . Let \mathcal{H} be a measurable, separable Hilbert space and let h_{θ} be a mapping from Θ to \mathcal{H} . Let $\{P_{\theta} : \theta \in \Theta\}$ be a family of probability measures over a measurable space (Ω, \mathcal{F}) , and let $y : \Omega \rightarrow \mathcal{Y}$ be a random variable, where \mathcal{Y} is a measurable Hilbert space. Assume regularity conditions P1–P6. Let $\hat{x} : \mathcal{Y} \rightarrow \mathcal{H}$ be an unbiased estimator of x from the measurements y such that*

$$\mathbb{E}\{\|\hat{x}(y)\|_{\mathcal{H}}^2\} < \infty. \quad (32)$$

Then, the MSE of \hat{x} satisfies

$$\text{MSE}(\hat{x}, x) \geq \text{Tr}\left[\left(\frac{\partial h_{\theta}}{\partial \theta}\right)^* \left(\frac{\partial h_{\theta}}{\partial \theta}\right) \mathbf{J}_{\theta}^{-1}\right] \quad (33)$$

where \mathbf{J}_{θ} is the FIM (30).

Theorem 1 enables us to obtain a lower bound on the estimation error of x based on the FIM for estimating θ . The latter can often be computed relatively easily since θ is a finite-dimensional vector. Even more conveniently, the trace on the right-hand side of (33) is taken over a $K \times K$ matrix, despite the involvement of continuous-time operators. Thus, the computation of (33) is often possible either analytically or numerically, a fact which will be used extensively in the sequel.

VI. CRB FOR CONTINUOUS-TIME MEASUREMENTS

We now apply Theorem 1 to the problem of estimating a deterministic signal x from continuous-time measurements y given by (6).

Theorem 2. *Let x be a deterministic function defined by (5), where $\theta \in \Theta$ is an unknown deterministic parameter and Θ is an open subset of \mathbb{R}^K . Suppose that Assumptions P6–P7 are satisfied. Then, the MSE of any unbiased, finite-variance estimator \hat{x} is bounded by*

$$\text{MSE}(\hat{x}, x) \geq K \sigma_c^2. \quad (34)$$

The bound of Theorem 2 can be translated to units of the rate of innovation ρ_{T_0} if we assume that the segment $[0, T_0]$ under analysis achieves the maximum (4), i.e., this is a segment containing the maximum allowed number of degrees of freedom. In this case $\rho_{T_0} = K/T_0$, and any unbiased estimator $\hat{x}(t)$ satisfies

$$\frac{\frac{1}{T_0} \mathbb{E}\left\{\int_0^{T_0} |x(t) - \hat{x}(t)|^2 dt\right\}}{\sigma_c^2} \geq \rho_{T_0}. \quad (35)$$

In the noisy setting, ρ_{T_0} loses its meaning as a lower bound on the sampling rate required for perfect recovery, since the

latter is no longer possible at any sampling rate. On the other hand, it follows from (35) that the rate of innovation gains an alternative meaning; namely, ρ_{T_0} is a lower bound on the ratio between the average MSE achievable by any unbiased estimator and the noise variance σ_c^2 , *regardless of the sampling method*.

Before formally proving Theorem 2, note that (34) has an intuitive geometric interpretation. Specifically, the constraint set \mathfrak{X} is a K -dimensional differential manifold in $L_2[0, T]$. In other words, for any point $x \in \mathfrak{X}$, there exists a K -dimensional subspace \mathcal{U} tangent to \mathfrak{X} at x . We refer to \mathcal{U} as the feasible direction subspace [35]: any perturbation of x which remains within the constraint set \mathfrak{X} must be in one of the directions in \mathcal{U} . Formally, \mathcal{U} can be defined as the range space of $\partial h_\theta / \partial \theta$.

If one wishes to use the measurements y to distinguish between x and its local neighborhood, then it suffices to observe the projection of y onto \mathcal{U} . Projecting the measurements onto \mathcal{U} removes most of the noise, retaining only K independent Gaussian components, each having a variance of σ_c^2 . Thus we have obtained an intuitive explanation for the bound of $K\sigma_c^2$ in Theorem 2. To formally prove this result, we apply Theorem 1 to the present setting, as follows.

Proof of Theorem 2: The problem of estimating the parameters θ from a continuous-time signal $y(t)$ of the form (6) was examined in [12, Example I.7.3], where the validity of Assumptions P1–P4 was demonstrated. It was further shown that the FIM $\mathbf{J}_\theta^{\text{cont}}$ for estimating θ from $y(t)$ is given by [12, *ibid.*]

$$\mathbf{J}_\theta^{\text{cont}} = \frac{1}{\sigma_c^2} \left(\frac{\partial h_\theta}{\partial \theta} \right)^* \left(\frac{\partial h_\theta}{\partial \theta} \right). \quad (36)$$

Our goal will be to use (36) and Theorem 1 to obtain a bound on estimators of the continuous-time function $x(t)$. To this end, observe that the FIM $\mathbf{J}_\theta^{\text{cont}}$ is finite since, by Assumption P6, the operator $\partial h_\theta / \partial \theta$ is a bounded operator into L_2 . Furthermore, by Assumption P7, $\partial h_\theta / \partial \theta$ has a trivial null space, and thus $\mathbf{J}_\theta^{\text{cont}}$ is invertible. Therefore, Assumption P5 has been demonstrated. We may consequently apply Theorem 1, which yields

$$\begin{aligned} \text{MSE}(\hat{x}, x) &\geq \sigma_c^2 \text{Tr} \left[\left(\frac{\partial h_\theta}{\partial \theta} \right)^* \left(\frac{\partial h_\theta}{\partial \theta} \right) \left(\left(\frac{\partial h_\theta}{\partial \theta} \right)^* \left(\frac{\partial h_\theta}{\partial \theta} \right) \right)^{-1} \right] \\ &= \sigma_c^2 \text{Tr}(\mathbf{I}_K) \\ &= K\sigma_c^2 \end{aligned} \quad (37)$$

thus completing the proof. \blacksquare

To illustrate the use of Theorem 2 in practice, let us consider as a simple example a signal $x(t)$ belonging to a finite-dimensional subspace \mathcal{G} . Specifically, assume that

$$x(t) = \sum_{k=1}^K a_k g_k(t) \quad (38)$$

for some coefficient vector $\theta = (a_1, \dots, a_K)^T$ and a given set of linearly independent functions $\{g_k\}$ spanning \mathcal{G} . This includes, for example, families of shift-invariant subspaces with a compactly supported generator (see Section III-A).

From Theorem 2, the MSE of any unbiased estimator of x is bounded by $K\sigma_c^2$, where K is the dimension of the subspace \mathcal{G} . We now demonstrate that this bound is achieved by the unbiased estimator

$$\hat{x} = \mathbf{P}_\mathcal{G} y \quad (39)$$

where $\mathbf{P}_\mathcal{G}$ is the orthogonal projector onto the subspace \mathcal{G} .

To verify that (39) achieves the CRB, let G denote the set transformation (2) associated with the functions $\{g_k\}_{k=1}^K$. One may then write $x = G\theta$ and $\mathbf{P}_\mathcal{G} = G(G^*G)^{-1}G^*$. Thus (39) becomes

$$\begin{aligned} \hat{x} &= G(G^*G)^{-1}G^*G\theta + G(G^*G)^{-1}G^*w \\ &= G\theta + \mathbf{P}_\mathcal{G} w \end{aligned} \quad (40)$$

and therefore

$$\mathbb{E}\{\|\hat{x} - x\|_{L_2}^2\} = \mathbb{E}\{\|\mathbf{P}_\mathcal{G} w\|_{L_2}^2\}. \quad (41)$$

Since \mathcal{G} is a K -dimensional subspace, it is spanned by a set of K orthonormal¹ functions $u_1, \dots, u_K \in L_2$. Thus

$$\mathbb{E}\{\|\mathbf{P}_\mathcal{G} w\|_{L_2}^2\} = \sum_{k=1}^K \mathbb{E}\{|\langle w, u_k \rangle|^2\} = K\sigma_c^2 \quad (42)$$

which demonstrates that \hat{x} indeed achieves the CRB in this case.

In practice, a signal is not usually estimated directly from its continuous-time measurements. Rather, the signal $y(t)$ is typically sampled and digitally manipulated. In the next section, we will compare the results of Theorem 2 with the performance achievable from sampled measurements, and demonstrate that in some cases, a finite-rate sampling scheme is sufficient to achieve the continuous-time bound of Theorem 2.

VII. CRB FOR SAMPLED MEASUREMENTS

In this section, we consider the problem of estimating $x(t)$ of (5) from a finite number of samples of the process $y(t)$ given by (6). Specifically, suppose our measurements are given by

$$c_n = \langle y, s_n \rangle + v_n = \int_0^{T_0} y(t) s_n^*(t) dt + v_n, \quad n = 1, \dots, N \quad (43)$$

where $\{s_n\}_{n=1}^N \subset L_2[0, T_0]$ are sampling kernels, and v_n is a discrete white Gaussian noise process, independent of $w(t)$, having mean zero and variance σ_d^2 . Note that the model (43) includes both continuous-time noise, which is present in the signal $y(t) = x(t) + w(t)$ prior to sampling, and digital noise v_n , which arises from the sampling process, e.g., as a result of quantization. In this section, we will separately examine the effect of each of these noise components.

From (6) and (43), it can be seen that the measurements c_1, \dots, c_N are jointly Gaussian with mean

$$\mu_n \triangleq \mathbb{E}\{c_n\} = \langle x, s_n \rangle \quad (44)$$

¹We require the new functions u_1, \dots, u_K since the functions g_1, \dots, g_K are not necessarily orthonormal. The choice of non-orthonormal functions g_1, \dots, g_K will prove useful in the sequel.

and covariance

$$\Gamma_{ij} \triangleq \text{Cov}(c_i, c_j) = \sigma_c^2 \langle s_i, s_j \rangle + \sigma_d^2 \delta_{ij}. \quad (45)$$

A somewhat unusual aspect of this estimation setting is that the choice of the sampling kernels $s_n(t)$ affects not only the measurements obtained, but also the statistics of the noise. One example of the impact of this fact is the following. Suppose first that no digital noise is present, i.e., $\sigma_d = 0$, and consider a modified set of sampling kernels $\{\tilde{s}_n(t)\}_{n=1}^N$ which are an invertible linear transformation of $\{s_n(t)\}_{n=1}^N$, so that

$$\tilde{s}_n(t) = \sum_{i=1}^N \mathbf{B}_{ni} s_i(t) \quad (46)$$

where $\mathbf{B} \in \mathbb{R}^{N \times N}$ is an invertible matrix. Then, the resulting measurements $\tilde{\mathbf{c}}$ are given by $\tilde{\mathbf{c}} = \mathbf{B}\mathbf{c}$, and similarly the original measurements \mathbf{c} can be recovered from $\tilde{\mathbf{c}}$. It follows that these settings are equivalent in terms of the accuracy with which x can be estimated. In particular, the FIM for estimating x in the two settings is identical [12, Th. I.7.2].

When digital noise is present in addition to continuous-time noise, the sampling schemes $\{s_n(t)\}$ and $\{\tilde{s}_n(t)\}$ are no longer necessarily equivalent, since the gain introduced by the transformation \mathbf{B} will alter the ratio between the energy of the signal and the digital noise. The two estimation problems are then equivalent if and only if \mathbf{B} is a unitary transformation.

How should one choose the space $\mathcal{S} = \text{span}\{s_1, \dots, s_N\}$ spanned by the sampling kernels? Suppose for a moment that there exist elements in the range space of $\partial h_\theta / \partial \theta$ which are orthogonal to \mathcal{S} . This implies that one can perturb x in such a way that the constraint set \mathcal{X} is not violated, without changing the distribution of the measurements \mathbf{c} . This situation occurs, for example, when the number of measurements N is smaller than the dimension K of the parametrization of \mathcal{X} . While it may still be possible to reconstruct some of the information concerning x from these measurements, this is an undesirable situation from an estimation point of view. Thus we will assume henceforth that

$$\mathcal{R}\left(\frac{\partial h_\theta}{\partial \theta}\right) \cap \mathcal{S}^\perp = \{\mathbf{0}\}. \quad (47)$$

As an example of the necessity of the condition (47), consider again the signal (38), which belongs to a K -dimensional subspace $\mathcal{G} \subset L_2$ spanned by the functions g_1, \dots, g_K . In this case it is readily seen that for any vector \mathbf{v}

$$\frac{\partial h_\theta}{\partial \theta} \mathbf{v} = \sum_{k=1}^K \mathbf{v}_k g_k(t). \quad (48)$$

Since the functions $\{g_k\}$ span the space \mathcal{G} , this implies that $\mathcal{R}(\partial h_\theta / \partial \theta) = \mathcal{G}$, and therefore the condition (47) can be written as

$$\mathcal{G} \cap \mathcal{S}^\perp = \{\mathbf{0}\} \quad (49)$$

which is a standard requirement in the design of a sampling system for signals belonging to a subspace \mathcal{G} [36].

By virtue of Theorem 1, a lower bound on unbiased estimation of x can be obtained by first computing the FIM $\mathbf{J}_\theta^{\text{samp}}$ for estimating θ from \mathbf{c} . This yields the following

result. For simplicity of notation, in this theorem we assume that the function h_θ and the sampling kernels s_n are real. If complex sampling kernels are desired (as will be required in the sequel), the result below can still be used by translating each measurement to an equivalent pair of real-valued samples.

Theorem 3. *Let x be a deterministic real function defined by (5), where $\theta \in \Theta$ is an unknown deterministic parameter and Θ is an open subset of \mathbb{R}^K . Assume regularity conditions P6–P7, and let \hat{x} be an unbiased estimator of x from the real measurements $\mathbf{c} = (c_1, \dots, c_N)^T$ of (43). Then, the FIM $\mathbf{J}_\theta^{\text{samp}}$ for estimating θ from \mathbf{c} is given by*

$$\mathbf{J}_\theta^{\text{samp}} = \left(\frac{\partial h_\theta}{\partial \theta}\right)^* S (\sigma_c^2 S^* S + \sigma_d^2 \mathbf{I}_N)^{-1} S^* \left(\frac{\partial h_\theta}{\partial \theta}\right) \quad (50)$$

where S is the set transformation corresponding to the functions $\{s_n\}_{n=1}^N$. If (47) holds, then $\mathbf{J}_\theta^{\text{samp}}$ is invertible. In this case, any finite-variance, unbiased estimator \hat{x} for estimating x from \mathbf{c} satisfies

$$\text{MSE}(\hat{x}, x) \geq \text{Tr} \left[\left(\frac{\partial h_\theta}{\partial \theta}\right)^* \left(\frac{\partial h_\theta}{\partial \theta}\right) (\mathbf{J}_\theta^{\text{samp}})^{-1} \right]. \quad (51)$$

Proof: In the present setting, the FIM $\mathbf{J}_\theta^{\text{samp}}$ is given by [31]

$$\mathbf{J}_\theta^{\text{samp}} = \left(\frac{\partial \boldsymbol{\mu}}{\partial \theta}\right)^* \boldsymbol{\Gamma}^{-1} \left(\frac{\partial \boldsymbol{\mu}}{\partial \theta}\right) \quad (52)$$

where the matrix $\boldsymbol{\Gamma} \in \mathbb{R}^{N \times N}$ is defined by (45) and the matrix $\partial \boldsymbol{\mu} / \partial \theta \in \mathbb{R}^{N \times K}$ is given by

$$\left(\frac{\partial \boldsymbol{\mu}}{\partial \theta}\right)_{nk} = \frac{\partial \mu_n}{\partial \theta_k} \quad (53)$$

with μ_n defined in (44).

By the definition of the set transformation, the ij th element of the $N \times N$ matrix $S^* S$ is given by

$$(S^* S)_{ij} = \langle \mathbf{S} \mathbf{e}_j, \mathbf{S} \mathbf{e}_i \rangle = \langle s_j, s_i \rangle \quad (54)$$

where \mathbf{e}_i is the i th column of the $N \times N$ identity matrix. Therefore, we have

$$\boldsymbol{\Gamma} = \sigma_c^2 S^* S + \sigma_d^2 \mathbf{I}_N. \quad (55)$$

Similarly, observe that

$$\left(S^* \frac{\partial h_\theta}{\partial \theta}\right)_{nk} = \left\langle \frac{\partial h_\theta}{\partial \theta} \tilde{\mathbf{e}}_k, \mathbf{S} \mathbf{e}_n \right\rangle = \left\langle \frac{\partial h_\theta}{\partial \theta_k}, s_n \right\rangle = \frac{\partial \mu_n}{\partial \theta_k} \quad (56)$$

where $\tilde{\mathbf{e}}_k$ is the k th column of the $K \times K$ identity matrix. Thus

$$\frac{\partial \boldsymbol{\mu}}{\partial \theta} = S^* \frac{\partial h_\theta}{\partial \theta}. \quad (57)$$

Substituting (55) and (57) into (52) yields the required expression (50).

We next demonstrate that if (47) holds, then $\mathbf{J}_\theta^{\text{samp}}$ is invertible. To see this, note that from (50) we have

$$\mathcal{N}(\mathbf{J}_\theta^{\text{samp}}) = \mathcal{N}\left(S^* \left(\frac{\partial h_\theta}{\partial \theta}\right)\right). \quad (58)$$

Now, consider an arbitrary function $f \in \mathcal{R}(\partial h_\theta / \partial \theta)$. If (47) holds, then f is not orthogonal to the subspace \mathcal{S} . Therefore,

$\langle f, s_n \rangle \neq 0$ for at least one value of n , and thus by (3), $S^* f \neq 0$. This implies that

$$\mathcal{N}\left(S^* \left(\frac{\partial h_\theta}{\partial \theta}\right)\right) = \mathcal{N}\left(\frac{\partial h_\theta}{\partial \theta}\right) = \{\mathbf{0}\}. \quad (59)$$

Combined with (58), we conclude that $\mathcal{N}(\mathbf{J}_\theta^{\text{samp}}) = \{\mathbf{0}\}$. This demonstrates that $\mathbf{J}_\theta^{\text{samp}}$ is invertible, proving Assumption P5. Moreover, in the present setting, Assumptions P1–P4 are fulfilled for any value of θ_0 [32]. Applying Theorem 1 yields (51) and completes the proof. ■

In the following subsections we draw several conclusions from Theorem 3.

A. Discrete-Time Noise

Suppose first that $\sigma_c^2 = 0$, so that only digital noise is present. This setting has been analyzed previously [9], [10], and therefore only briefly examine the contrast with continuous-time noise. When only digital noise is present, its effects can be surmounted either by increasing the gain of the sampling kernels, or by increasing the number of measurements. These intuitive conclusions can be verified from Theorem 3 as follows. Assume that condition (47) holds, and consider the modified kernels $\tilde{s}_n(t) = 2s_n(t)$. The set transformation \tilde{S} corresponding to the modified kernels is $\tilde{S} = 2S$, and since $\sigma_c^2 = 0$, this implies that the FIM obtained from the modified kernels is given by $\tilde{\mathbf{J}}_\theta^{\text{samp}} = 4\mathbf{J}_\theta^{\text{samp}}$. Thus, a sufficient increase in the sampling gain can arbitrarily increase $\mathbf{J}_\theta^{\text{samp}}$ and consequently reduce the bound (51) arbitrarily close to zero.

Of course, from a practical point of view, increasing the gain also increases the likelihood that the sampled signal will exceed the dynamic range of the quantizer. It is therefore not feasible to arbitrarily increase the sampling gain. As an alternative, it is possible to increase the number of measurements. For example, suppose one simply repeats each measurement twice. Let S and \tilde{S} denote the transformations corresponding to the original and doubled sets of measurements. It can then readily be seen from the definition of the set transformation (2) and its adjoint (3) that $\tilde{S}\tilde{S}^* = 4SS^*$. Consequently, by the same argument, in the absence of continuous-time noise one can achieve arbitrarily low error by repeated measurements, or more generally, by increasing the sampling rate.

B. Continuous-Time Noise

As we have seen, sampling noise can be mitigated by increasing the sampling rate. Furthermore, digital noise is inherently dependent on the sampling scheme being used. Since our goal is to determine the fundamental performance limits regardless of the sampling technique, we will focus here and in subsequent sections on continuous-time noise. Thus, suppose that $\sigma_d^2 = 0$, so that only continuous-time noise is present. In this case, as we now show, it is generally impossible to achieve arbitrarily low reconstruction error, regardless of the sampling kernels used; indeed, it is never possible to outperform the continuous-time CRB of Section VI, which is typically nonzero. To see this formally, observe first that in

the absence of digital noise, the FIM for estimating θ can be written as

$$\begin{aligned} \mathbf{J}_\theta^{\text{samp}} &= \frac{1}{\sigma_c^2} \left(\frac{\partial h_\theta}{\partial \theta}\right)^* S(S^*S)^{-1} S^* \left(\frac{\partial h_\theta}{\partial \theta}\right) \\ &= \frac{1}{\sigma_c^2} \left(\frac{\partial h_\theta}{\partial \theta}\right)^* \mathbf{P}_S \left(\frac{\partial h_\theta}{\partial \theta}\right) \end{aligned} \quad (60)$$

where \mathbf{P}_S is the orthogonal projection onto the subspace \mathcal{S} . It is insightful to compare this expression with the FIM $\mathbf{J}_\theta^{\text{cont}}$ obtained from continuous-time measurements in (36). In both cases, a lower bound on the MSE for unbiased estimation of x was obtained from \mathbf{J}_θ by applying Theorem 1. Consequently, if it happens that $\mathbf{J}_\theta^{\text{cont}} = \mathbf{J}_\theta^{\text{samp}}$, then the continuous-time bound of Theorem 2 and the sampled bound of Theorem 3 coincide. Thus, if no digital noise is added, then it is possible (at least in terms of the performance bounds) that estimators based on the samples \mathbf{c} will suffer no degradation compared with the “ideal” estimator based on the entire set of continuous-time measurements. This occurs if and only if $\mathcal{R}(\partial h_\theta / \partial \theta) \subseteq \mathcal{S}$; in this case, the projection \mathbf{P}_S will have no effect on the FIM $\mathbf{J}_\theta^{\text{samp}}$, which will then coincide with $\mathbf{J}_\theta^{\text{cont}}$ of (36). In the remainder of this section, we will discuss several cases in which this fortunate circumstance arises.

C. Example: Sampling in a Subspace

The simplest situation in which samples provide all of the information present in the continuous-time signal is the case in which $x(t)$ belongs to a K -dimensional subspace \mathcal{G} of L_2 . This is the case, for example, when the signal lies in a shift-invariant subspace having a compactly supported generator (see Section III-A). As we have seen above (cf. (48)), in this scenario $\partial h_\theta / \partial \theta$ is a mapping onto the subspace \mathcal{G} . Assuming that there is no discrete-time noise, it follows from (60) that the optimal choice of a sampling space \mathcal{S} is \mathcal{G} itself. Such a choice requires $N = K$ samples and yields $\mathbf{J}_\theta^{\text{cont}} = \mathbf{J}_\theta^{\text{samp}}$. Of course, such an occurrence is not possible if the sampling process contributes additional noise to the measurements.

In some cases, it may be difficult to implement a set of sampling kernels spanning the subspace \mathcal{G} . It may then be desirable to choose a K -dimensional subspace \mathcal{S} which is close to \mathcal{G} but does not equal it. We will now compute the CRB for this setting and demonstrate that it can be achieved by a practical estimation technique. This will also demonstrate achievability of the CRB in the special case $\mathcal{S} = \mathcal{G}$. We first note from (2) and (48) that $\partial h_\theta / \partial \theta = G$, where G is the set transformation corresponding to the generators $\{g_k\}_{k=1}^K$. Furthermore, it follows from (49) that S^*G and G^*S are invertible $K \times K$ matrices [36]. Using Theorem 3, we thus find that the CRB is given by

$$\begin{aligned} \text{MSE}(\hat{x}, x) &\geq \sigma_c^2 \text{Tr}\left(G(G^*S(S^*S)^{-1}S^*G)^{-1}G^*\right) \\ &= \sigma_c^2 \text{Tr}\left(G(S^*G)^{-1}S^*S(G^*S)^{-1}G^*\right). \end{aligned} \quad (61)$$

It is readily seen that when $\mathcal{S} = \mathcal{G}$, the bound (61) reduces to $K\sigma_c^2$, which is (as expected) the continuous-time bound of Theorem 2. When $\mathcal{S} \neq \mathcal{G}$, the bound (61) will generally be higher than $K\sigma_c^2$, since $\mathbf{J}_\theta^{\text{samp}}$ of (60) will exceed $\mathbf{J}_\theta^{\text{cont}}$ of

(36). In this case, it is common to use the consistent, unbiased estimator [15], [36]

$$\hat{x} = G(S^*G)^{-1}\mathbf{c}. \quad (62)$$

As we now show, the bound (61) is achieved by this estimator. Indeed, observe that $\mathbf{c} = S^*y = S^*G\theta + S^*w$, and thus

$$\begin{aligned} \mathbb{E}\{\|\hat{x} - x\|_{L_2}^2\} &= \mathbb{E}\{\|G(S^*G)^{-1}S^*w\|_{L_2}^2\} \\ &= \mathbb{E}\{\text{Tr}(G(S^*G)^{-1}S^*ww^*S(G^*S)^{-1}G^*)\} \\ &= \text{Tr}(G(S^*G)^{-1}\text{Cov}(S^*w)(G^*S)^{-1}G^*). \end{aligned} \quad (63)$$

Note that $\text{Cov}(S^*w) = \text{Cov}(\mathbf{c})$, which by (55) is equal to $\sigma_c^2 S^*S$. Substituting this result into (63) and comparing with (61) verifies that \hat{x} achieves the CRB.

D. Nyquist-Equivalent Sampling

We refer to situations in which the dimension of the sampling space equals the dimension of the signal space as ‘‘Nyquist-equivalent’’ sampling schemes. In the previous section, we saw that Nyquist-equivalent sampling is possible using K samples when the signal lies in a K -dimensional subspace \mathfrak{X} , and that the resulting system achieves the continuous-time CRB. A similar situation occurs when the set of possible signals \mathfrak{X} is a subset of an M -dimensional subspace \mathcal{M} of L_2 with $M > K$. In this case, it can be readily shown that $\mathcal{R}(\partial h_\theta / \partial \theta) \subseteq \mathcal{M}$. Thus, by choosing $N = M$ sampling kernels such that $\mathcal{S} = \mathcal{M}$, we again achieve $\mathbf{J}_\theta^{\text{cont}} = \mathbf{J}_\theta^{\text{samp}}$, demonstrating that all of the information content in x has been captured by the samples. This is again a Nyquist-equivalent scheme, but the number of samples it requires is higher than the number of parameters K defining the signals. Therefore, in this case it is not possible to sample at the rate of innovation without losing some of the information content of the signal.

In general, the constraint set \mathfrak{X} will not be contained in any finite-dimensional subspace of L_2 . In such cases, it will generally not be possible to achieve the performance of the continuous-time bound using *any* finite number of samples, even in the absence of digital noise. This implies that in the most general setting, sampling above the rate of innovation can often improve the performance of estimation schemes. This conclusion will be verified by simulation in Section IX.

VIII. OPTIMAL SAMPLING KERNELS: A BAYESIAN VIEWPOINT

In this section, we address the problem of designing a sampling method which minimizes the MSE. One route towards this goal could be to minimize the sampled CRB of Theorem 3 with respect to the sampling space \mathcal{S} . However, the CRB is a function of the unknown parameter vector θ . Consequently, for each value of θ , there may be a different sampling space \mathcal{S} which minimizes the bound. To obtain a sampling method which is optimal on average over all possible choices of θ , we now make the additional assumption that the parameter vector θ is random and has a known distribution. Our goal, then, is to determine the sampling space \mathcal{S} that minimizes the MSE $\mathbb{E}\{\|\hat{x} - x\|_{L_2}^2\}$ within a class of allowed estimators. Note that

the mean is now taken over realizations of both the noise $w(t)$ and the parameter θ .

Since θ is random, the signal $x(t)$ is random as well. To make our discussion general, we will derive the optimal sampling functions for estimating a general random process $x(t)$ (not necessarily having realizations in \mathfrak{X} of (5)) from samples of the noisy process $y(t) = x(t) + w(t)$. We will then specialize the results to several specific types of FRI signals and obtain explicit expressions for the optimal sampling kernels in these scenarios.

Let $x(t)$ denote a zero-mean random process defined over $t \in [0, T_0]$, and suppose that its autocorrelation function

$$R_X(t, \eta) \triangleq \mathbb{E}\{x(t)x^*(\eta)\} \quad (64)$$

is continuous in t and η . Our goal is to estimate $x(t)$ based on a finite number N of samples of the signal $y(t) = x(t) + w(t)$, $t \in [0, T_0]$, where $w(t)$ is a white noise process (not necessarily Gaussian) with variance σ_c^2 which is uncorrelated with $x(t)$. We focus our attention on *linear sampling* schemes, i.e., we assume the samples are given by

$$c_n = \langle y, s_n \rangle. \quad (65)$$

Finally, we restrict the discussion to *linear estimation* methods, namely those techniques in which the estimate $\hat{x}(t)$ is constructed as

$$\hat{x}(t) = \sum_{n=1}^N c_n v_n(t), \quad (66)$$

for some set of reconstruction functions $\{v_n(t)\}_{n=1}^N$. It is important to note that for any given set of sampling functions $\{s_n(t)\}_{n=1}^N$, the minimum MSE (MMSE) estimator of $x(t)$ is often a nonlinear function of the measurements $\{c_n\}_{n=1}^N$. Indeed, typical FRI reconstruction techniques involve a nonlinear stage. Consequently, restricting the discussion to linear recovery schemes may seem inadequate. However, this choice has two advantages. First, as we will see, the optimal linear scheme is determined only by the second-order statistics of $x(t)$ and $w(t)$, whereas the analysis of nonlinear methods necessitates exact knowledge of their entire distribution functions. Second, it is not the final estimate $\hat{x}(t)$ that interests us in this discussion, but merely the set of optimal sampling functions. Once such a set is determined, albeit from a linear recovery perspective, it can be used in conjunction with existing nonlinear FRI techniques. As we will see in Section IX, the conclusions obtained through our analysis appear to apply to FRI techniques in general. Under the above assumptions, our goal is to design the sampling kernels $\{s_n(t)\}_{n=1}^N$ and reconstruction functions $\{v_n(t)\}_{n=1}^N$ such that the MSE (8) is minimized.

As can be seen from (65), we assume henceforth that only continuous-time noise is present in the sampling system. The situation is considerably more complicated in the presence of digital noise. First, without digital noise, one must choose only the subspace spanned by the sampling kernels, as the kernels themselves do not affect the performance; this is no longer the case when digital noise is added. Second, digital noise may give rise to a requirement that a particular measurement be repeated in order to average out the noise.

This is undesirable in the continuous noise regime, since the repeated measurement will contain the exact same noise realization.

A. Relation to the Karhunen–Loève Expansion and Finite-Dimensional Generalizations

The problem posed above is closely related to the Karhunen–Loève transform (KLT) [37], [38], which is concerned with the reconstruction of a random signal $x(t)$ from its noiseless samples. Specifically, one may express $x(t)$ in terms of a complete orthonormal basis $\{\psi_k(t)\}_{k=1}^{\infty}$ for L_2 as

$$x(t) = \sum_{k=1}^{\infty} \langle x, \psi_k \rangle \psi_k(t). \quad (67)$$

The goal of the KLT is to choose the functions $\{\psi_k(t)\}_{k=1}^{\infty}$ such that the MSE resulting from the truncation of this series after N terms is minimal. It is well known that the solution to this problem is given by the N -term truncation of the Karhunen–Loève expansion [37], [39].

Since $R_X(t, \eta)$ is assumed to be continuous in our setting, by Mercer’s theorem [39] it possesses a discrete set of eigenfunctions $\{\psi_k(t)\}_{k=1}^{\infty}$, which constitute an orthonormal basis for L_2 . These functions satisfy the equations

$$\lambda_k \psi_k(t) = \int_0^{T_0} R_X(t, \eta) \psi_k(\eta) d\eta, \quad (68)$$

in which the corresponding eigenvalues $\lambda_1 \geq \lambda_2 \geq \dots \geq 0$ are nonnegative and are assumed to be arranged in descending order. With these functions, (67) is known as the Karhunen–Loève expansion. It can be easily shown that the first N terms in this series constitute the best N -term approximation of $x(t)$ in an MSE sense [39]. In other words, in the noiseless case, the optimal sampling and reconstruction functions are $s_n(t) = v_n(t) = \psi_n(t)$.

In our setting, we do not have access to samples of $x(t)$ but rather only to samples of the noisy process $y(t)$. In this case, it is not clear *a priori* whether the optimal sampling and reconstruction filters coincide or whether they match the Karhunen–Loève expansion of $x(t)$.

The finite-dimensional analogue of our problem, in which x , y , and w are random vectors taking values in \mathbb{C}^M , was treated in [40], [41]. The derivation in these works, however, relied on the low-rank approximation property of the singular-value decomposition (SVD) of a matrix. The generalization of this concept to infinite-dimensional operators is subtle and will thus be avoided here. Instead, we provide a conceptually simple (if slightly cumbersome) derivation of the optimal linear sampling and reconstruction method for noisy signals. As we will see, it still holds that $s_n(t) = v_n(t) = \psi_n(t)$.

B. Optimal Sampling in Noisy Settings

As explained in Section VII, in the absence of discrete-time noise, the MSE is not affected by modifications of the sampling kernels which leave the set $\mathcal{S} = \text{span}\{s_1(t), \dots, s_N(t)\}$

unchanged. Thus, without loss of generality, we constrain $\{s_n(t)\}_{n=1}^N$ to satisfy

$$\left\langle s_n, \sigma_c^2 s_m + \int_0^T R_X(\cdot, \tau) s_m(\tau) d\tau \right\rangle = \delta_{m,n} \quad (69)$$

for every $m, n = 1, \dots, N$. This can always be done since the operator $R_Y : L_2 \rightarrow L_2$ defined by $(R_Y f)(t) = \int_0^T R_X(t, \tau) f(\tau) d\tau + \sigma_c^2 f(t)$ is positive definite. This choice is particularly convenient as it results in a set of uncorrelated samples $\{c_n\}$. Indeed

$$\begin{aligned} \mathbb{E}\{c_m c_n^*\} &= \mathbb{E}\left\{ \left(\int_0^T s_m^*(\tau) y(\tau) d\tau \right) \left(\int_0^T s_n^*(\eta) y(\eta) d\eta \right)^* \right\} \\ &= \iint_0^T s_m^*(\tau) \mathbb{E}\{y(\tau) y^*(\eta)\} s_n(\eta) d\tau d\eta \\ &= \iint_0^T s_m^*(\tau) R_X(\tau, \eta) s_n(\eta) d\tau d\eta \\ &\quad + \mathbb{E}\{s_n, w\} \langle s_m, w \rangle^* \\ &= \iint_0^T s_m^*(\tau) R_X^*(\eta, \tau) s_n(\eta) d\tau d\eta + \sigma_c^2 \langle s_n, s_m \rangle \\ &= \delta_{m,n}. \end{aligned} \quad (70)$$

We are now ready to determine the optimal sampling method. We begin by expressing the MSE (8) as

$$\begin{aligned} \int_0^T \mathbb{E}\{|x(t) - \hat{x}(t)|^2\} dt &= \int_0^T \mathbb{E}\{|x(t)|^2\} dt \\ &\quad - 2 \int_0^T \Re\{\mathbb{E}\{x^*(t) \hat{x}(t)\}\} dt + \int_0^T \mathbb{E}\{|\hat{x}(t)|^2\} dt. \end{aligned} \quad (71)$$

The first term in this expression does not depend on the choice of $\{s_n(t)\}_{n=1}^N$ and $\{v_n(t)\}_{n=1}^N$, and is therefore irrelevant for our purpose. Substituting (66) and (65), and using the fact that $w(t)$ is uncorrelated with $x(t)$, the second term can be written as

$$\begin{aligned} &\int_0^T 2\Re\left\{ \mathbb{E}\left\{ x^*(t) \sum_{n=1}^N c_n v_n(t) \right\} \right\} dt \\ &= \sum_{n=1}^N 2 \int_0^T \Re\left\{ \mathbb{E}\left\{ x^*(t) \int_0^T y(\tau) s_n^*(\tau) d\tau \right\} v_n(t) \right\} dt \\ &= \sum_{n=1}^N 2 \iint_0^T \Re\{\mathbb{E}\{x^*(t)(x(\tau) + w(\tau))\} s_n^*(\tau) v_n(t)\} d\tau dt \\ &= \sum_{n=1}^N 2 \iint_0^T \Re\{s_n^*(\tau) R_X(\tau, t) v_n(t)\} d\tau dt \\ &= \sum_{n=1}^N 2\Re\left\{ \left\langle v_n, \int_0^T R_X(\cdot, \tau) s_n(\tau) d\tau \right\rangle \right\}. \end{aligned} \quad (72)$$

Similarly, using the fact that $\{c_n\}_{n=1}^N$ are uncorrelated and

have unit variance (see (70)), the last term in (71) becomes

$$\begin{aligned} \int_0^T \mathbb{E} \left\{ \left| \sum_{n=1}^N c_n v_n(t) \right|^2 \right\} dt &= \sum_{m=1}^N \sum_{n=1}^N \mathbb{E} \{ c_m^* c_n \} \langle v_m, v_n \rangle \\ &= \sum_{n=1}^N \|v_n\|^2. \end{aligned} \quad (73)$$

Substituting (72) and (73) back into (71), we conclude that minimization of the MSE is equivalent to minimization of

$$\sum_{n=1}^N \left(\|v_n\|^2 - 2\Re \left\{ \left\langle v_n, \int_0^T R_X(\cdot, \tau) s_n(\tau) d\tau \right\rangle \right\} \right) \quad (74)$$

with respect to $\{s_n(t)\}_{n=1}^N$ and $\{v_n(t)\}_{n=1}^N$, subject to the set of constraints (69).

As a first stage, we minimize (74) with respect to the reconstruction functions $\{v_n(t)\}_{n=1}^N$. To this end, we note that the n th summand in (74) is lower bounded by

$$\begin{aligned} \|v_n\|^2 - 2\|v_n\| \left\| \int_0^T R_X(\cdot, \tau) s_n(\tau) d\tau \right\| \\ \geq - \left\| \int_0^T R_X(\cdot, \tau) s_n(\tau) d\tau \right\|^2, \end{aligned} \quad (75)$$

where we used the Cauchy–Schwarz inequality and the fact that $\min_z \{z^2 - 2bz\} = -b^2$. This bound is achieved by choosing

$$v_n(t) = \int_0^T R_X(t, \tau) s_n(\tau) d\tau, \quad (76)$$

thus identifying the optimal reconstruction functions.

Substituting (76) into (74), our goal becomes to maximize

$$\sum_{n=1}^N \left\| \int_0^T R_X(\cdot, \tau) s_n(\tau) d\tau \right\|^2 \quad (77)$$

with respect to the sampling functions $\{s_n(t)\}_{n=1}^N$. As we show in Appendix B, the maximum of this expression is achieved by any set of kernels of the form

$$s_n(t) = \sum_{k=1}^N \mathbf{A}_{k,n} (\lambda_k + \sigma_c^2)^{-\frac{1}{2}} \psi_k(t), \quad (78)$$

where \mathbf{A} is a unitary $N \times N$ matrix and λ_k and $\psi_k(t)$ are the eigenvalues and eigenfunctions of $R_x(t, \eta)$ respectively (see (68)). In particular, we can choose $\mathbf{A} = \mathbf{I}_N$, leading to

$$s_n(t) = \frac{1}{\sqrt{\lambda_n + \sigma_c^2}} \psi_n(t), \quad n = 1, \dots, N. \quad (79)$$

From (76), the optimal reconstruction kernels are given by

$$v_n(t) = \frac{\lambda_n}{\sqrt{\lambda_n + \sigma_c^2}} \psi_n(t), \quad n = 1, \dots, N. \quad (80)$$

The following theorem summarizes the result.

Theorem 4. *Let $x(t)$, $t \in [0, T]$ be a random process whose autocorrelation function $R_X(t, \eta)$ is jointly continuous in t and η . Assume that $y(t) = x(t) + w(t)$, where $w(t)$ is a*

white noise process uncorrelated with $x(t)$. Then, among all estimates $\hat{x}(t)$ of $x(t)$ having the form

$$\hat{x}(t) = \sum_{n=1}^N v_n(t) \int_0^T s_n^*(\tau) y(\tau) d\tau \quad (81)$$

the MSE (8) is minimized with $\{s_n(t)\}_{n=1}^N$ and $\{v_n(t)\}_{n=1}^N$ of (79) and (80) respectively. In these expressions, λ_n and $\psi_n(t)$ are the eigenvalues and eigenfunctions of $R_x(t, \eta)$ respectively (see (68)).

Interestingly, the optimal sampling and reconstruction functions in our noisy setting are similar to those dictated by the KLT. The only difference is that in the present scenario, the n th sample is shrunk by a factor of $\lambda_n/(\lambda_n + \sigma_c^2)$ prior to reconstruction. This ensures that the low-SNR measurements do not contribute to the recovery as much as their high-SNR counterparts. From the viewpoint of designing the sampling mechanism, however, this difference is of no importance.

As stated above, in practice one would generally favor nonlinear processing of the samples (namely, applying standard nonlinear FRI techniques) rather than a simple element-wise shrinkage. Thus, the importance of Theorem 4 for our purposes is in identifying that the eigenfunctions of $R_X(t, \tau)$ remain the optimal sampling kernels even in the noisy setting.

C. Example: Sampling in a Subspace

To demonstrate the utility of Theorem 4, we now revisit the situation in which $x(t)$ is given by (38) for some set of linearly independent functions $\{g_k(t)\}_{k=1}^K$ spanning a subspace $\mathcal{G} \in L_2$. We assume that the coefficients $\boldsymbol{\theta} = \{a_1, \dots, a_K\}^T$ form a zero-mean random vector and denote its autocorrelation matrix by \mathbf{R}_θ . In this case, the signal's autocorrelation function is given by

$$\begin{aligned} R_X(t, \eta) &= \mathbb{E}\{x(t)x^*(\eta)\} \\ &= \mathbb{E} \left\{ \sum_{k=1}^K a_k g_k(t) \sum_{\ell=1}^K a_\ell^* g_\ell^*(\eta) \right\} \\ &= \sum_{k=1}^K \sum_{\ell=1}^K g_k(t) g_\ell^*(\eta) (\mathbf{R}_\theta)_{k,\ell}. \end{aligned} \quad (82)$$

Consequently, the operator $R_X : L_2 \rightarrow L_2$ defined by $(R_X h)(t) = \int_0^T R_X(t, \eta) h(\eta) d\eta$ can be expressed as

$$R_x = G \mathbf{R}_\theta G^*, \quad (83)$$

where G is the set transformation (2) associated with $\{g_k\}_{k=1}^K$.

Now, let \mathbf{U} be a unitary matrix and let \mathbf{D} be a diagonal matrix, such that

$$\mathbf{U} \mathbf{D} \mathbf{U}^* = (G^* G)^{1/2} \mathbf{R}_\theta (G^* G)^{1/2}. \quad (84)$$

Since the dimension of $\mathcal{R}(G)$ is K , the operator R_X has at most K nonzero eigenvalues $\{\lambda_k\}_{k=1}^K$. Let Ψ denote the set transformation associated with the N eigenfunctions $\{\psi_n\}_{n=1}^N$ corresponding to the N largest eigenvalues, for some $N \leq K$. Then, it can be shown that

$$\Psi = G(G^* G)^{-1/2} \mathbf{U} \quad (85)$$

and the corresponding eigenvalues are

$$\lambda_n = \mathbf{D}_{n,n}. \quad (86)$$

To see this, note that according to (85), Ψ is an isometry, since

$$\Psi^* \Psi = \mathbf{U}^* (G^* G)^{-1/2} G^* G (G^* G)^{-1/2} \mathbf{U} = \mathbf{U}^* \mathbf{U} = \mathbf{I}_K. \quad (87)$$

Furthermore, (83) and (84) imply that $R_X = \Psi \mathbf{D} \Psi^*$. Consequently

$$R_X \Psi = \Psi \mathbf{D} \Psi^* \Psi = \Psi \mathbf{D}, \quad (88)$$

which proves the claim.

It is important to emphasize that the K functions $\{\psi_n(t)\}_{n=1}^K$ span \mathcal{G} . Therefore, if one is allowed to take $N = K$ samples, then the optimal choice is a set of kernels that span \mathcal{G} . This conclusion is compatible with the CRB analysis of the previous sections. However, the advantage of the Bayesian viewpoint is that it allows us to identify the optimal sampling space when less than K samples are allowed. For example, suppose that $\{g_n(t)\}$ are orthonormal, and the coefficients $\{a_n\}$ are uncorrelated. Then the optimal sampling space is the one spanned by the N functions $\{g_n(t)\}$ corresponding to the N largest-variance coefficients $\{a_n\}$.

A second example demonstrating the derivation of the optimal sampling kernels will be given in the next section.

IX. APPLICATION: CHANNEL ESTIMATION

In this section, we focus on a specific application of FRI signals, namely, that of estimating a signal consisting of a number of pulses having unknown positions and amplitudes [5]–[7]. More precisely, we consider periodic signals $x(t)$ of the form (16), which were discussed in Section III-C. These are T -periodic pulse sequences, in which the pulse shape $g(t)$ is known, but the amplitudes $\{a_\ell\}$ and delays $\{t_\ell\}$ are unknown. After analyzing periodic signals of this type, we will also compare estimation performance in this case with the semi-periodic family (17), and attempt to explain the empirically observed differences in stability between these two cases.

By defining the T -periodic function $h(t) = \sum_{n \in \mathbb{Z}} g(t - nT)$, we can write $x(t)$ of (16) as

$$x(t) = \sum_{\ell=1}^L a_\ell h(t - t_\ell). \quad (89)$$

Our goal is now to estimate $x(t)$ from samples of the noisy process $y(t)$ of (6). As before, we will assume that only continuous-time noise is present in the system. Since $x(t)$ is T -periodic, it suffices to recover the signal in the region $[0, T]$. In particular, we would like to identify the optimal sampling kernels for this setting, and to compare existing algorithms with the resulting CRB in order to determine when the optimal estimation performance is achieved.

Let

$$\tilde{h}_k \triangleq \frac{1}{T} \langle h, \varphi_k \rangle, \quad k \in \mathbb{Z} \quad (90)$$

be the Fourier series of $h(t)$, where $\varphi_k(t) = e^{j2\pi kt/T}$. The Fourier series of $x(t)$ is then given by

$$\tilde{x}_k \triangleq \frac{1}{T} \langle x, \varphi_k \rangle = \tilde{h}_k \sum_{\ell=1}^L a_\ell e^{-j\frac{2\pi}{T} k t_\ell}, \quad k \in \mathbb{Z}. \quad (91)$$

Let $\mathcal{K} = \{k \in \mathbb{Z} : \tilde{h}_k \neq 0\}$ denote the indices of the nonzero Fourier coefficients of $h(t)$. Suppose for a moment that \mathcal{K} is finite. It then follows from (91) that $x(t)$ also has a finite number of nonzero Fourier coefficients. Consequently, the set \mathfrak{X} of possible signals $x(t)$ is contained in the finite-dimensional subspace $\mathcal{M} = \text{span}\{\varphi_k\}_{k \in \mathcal{K}}$. Therefore, as explained in Section VII-D, choosing the $N = |\mathcal{K}|$ sampling kernels $\{s_n(t) = e^{-j2\pi n t/T}\}_{n \in \mathcal{K}}$ results in a sampled CRB which is equivalent to the continuous-time bound. This result is compatible with recent work demonstrating successful performance of FRI recovery algorithms using exponentials as sampling kernels [6].

Note, however, that this is a ‘‘Nyquist-equivalent’’ sampling scheme, i.e., the number of samples required $N = |\mathcal{K}|$ is potentially much higher than the number of degrees of freedom $2L$ in the signal $x(t)$ (see Section VII-D). This provides a theoretical explanation of the empirically recognized fact that sampling above the rate of innovation improves the performance of FRI techniques in the presence of noise [7], a fact which stands in contrast to the noise-free performance guarantees of many FRI algorithms.

Moreover, if there exists an infinite number of nonzero coefficients \tilde{h}_k , then in general the set \mathfrak{X} will not belong to any finite-dimensional subspace. Consequently, it will not be possible in this case for an algorithm based on a finite number of samples to achieve the performance obtainable from the complete signal $y(t)$. This occurs, for example, whenever the pulse $g(t)$ of (16) is time-limited. In such cases, any increase in the sampling rate will potentially continue to reduce the CRB, although the sampled CRB will converge to the asymptotic value of $\rho_{T_0} \sigma_c^2$ in the limit as the sampling rate increases.

A. Choosing the Sampling Kernels

An important question in the current setting is how to choose the sampling kernels so as to achieve the best possible performance under a limited budget of samples. This can be done via the Bayesian analysis provided in Section VIII. Assume, for example, that the time delays $\{t_\ell\}_{\ell=1}^L$ are independently drawn from a uniform distribution over the interval $[0, T]$. Furthermore, suppose that the amplitudes $\{a_\ell\}_{\ell=1}^L$ are mutually uncorrelated zero-mean random variables which are

independent of the time delays and have variance σ_a^2 . Then,

$$\begin{aligned}
R_X(t, \tau) &= \mathbb{E}\{x(t)x^*(\tau)\} \\
&= \sum_{k=1}^L \sum_{\ell=1}^L \mathbb{E}\{a_k a_\ell^*\} \mathbb{E}\{h(t-t_k)h^*(\tau-t_\ell)\} \\
&= \sigma_a^2 \sum_{\ell=1}^L \mathbb{E}\{h(t-t_\ell)h^*(\tau-t_\ell)\} \\
&= \sigma_a^2 L \frac{1}{T} \int_0^T h(t-t_\ell)h^*(\tau-t_\ell) dt_\ell \\
&= \sigma_a^2 L \sum_{k \in \mathbb{Z}} |\tilde{h}_k|^2 e^{j\frac{2\pi}{T}k(t-\tau)}, \tag{92}
\end{aligned}$$

where we used Parseval's theorem. It is easily verified that the eigenfunctions of $R_X(t, \tau)$ are given by

$$\psi_n(t) = \frac{1}{\sqrt{T}} e^{j\frac{2\pi}{T}nt}, \quad n \in \mathbb{Z} \tag{93}$$

and the corresponding eigenvalues are

$$\lambda_n = L\sigma_a^2 T |\tilde{h}_n|^2, \quad n \in \mathbb{Z}. \tag{94}$$

Therefore, the optimal set of N sampling functions is

$$s_n(t) = e^{j\frac{2\pi}{T}p_n t}, \quad n = 1, \dots, N \tag{95}$$

where p_n is the index of the n th largest Fourier coefficient $|\tilde{h}_{p_n}|$. The optimal linear recovery of $x(t)$ from the resulting samples is given by

$$\hat{x}(t) = \sum_{n=1}^N c_n \frac{L\sigma_a^2 |\tilde{h}_{p_n}|^2}{L\sigma_a^2 T |\tilde{h}_{p_n}|^2 + \sigma_c^2} e^{j\frac{2\pi}{T}p_n t}. \tag{96}$$

The performance of this estimator is poorer than state-of-the-art techniques, due to the restriction to linear reconstruction schemes. We recall that this technique is intended only for selecting the sampling kernels.

The above analysis again lends credence to the recently proposed time-delay estimation technique of Gedalyahu et al. [6], which makes use of complex exponentials as sampling functions. A disadvantage of this algorithm is that it can only handle a set of exponents with successive frequencies, while for general pulses, the indices of the N largest Fourier coefficients may be sporadic. As we will see in Section IX-C, this limitation may result in deteriorated performance of the algorithm in some cases.

B. Computing the CRB

Having identified the optimal sampling kernels (95), we would now like to compute the CRB for estimating $x(t)$ from the resulting samples. In order to compare these results with the continuous-time CRB, we assume that no digital noise is added in the sampling process. However, the calculations described below can be adapted without difficulty to situations containing both continuous-time and digital noise.

We assume for simplicity that $h(t)$ and $\{a_\ell\}$ are real-valued. Nonetheless, the sampling kernels chosen above are complex-valued, implying that Theorem 3 cannot be directly applied. Yet since $h(t)$ is real-valued, we have $|\tilde{h}_k| = |\tilde{h}_{-k}|$, and consequently the optimal sampling kernels consist of complex

conjugate pairs $e^{\pm j2\pi nt/T}$. Recall that the sampling kernels can be changed without affecting the CRB, as long as the subspace they span remains constant. Consequently, the CRB can be computed for the equivalent sampling kernels $\sin(2\pi nt/T)$ and $\cos(2\pi nt/T)$, which are real and can therefore be used in conjunction with the results of Section VII. We note that since the transition to these real-valued kernels is unitary, the CRB will not change even if digital noise is added. To be specific without complicating the notation, we assume that N is odd and that the DC component is included among the sampling kernels chosen in (95). We can then define the equivalent set of kernels

$$\begin{aligned}
\tilde{s}_0(t) &= 1, \\
\tilde{s}_n(t) &= \cos(2\pi p_n t/T), \quad n = 1, \dots, \frac{N-1}{2}, \\
\tilde{s}_{n+\frac{N+1}{2}}(t) &= \sin(2\pi p_n t/T), \quad n = 1, \dots, \frac{N-1}{2}. \tag{97}
\end{aligned}$$

We further define the parameter vector

$$\boldsymbol{\theta} = (a_1, \dots, a_L, t_1, \dots, t_L)^T \tag{98}$$

whose length is $K = 2L$.

Theorem 3 provides a two-step process for computing the CRB of the signal $x(t)$ from its samples. First, the FIM $\mathbf{J}_\theta^{\text{samp}}$ for estimating $\boldsymbol{\theta}$ is determined. Second, the formula (51) is applied to compute the CRB. While Theorem 3 also provides a means for calculating $\mathbf{J}_\theta^{\text{samp}}$, it is more convenient in the present setting to derive the FIM directly. This can be done by calculating the expectations μ_n of (44) and applying (52). In our setting, $\mu_n = \langle x, \tilde{s}_n \rangle$ are given by

$$\begin{aligned}
\mu_n &= \frac{\tilde{x}_{p_n} + \tilde{x}_{-p_n}}{2}, \quad n = 0, \dots, \frac{N-1}{2} \\
\mu_{n+\frac{N+1}{2}} &= \frac{\tilde{x}_{p_n} - \tilde{x}_{-p_n}}{2j}, \quad n = 1, \dots, \frac{N-1}{2} \tag{99}
\end{aligned}$$

where $\{\tilde{x}_n\}$ are the Fourier coefficients of $x(t)$. These coefficients depend in turn on the parameter vector $\boldsymbol{\theta}$, as shown in (91). Substituting μ_n into (52) yields a closed-form expression for $\mathbf{J}_\theta^{\text{samp}}$. Since the resulting formula is cumbersome and not very insightful, it is not explicitly written herein.

To obtain the sampled CRB, our next step is to compute the $2L \times 2L$ matrix

$$\mathbf{M} \triangleq \left(\frac{\partial h_\theta}{\partial \boldsymbol{\theta}} \right)^* \left(\frac{\partial h_\theta}{\partial \boldsymbol{\theta}} \right). \tag{100}$$

The function $h_\theta : \mathbb{R}^{2L} \rightarrow L_2$ maps a given parameter vector $\boldsymbol{\theta}$ to the resulting signal $x(t)$ as defined by (89). Differentiating this function with respect to $\boldsymbol{\theta}$, we find that the operator $\partial h_\theta / \partial \boldsymbol{\theta} : \mathbb{R}^{2L} \rightarrow L_2$ is defined by

$$\begin{aligned}
\left(\frac{\partial h_\theta}{\partial \boldsymbol{\theta}} \right) \mathbf{v} &= \mathbf{v}_1 h(t-t_1) + \dots + \mathbf{v}_L h(t-t_L) \\
&\quad - \mathbf{v}_{L+1} a_1 h'(t-t_1) - \dots - \mathbf{v}_{2L} a_L h'(t-t_L) \tag{101}
\end{aligned}$$

for any vector $\mathbf{v} \in \mathbb{R}^{2L}$.

One may now compute the ik th element of \mathbf{M} as

$$M_{ik} = \mathbf{e}_i^* \left(\frac{\partial h_{\boldsymbol{\theta}}}{\partial \boldsymbol{\theta}} \right)^* \left(\frac{\partial h_{\boldsymbol{\theta}}}{\partial \boldsymbol{\theta}} \right) \mathbf{e}_k = \left\langle \frac{\partial h_{\boldsymbol{\theta}}}{\partial \boldsymbol{\theta}} \mathbf{e}_i, \frac{\partial h_{\boldsymbol{\theta}}}{\partial \boldsymbol{\theta}} \mathbf{e}_k \right\rangle. \quad (102)$$

Thus, each element of \mathbf{M} is an inner product between two of the terms in (101). To calculate this inner product numerically for a given function $h(t)$, it is more convenient to use Parseval's theorem in order to convert the (continuous-time) inner product to a sum over Fourier coefficients. For example, in the case $1 \leq i, k \leq L$, we have

$$\begin{aligned} M_{ik} &= \int_0^T h(t-t_i) h^*(t-t_k) dt \\ &= T \sum_{n \in \mathbb{Z}} \tilde{h}_n e^{-j2\pi t_i n/T} \tilde{h}_n^* e^{j2\pi t_k n/T} \\ &= T \sum_{n \in \mathbb{Z}} |\tilde{h}_n|^2 e^{-j2\pi(t_k-t_i)n/T}, \quad 1 \leq i, k \leq L. \end{aligned} \quad (103)$$

An analogous derivation can be carried out when i or k are in the complementary range $L+1, \dots, 2L$.

Finally, having calculated the matrices $\mathbf{J}_{\boldsymbol{\theta}}^{\text{samp}}$ and \mathbf{M} , the CRB for sampled measurements is obtained using (51). We are now ready to compare this bound to the performance of practical estimators in some specific scenarios.


C. Effect of the Pulse Shape

In Fig. 1, we document several experiments comparing the CRB with the time-delay estimation technique of Gedalyahu et al. [6]. Specifically, we sampled the signal $x(t)$ of (16) using a set of exponential kernels, and used the matrix pencil method [42] to estimate $x(t)$ from the resulting measurements. Since we are considering only continuous-time noise, applying an invertible linear transformation to the sampling kernels has no effect on our performance bounds (see Section VII). The various kernels suggested in [6] amount to precisely such an invertible linear transformation, and the same performance bound applies to all of these approaches. Moreover, under the continuous-time noise model, it can be shown that these techniques also exhibit the same performance. For the same reason, the performance reported here is also identical to the method of Vetterli et al. [2].


In our experiments, a signal containing $L = 2$ pulses was constructed. The delays and amplitudes of the pulses were chosen randomly and are given by

$$\begin{aligned} a_1 &= 0.3204, & t_1 &= 0.6678, \\ a_2 &= 0.6063, & t_2 &= 0.9863. \end{aligned} \quad (104)$$


Modifications of these parameters does not appear to significantly affect the reported results, except when the time delays are close to one another, a situation which will be discussed in depth in Section IX-D. The pulse $h(t)$ consisted of $|\mathcal{K}| = 401$ nonzero Fourier coefficients at positions $\mathcal{K} = \{-200, \dots, 200\}$. The CRB is plotted as a function of the number of samples N , where the sampling kernels are given by $s_n(t) = e^{j2\pi n t/T}$ with $n \in \{-\lfloor N/2 \rfloor, \dots, \lfloor N/2 \rfloor\}$. This is done because the matrix pencil method requires the sampling kernels to have contiguous frequencies.



(a) The pulse $g(t)$ is a filtered Dirac with 401 Fourier coefficients.



(b) The pulse $g(t)$ contains 401 nonzero Fourier coefficients which decrease monotonically with the frequency.



(c) The pulse $g(t)$ is a filtered $\text{rect}(\cdot)$ with 401 Fourier coefficients.

Fig. 1. Comparison of the CRB and the performance of a practical estimator, as a function of the number of samples.

In Fig. 1(a), we chose $\tilde{h}_k = 1$ for $-200 \leq k \leq 200$ and $\tilde{h}_k = 0$ elsewhere; these are the low-frequency components of a Dirac delta function. The noise standard deviation was $\sigma_c = 10^{-5}$. In this case, for a fixed budget of N samples, any choice of N exponentials having frequencies in the range $-200 \leq k \leq 200$ is optimal according to the criterion of Section VIII. As expected, the sampled CRB achieves the continuous-time bound $K\sigma_c^2$ when $N \geq |\mathcal{K}|$. However, the CRB obtained at low sampling rates is higher by several orders of magnitude than the continuous-time limit. This indicates that the maxim of FRI theory, whereby sampling at the rate of innovation suffices for reconstruction, may not always hold in the presence of mild levels of noise. Indeed, if no noise is added in the present setting, then perfect recovery can be guaranteed using as few as $N = 4$ samples; yet even in the presence of mild noise, our bounds demonstrate that performance is quite poor unless the number of samples is increased substantially. This result may provide an explanation for the previously observed numerical instability of FRI techniques [2], [7].

As a further observation, we note that in this scenario, existing algorithms come very close to the CRB. Thus, the previously observed improvements achieved by oversampling are a result of fundamental limitations of low-rate sampling, rather than drawbacks of the specific technique used.

The same experiment is repeated in Fig. 1(b) with a pulse having Fourier coefficients $\tilde{h}_k = 1/(1 + 0.01k^2)$. Since the Fourier coefficients decrease with $|k|$, in this case our choice of low-frequency sampling kernels is optimal. However, the SNR of the measurements c_n decreases with n . As can be seen, this has a negative effect on the performance of the algorithm, which is not designed for high noise levels. Indeed, including low-SNR measurements causes the MSE not only to depart from the CRB, but eventually even to increase as more noisy samples are provided. In other words, one would do better to ignore the high-frequency measurements than to feed them to the recovery algorithm. Yet information is clearly present in these high-frequency samples, as indicated by the continual decrease of the CRB with increasing N . Thus, our analysis indicates that improved estimation techniques should be achievable in this case, in particular by careful utilization of low-SNR measurements.

The adverse effect of low-SNR measurements is exacerbated if, for a given N , one does not choose the N largest Fourier coefficients. This is demonstrated in Fig. 1(c). Here, the results of a similar experiment are plotted, in which $\tilde{h}_k = P \text{sinc}(nP/T)$, $-200 \leq k \leq 200$. These are the 401 lowest-frequency Fourier coefficients of a rectangular pulse having width P . In this case, the Fourier coefficients are no longer monotonically decreasing with $|k|$. Consequently, the sampling kernels $s_n(t) = e^{j2\pi nt/T}$ with $n \in \{-\lfloor N/2 \rfloor, \dots, \lfloor N/2 \rfloor\}$ do not correspond to the N largest Fourier coefficients, and thus are not optimal. In particular, for the chosen parameters, $|\tilde{h}_{25}| = |\tilde{h}_{-25}|$ are considerably smaller than the rest of the coefficients. When $N \geq 50$, the corresponding measurements are included, causing the MSE to deteriorate significantly.




Fig. 2. Comparison between the CRB and the performance of a practical estimator as a function of the pulse positions. The signal contains $L = 2$ pulses, the first of which is located at $t_1 = 0.5$. The MSE is plotted as a function of the position of the second pulse.

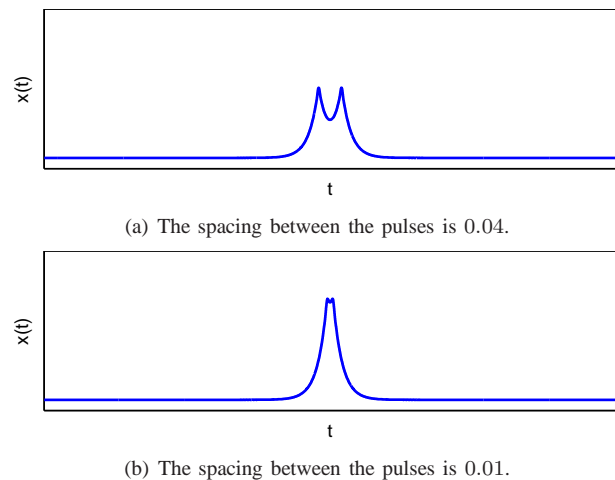


Fig. 3. Demonstration of the different levels of overlap between pulses.

D. Closely-Spaced Pulses

It is well-known that the estimation of pulse positions becomes ill-conditioned when several of the pulses are located close to one another. Intuitively, this is a consequence of the overlap between the pulses, which makes it more difficult to identify the precise location of each pulse. However, our goal is to estimate the signal $x(t)$ itself, rather than the positions of its constituent pulses. As we will see, for this purpose the effect of closely-spaced pulses is less clear-cut.

To study the effect of pulse position on the estimation error, we used a setup similar to the one of Fig. 1(b), with the following differences. First, a higher noise level of $\sigma_c = 10^{-3}$ was chosen. Second, the signal consisted of $L = 2$ pulses, with the first pulse at position $t_1 = 0.5$. The position of the second pulse was varied in the range $[0.3, 0.7]$ to demonstrate the effect of pulse proximity on the performance. The setting was otherwise identical to that of Section IX-C. In particular, recall that $T = 1$.

The results of this experiment are plotted in Fig. 2, which documents both the values of the sampled CRB and the actual

MSE obtained by the estimator of Gedalyahu et al. [6]. The continuous-time CRB is also plotted, although, as is evident from Theorem 2, this bound is a function only of the number of parameters determining the signal, and is therefore unaffected by the proximity of the pulses.

Several different effects are visible in Fig. 2. First, as the two pulses begin to come closer, both the CRB and the observed MSE increase by several orders of magnitude; this occurs when $|t_1 - t_2|$ is between about 0.15 and 0.03. (Of course, the precise distances at which these effects occur depend on the pulse width and other parameters of the experiment.) This level of proximity is demonstrated in Fig. 3(a). At this stage, the overlap between the pulses is sufficient to make it more difficult to estimate their positions accurately, but the separation between the pulses is still large, so that they are not mistaken for a single pulse.

As the pulses draw nearer each other, they begin to resemble a single pulse located at $(t_1 + t_2)/2$ (see Fig. 3(b)). Depending on the noise level, at some point the estimation algorithm will indeed identify the two pulses as one. Since our goal is to estimate $x(t)$ and not the pulse positions, such an “error” causes little deterioration in MSE. This is visible in Fig. 2 as the region in which the MSE of the practical algorithm ceases to deteriorate and ultimately decreases.

Interestingly, the CRB does not capture this improvement in performance. This failure is due to the fact that the CRB applies only to unbiased estimators, while the strategy utilized in [6] becomes biased for closely-spaced pulses. For an estimator to be unbiased, it is required that the mean estimate, averaged over noise realizations, will converge to the true value of $x(t)$, which has a form similar to that of Fig. 3(b). The expectation of an estimator reconstructing a single pulse will not have the form of two closely-spaced pulses; such an estimator is thus necessarily biased. In other words, the discrepancy observed here results from the fact that in this case, biased techniques outperform the best unbiased approach.

E. Non-Periodic and Semi-Periodic Signal Models

As we have seen above, the reconstruction of signals of the form (16) in the presence of noise is often severely hampered when sampled at or slightly above the rate of innovation. Rather than indicating a lack of appropriate algorithms, in many cases this phenomenon results from fundamental limits on the ability to recover such signals from noisy measurements. A similar effect was demonstrated [7] in the non-periodic (or finite) pulse stream model (14). In fact, if one is allowed to sample a non-periodic pulse stream with arbitrary sampling kernels, then by designing kernels having sufficiently large time-domain support, one can capture all or most of the energy in the signal. This setting then essentially becomes equivalent to a periodic signal model (16) in which the period is larger than the effective support of the pulse stream: One can imagine that the signal repeats itself beyond the sampled region, as this would not affect the measurements. Consequently, it is not surprising that the non-periodic model demonstrates substantial improvement in the presence of oversampling [7].

On the other hand, some types of FRI and union of subspace signals exhibit remarkable noise resilience, and do not appear

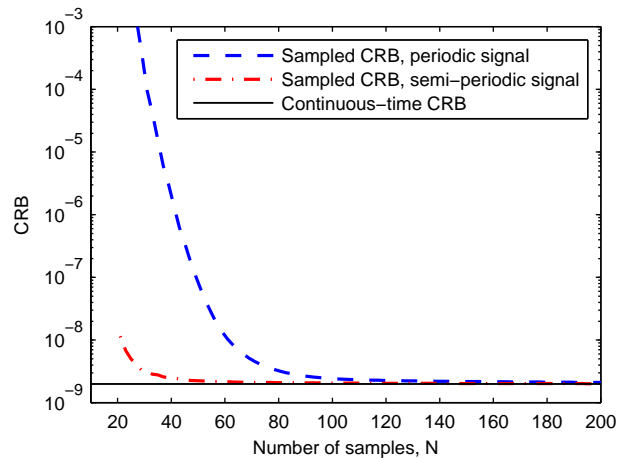


Fig. 4. Comparison between the CRB for a periodic signal (16) and a semi-periodic signal (17).

to require substantial oversampling in the presence of noise [5], [24]. As we now show, the CRB can be used to verify that such phenomena arise from a fundamental difference between families of FRI signals.

As an example, we compare the CRB for reconstructing the periodic signal (16) with the semi-periodic signal (17). Recall that in the former case, each period consists of pulses having unknown amplitudes and time shifts. By contrast, in the latter signal, the time delays are identical throughout all periods, but the amplitudes can change from one period to the next.

While these are clearly different types of signals, an effort was made to form a fair comparison between the reconstruction capabilities in the two cases. To this end, we chose an identical pulse $g(t)$ in both cases. We selected the signal segment $[0, T_0]$, where $T_0 = 1$, and chose the signal parameters so as to guarantee an identical T_0 -local rate of innovation. We also used identical sampling kernels in both settings: specifically, we chose the kernels (97) which measure the N lowest frequency components of the signal.

To simplify the analysis and focus on the fundamental differences between these settings, we will assume in this section that the pulses $g(t)$ are compactly supported, and that the time delays are chosen such that pulses from one period do not overlap with other periods. In other words, if the support of $g(t)$ is given by $[t_a, t_b]$, then we require

$$\min_{\ell} \{t_{\ell}\} > t_a \quad \text{and} \quad \max_{\ell} \{t_{\ell}\} < T - t_b. \quad (105)$$

Specifically, we chose the pulse $g(t)$ used in Fig. 1(b), which is compactly supported to a high approximation.

For the periodic signal, we chose $L = 10$ pulses with random delays and amplitudes, picked so as to satisfy the condition (105). A period of $T = 1$ was selected. This implies that the signal of interest is determined by $2L = 20$ parameters (L amplitudes and L time delays).

To construct a semi-periodic signal with the same number of parameters, we chose a period of $T = 1/9$ containing $L = 2$ pulses. The segment $[0, T_0]$ then contains precisely $M = 9$ periods, for a total of 20 parameters. While it may seem plausible to require the same number of periods for

both signals, this would actually disadvantage the periodic approach, as it would require the estimation of much more closely-spaced pulses.

The CRB for the periodic signal was computed as explained in Section IX-B, and the CRB for the semi-periodic signal can be calculated in a similar fashion. The results are compared with the continuous-time CRB in Fig. 4. Note that since the number of parameters to be estimated is identical in both signal models, the continuous-time CRB for the two settings coincides. Consequently, for a large number of measurements, the sampled bounds also converge to the same values. However, when the number of samples is closer to the rate of innovation, the bound on the reconstruction error for the semi-periodic signal is much lower than that of the periodic signal. As mentioned above, this is in agreement with previously reported findings for the two types of signals [2], [5], [6].

To find an explanation for this difference, it is helpful to recall that both signals can be described using the union of subspaces viewpoint (see Section III-C). Each of the signals in this experiment is defined by precisely 20 parameters, which determine the subspace to which the signal belongs and the position within this subspace. Specifically, the values of the time delays select the subspace, and the pulse amplitudes define a point within this subspace. Thus, in the above setting, the periodic signal contains 10 parameters for selecting the subspace and 10 additional parameters determining the position within it; whereas for the semi-periodic signal, only 2 parameters determine the subspace while the remaining 18 parameters set the location in the subspace. Evidently, identification of the subspace is challenging, especially in the presence of noise, but once the subspace is determined, the remaining parameters can be estimated using a simple linear operation (a projection onto the chosen subspace). Consequently, if many of the unknown parameters identify the position within a subspace, estimation can be performed more accurately. This may provide an explanation for the difference between the two examined signal models.

As further evidence in support of this explanation, we recall from Section III-C that the multiband signal model (19) can also be viewed as a union of subspaces. Here, again, the parameters $\{\omega_\ell\}_{\ell=1}^L$ determining the subspace (i.e., the utilized frequency bands) are far fewer than the parameters $\{a_\ell[n]\}$ selecting the point within the subspace (i.e., the content of each frequency band). In support of the proposed explanation, highly noise resistant algorithms can be constructed for the recovery of multiband signals [24], [26]. An even more extreme case is the single subspace setting, exemplified by shift-invariant signals (Section III-A). In this case, all of the signal parameters are used to determine the position within the one possible subspace. As we have seen in Section VII-C, in this case Nyquist-equivalent sampling at the rate of innovation achieves the continuous-time CRB.

X. CONCLUSION

In this paper, we studied the inherent limitations in recovering FRI signals from noisy measurements. We derived a continuous-time CRB which provides a lower bound on the

achievable MSE of any unbiased estimation method, regardless of the sampling mechanism. We showed that the rate of innovation ρ_{T_0} is a lower bound on the ratio between the average MSE achievable by any unbiased estimator and the noise variance σ_c^2 , *regardless of the sampling method*. This stands in contrast to the noise-free interpretation of ρ_{T_0} as the minimum sampling rate required for perfect recovery.

We next examined the CRB for estimating an FRI signal from a discrete set of noisy samples. We showed that the sampled bound is in general higher than the continuous-time CRB, and approaches it as the sampling rate increases. In general, the rate which is needed in order to achieve the continuous-time CRB is equal to the rate associated with the smallest subspace that encompasses all possible signal realizations. In particular, if a signal belongs to a union of subspaces, then the rate required to achieve the continuous-time bound is that associated with the sum of the subspaces. In some cases, this rate is finite, but in other cases the sum covers the entire space L_2 and no finite-rate technique achieves the CRB.

A consequence of these results is that oversampling can generally improve estimation performance. Indeed, our experiments demonstrate that sampling rates much higher than ρ_{T_0} are required in certain settings in order to approach the optimal performance. Furthermore, these gains can be substantial: In some cases, oversampling can improve the MSE by several orders of magnitude. We showed that the CRB can be used to determine which estimation problems require substantial oversampling to achieve stable performance. As a rule of thumb, it appears that for union of subspace signals, performance is improved at low rates if most of the parameters identify the position within the subspace, rather than the subspace itself. Our analysis can also be used to identify cases in which no existing algorithm comes close to the CRB, implying that better approaches can be constructed. In particular, it seems that existing algorithms do not deal well with measurement sets having a wide dynamic range.

Lastly, we addressed the problem of choosing the sampling kernels. This was done by adopting a Bayesian framework, so that an optimality criterion can be rigorously defined. Using a generalization of the KLT, we showed that the optimal kernels are the eigenfunctions of the autocorrelation function of the signal. In the context of time-delay estimation, these kernels are exponentials with appropriately chosen frequencies. This choice coincides with recent FRI techniques [5].

APPENDIX A PROOF OF THEOREM 1

The following notation will be used within this appendix. Let \mathcal{H}_1 and \mathcal{H}_2 be two measurable Hilbert spaces, and let (Ω, \mathcal{F}, P) be a probability space. Consider two random variables $u : \Omega \rightarrow \mathcal{H}_1$ and $v : \Omega \rightarrow \mathcal{H}_2$. Then, the notation $\mathbb{E}\{uv^*\}$ will be used to denote the linear operator $\mathcal{H}_2 \rightarrow \mathcal{H}_1$ such that, for any $h_1 \in \mathcal{H}_1$ and $h_2 \in \mathcal{H}_2$,

$$\langle h_1, \mathbb{E}\{uv^*\} h_2 \rangle_{\mathcal{H}_1} = \mathbb{E}\{\langle h_1, u \rangle_{\mathcal{H}_1} \langle v, h_2 \rangle_{\mathcal{H}_2}\} \quad (106)$$

if the expectation exists for all h_1 and h_2 .

We begin by stating two general lemmas which will be of use in the proof of Theorem 1.

Lemma 1. *Let \mathcal{H}_1 and \mathcal{H}_2 be two Hilbert spaces, and consider the operators*

$$\begin{aligned} A &: \mathcal{H}_1 \rightarrow \mathcal{H}_1, \\ B &: \mathcal{H}_2 \rightarrow \mathcal{H}_1, \\ C &: \mathcal{H}_2 \rightarrow \mathcal{H}_2. \end{aligned} \quad (107)$$

Suppose C is self-adjoint and invertible. Define the product Hilbert space $\mathcal{H}_1 \times \mathcal{H}_2$ in the usual manner, and suppose the operator $M : \mathcal{H}_1 \times \mathcal{H}_2 \rightarrow \mathcal{H}_1 \times \mathcal{H}_2$ defined by

$$M \begin{pmatrix} h_1 \\ h_2 \end{pmatrix} = \begin{pmatrix} Ah_1 + Bh_2 \\ B^*h_1 + Ch_2 \end{pmatrix} \quad (108)$$

is positive semidefinite (psd). Then,

$$A \succeq BC^{-1}B^* \quad (109)$$

in the sense that the $\mathcal{H}_1 \rightarrow \mathcal{H}_1$ operator $A - BC^{-1}B^$ is psd.*

Proof: Since M is psd, we have for any $h_1 \in \mathcal{H}_1$ and $h_2 \in \mathcal{H}_2$

$$\left\langle \begin{pmatrix} h_1 \\ h_2 \end{pmatrix}, M \begin{pmatrix} h_1 \\ h_2 \end{pmatrix} \right\rangle_{\mathcal{H}_1 \times \mathcal{H}_2} \geq 0 \quad (110)$$

which implies

$$\langle h_1, Ah_1 \rangle_{\mathcal{H}_1} + 2\Re[\langle h_1, Bh_2 \rangle_{\mathcal{H}_2}] + \langle b, Cb \rangle_{\mathcal{H}_2} \geq 0. \quad (111)$$

Choosing $h_2 = -C^{-1}B^*h_1$, we have that $\langle h_1, Bh_2 \rangle_{\mathcal{H}_2} = -\langle B^*h_1, C^{-1}B^*h_1 \rangle_{\mathcal{H}_1}$, which is real since C^{-1} is self-adjoint. It follows from (111) that

$$\langle h_1, Ah_1 \rangle_{\mathcal{H}_1} - \langle h_1, BC^{-1}B^*h_1 \rangle_{\mathcal{H}_1} \geq 0 \quad (112)$$

which leads to (109), as required. \blacksquare

Lemma 2. *Let \mathcal{H}_1 and \mathcal{H}_2 be two Hilbert spaces and let (Ω, \mathcal{F}, P) be a probability space. Let $u : \Omega \rightarrow \mathcal{H}_1$ and $v : \Omega \rightarrow \mathcal{H}_2$ be random variables, and suppose the expectations $\mathbb{E}\{uu^*\}$, $\mathbb{E}\{uv^*\}$, and $\mathbb{E}\{vv^*\}$ exist as linear operators as defined in (106). If $\mathbb{E}\{vv^*\}$ is invertible, then*

$$\mathbb{E}\{uu^*\} \succeq \mathbb{E}\{uv^*\} (\mathbb{E}\{vv^*\})^{-1} \mathbb{E}\{vu^*\}. \quad (113)$$

Proof: Let us denote $A = \mathbb{E}\{uu^*\}$, $B = \mathbb{E}\{uv^*\}$, and $C = \mathbb{E}\{vv^*\}$ and define the linear operator $M : \mathcal{H}_1 \times \mathcal{H}_2 \rightarrow \mathcal{H}_1 \times \mathcal{H}_2$ as in (108). From (106), for any $h_1 \in \mathcal{H}_1$ and $h_2 \in \mathcal{H}_2$ we have

$$\begin{aligned} & \left\langle \begin{pmatrix} h_1 \\ h_2 \end{pmatrix}, M \begin{pmatrix} h_1 \\ h_2 \end{pmatrix} \right\rangle_{\mathcal{H}_1 \times \mathcal{H}_2} \\ &= \langle h_1, Ah_1 \rangle_{\mathcal{H}_1} + 2\Re[\langle h_1, Bh_2 \rangle_{\mathcal{H}_1}] + \langle b, Cb \rangle_{\mathcal{H}_2} \\ &= \mathbb{E}\left\{ |\langle h_1, u \rangle_{\mathcal{H}_1}|^2 + 2\Re[\langle h_1, u \rangle_{\mathcal{H}_1} \langle v, h_2 \rangle_{\mathcal{H}_2}] \right. \\ & \quad \left. + |\langle h_2, v \rangle_{\mathcal{H}_2}|^2 \right\} \\ &= \mathbb{E}\left\{ |\langle h_1, u \rangle_{\mathcal{H}_1} + \langle v, h_2 \rangle_{\mathcal{H}_2}|^2 \right\} \\ & \geq 0. \end{aligned} \quad (114)$$

Thus M is a psd operator. Invoking Lemma 1 yields (113), as required. \blacksquare

We are now ready to prove Theorem 1.

Proof of Theorem 1: Throughout the proof, let θ be a fixed parameter and consider all functions as implicitly dependent on θ . Define the random variables

$$u : \Omega \rightarrow \mathcal{H} : \quad u(\omega) = \hat{x}(y(\omega)) - h_\theta, \quad (115)$$

$$v : \Omega \rightarrow \mathbb{R}^K : \quad v(\omega) = \frac{\partial \log p_\theta(y(\omega))}{\partial \theta}. \quad (116)$$

We then have the linear operators $\mathbb{E}\{vv^*\} : \mathbb{R}^K \rightarrow \mathbb{R}^K$, $\mathbb{E}\{uu^*\} : \mathcal{H} \rightarrow \mathcal{H}$, and $\mathbb{E}\{uv^*\} : \mathbb{R}^K \rightarrow \mathcal{H}$, which satisfy

$$\mathbb{E}\{vv^*\} = \mathbf{J}_\theta, \quad (117)$$

$$\langle \varphi_i, \mathbb{E}\{uu^*\} \varphi_j \rangle = \mathbb{E}\{\langle \varphi_i, u \rangle \langle u, \varphi_j \rangle\}, \quad (118)$$

$$\langle \varphi_i, \mathbb{E}\{uv^*\} \mathbf{e}_j \rangle = \mathbb{E}\left\{ \langle \varphi_i, u \rangle \frac{\partial \log p_\theta(y)}{\partial \theta_j} \right\}, \quad (119)$$

where $\{\varphi_n\}_{n \in \mathbb{Z}}$ denotes a complete orthonormal basis for \mathcal{H} . The operator $\mathbb{E}\{uu^*\}$ can be thought of as the covariance of \hat{x} , and is well-defined since, by (32), \hat{x} has finite variance. Indeed, we have

$$\sum_{i \in \mathbb{Z}} \langle \varphi_i, \mathbb{E}\{uu^*\} \varphi_i \rangle = \mathbb{E}\{\|\hat{x} - h_\theta\|_{L_2}^2\} < \infty \quad (120)$$

so that $\mathbb{E}\{uu^*\}$ is not only well-defined, but a trace class operator. Furthermore, $\mathbb{E}\{vv^*\} = \mathbf{J}_\theta$ is well-defined and invertible by Assumption P5. The operator $\mathbb{E}\{uv^*\}$ is thus also well-defined by virtue of the Cauchy–Schwarz inequality.

To prove the theorem, we will show that

$$\mathbb{E}\{uv^*\} = \frac{\partial h_\theta}{\partial \theta} \quad (121)$$

and then obtain the required result by applying Lemma 2. To demonstrate (121), observe that

$$\begin{aligned} \langle \varphi_i, \mathbb{E}\{uv^*\} \mathbf{e}_j \rangle &= \mathbb{E}\left\{ \langle \varphi_i, u \rangle \frac{\partial \log p_\theta(y)}{\partial \theta_j} \right\} \\ &= \int \langle \varphi_i, \hat{x}(y) - h_\theta \rangle \frac{1}{p(y; \theta)} \frac{\partial p(y; \theta)}{\partial \theta_j} p(y; \theta) P_{\theta_0}(dy) \\ &= \int \langle \varphi_i, \hat{x}(y) - h_\theta \rangle \lim_{\Delta \rightarrow 0} \frac{p(y; \theta + \Delta \mathbf{e}_j) - p(y; \theta)}{\Delta} P_{\theta_0}(dy). \end{aligned} \quad (122)$$

By Assumption P4, for any sufficiently small $\Delta > 0$ we have

$$\begin{aligned} & \left| \langle \varphi_i, \hat{x}(y) - h_\theta \rangle \frac{p(y; \theta + \Delta \mathbf{e}_j) - p(y; \theta)}{\Delta} \right| \\ & \leq |\langle \varphi_i, \hat{x}(y) - h_\theta \rangle| q(y, \theta). \end{aligned} \quad (123)$$

Let us demonstrate that the right-hand side of (123) is absolutely integrable. By the Cauchy–Schwarz inequality,

$$\begin{aligned} & \left(\int |\langle \varphi_i, \hat{x}(y) - h_\theta \rangle| q(y, \theta) P_{\theta_0}(dy) \right)^2 \\ & \leq \int |\langle \varphi_i, \hat{x}(y) - h_\theta \rangle|^2 P_{\theta_0}(dy) \cdot \int q^2(y, \theta) P_{\theta_0}(dy). \end{aligned} \quad (124)$$

The rightmost integral in (124) is finite by virtue of (29). As for the remaining integral, we have

$$\begin{aligned}
& \int |\langle \varphi_i, \hat{x}(y) - h_\theta \rangle|^2 P_{\theta_0}(dy) \\
& \leq \int \|\hat{x}(y) - h_\theta\|^2 P_{\theta_0}(dy) \\
& \stackrel{(a)}{\leq} \int (\|\hat{x}(y)\| + \|h_\theta\|)^2 P_{\theta_0}(dy) \\
& \stackrel{(b)}{\leq} \int \|\hat{x}(y)\|^2 P_{\theta_0}(dy) + \|h_\theta\|^2 \int P_{\theta_0}(dy) \\
& \quad + 2\|h_\theta\| \left(\int \|\hat{x}(y)\|^2 P_{\theta_0}(dy) \right)^{1/2} \\
& \stackrel{(c)}{<} \infty
\end{aligned} \tag{125}$$

where we have used the triangle inequality in (a), the Cauchy–Schwarz inequality in (b), and the assumption (32) that \hat{x} has finite energy in (c). We conclude that (123) is bounded by an absolutely integrable function, and we can thus apply the dominated convergence theorem to (122), obtaining

$$\begin{aligned}
\langle \varphi_i, \mathbb{E}\{uv^*\} \mathbf{e}_j \rangle &= \frac{\partial}{\partial \theta_j} \int \langle \varphi_i, \hat{x}(y) \rangle p(y; \boldsymbol{\theta}) P_{\theta_0}(dy) \\
&\quad - \langle \varphi_i, h_\theta \rangle \frac{\partial}{\partial \theta_j} \int p(y; \boldsymbol{\theta}) P_{\theta_0}(dy). \tag{126}
\end{aligned}$$

The second integral in (126) equals 1 and its derivative is therefore 0. Thus we have

$$\langle \varphi_i, \mathbb{E}\{uv^*\} \mathbf{e}_j \rangle = \frac{\partial \mathbb{E}\{\langle \varphi_i, \hat{x}(y) \rangle\}}{\partial \theta_j} = \frac{\partial \langle \varphi_i, h_\theta \rangle}{\partial \theta_j}. \tag{127}$$

On the other hand, note that the Fréchet derivative $\partial h_\theta / \partial \boldsymbol{\theta}$ of (31) coincides with the Gâteaux derivative of h_θ . In other words, for any vector $\mathbf{v} \in \mathbb{R}^K$, we have

$$\frac{\partial h_\theta}{\partial \boldsymbol{\theta}} \mathbf{v} = \lim_{\varepsilon \rightarrow 0} \frac{h_{\theta + \varepsilon \mathbf{v}} - h_\theta}{\varepsilon}. \tag{128}$$

It follows that

$$\left\langle \varphi_i, \frac{\partial h_\theta}{\partial \boldsymbol{\theta}} \mathbf{e}_j \right\rangle = \frac{\partial \langle \varphi_i, h_\theta \rangle}{\partial \theta_j}. \tag{129}$$

Since $\mathbb{E}\{uv^*\}$ and $(\partial h_\theta / \partial \boldsymbol{\theta})^*$ are both linear operators, (127) and (129) imply that the two operators are equal, demonstrating (121). Applying Lemma 2 and using the results (117) and (121), we have

$$\mathbb{E}\{(\hat{x} - h_\theta)(\hat{x} - h_\theta)^*\} \succeq \left(\frac{\partial h_\theta}{\partial \boldsymbol{\theta}} \right) \mathbf{J}_\theta^{-1} \left(\frac{\partial h_\theta}{\partial \boldsymbol{\theta}} \right)^*. \tag{130}$$

As we have seen, the left-hand side of (130) is trace class, and thus so is the right-hand side. Taking the trace of both sides of the equation, we obtain

$$\mathbb{E}\{\|\hat{x} - h_\theta\|^2\} \geq \text{Tr} \left(\left(\frac{\partial h_\theta}{\partial \boldsymbol{\theta}} \right) \mathbf{J}_\theta^{-1} \left(\frac{\partial h_\theta}{\partial \boldsymbol{\theta}} \right)^* \right) \tag{131}$$

which is equivalent to (33), as required. \blacksquare

APPENDIX B MAXIMIZATION OF (77)

The task of maximizing (77) is most easily accomplished by optimizing the coordinates of $s_n(t)$ in the orthonormal basis for $L_2[0, T]$ generated by the eigenfunctions of $R_X(t, \eta)$. Specifically, the function $s_n(t)$ can be written as

$$s_n(t) = \sum_{k=1}^{\infty} \alpha_k^n (\lambda_k + \sigma_c^2)^{-\frac{1}{2}} \psi_k(t), \tag{132}$$

with $\{\psi_k(t)\}_{k=1}^{\infty}$ and $\{\lambda_k\}_{k=1}^{\infty}$ of (68). (The coefficients $(\lambda_k + \sigma_c^2)^{-1/2}$ are inserted since they simplify the subsequent analysis.) Now, by Mercer's theorem, $R_X(t, \eta)$ can be expressed as

$$R_X(t, \eta) = \sum_{\ell=1}^{\infty} \lambda_\ell \psi_\ell(t) \psi_\ell^*(\eta), \tag{133}$$

where the convergence is absolute and uniform. Therefore

$$\begin{aligned}
& \int_0^T R_X(t, \tau) s_n(\tau) d\tau \\
&= \int_0^T \sum_{k=1}^{\infty} \alpha_k^n (\lambda_k + \sigma_c^2)^{-\frac{1}{2}} \psi_k(\tau) \sum_{\ell=1}^{\infty} \lambda_\ell \psi_\ell(t) \psi_\ell^*(\tau) d\tau \\
&= \sum_{k=1}^{\infty} \alpha_k^n \frac{\lambda_k}{(\lambda_k + \sigma_c^2)^{\frac{1}{2}}} \psi_k(t), \tag{134}
\end{aligned}$$

and consequently, by Parseval's theorem, (77) is given by

$$\begin{aligned}
& \sum_{n=1}^N \int_0^T \left| \int_0^T R_X(t, \tau) s_n(\tau) d\tau \right|^2 dt \\
&= \sum_{n=1}^N \int_0^T \left| \sum_{k=1}^{\infty} \alpha_k^n \frac{\lambda_k}{(\lambda_k + \sigma_c^2)^{\frac{1}{2}}} \psi_k(t) \right|^2 dt \\
&= \sum_{n=1}^N \sum_{k=1}^{\infty} |\alpha_k^n|^2 \frac{\lambda_k^2}{\lambda_k + \sigma_c^2}. \tag{135}
\end{aligned}$$

Similarly, using (132) and (134), we have

$$\iint_0^T s_m^*(t) R_X(t, \tau) s_n(\tau) d\tau dt = \sum_{k=1}^{\infty} \alpha_k^n (\alpha_k^m)^* \frac{\lambda_k}{\lambda_k + \sigma_c^2} \tag{136}$$

and by (132)

$$\sigma_c^2 \int_0^T s_m^*(t) s_n(\tau) dt = \sum_{k=1}^{\infty} \alpha_k^n (\alpha_k^m)^* \frac{\sigma_c^2}{\lambda_k + \sigma_c^2}. \tag{137}$$

Combining (136) and (137), the set of constraints (69) is translated to

$$\sum_{k=1}^{\infty} \alpha_k^m (\alpha_k^n)^* = \delta_{m,n} \tag{138}$$

for every $m, n = 1, \dots, N$. Consequently, our problem has now been reduced to

$$\max_{\{\alpha_k^n\}} \sum_{n=1}^N \sum_{k=1}^{\infty} |\alpha_k^n|^2 \frac{\lambda_k^2}{\lambda_k + \sigma_c^2} \quad \text{s.t.} \quad \sum_{k=1}^{\infty} \alpha_k^m (\alpha_k^n)^* = \delta_{m,n}. \tag{139}$$

We now show that the sequences $\{\alpha_k^n\}$ which solve (139) must satisfy $\alpha_k^n = 0$ for every $k > N$ and $n = 1, \dots, N$. To see this, assume to the contrary that the n th sequence satisfies $\alpha_\ell^n \neq 0$ for some $\ell > N$. We can then replace this sequence by a sequence $\{\tilde{\alpha}_k^n\}_{k \in \mathbb{Z}}$ satisfying

$$|\tilde{\alpha}_k^n|^2 = \begin{cases} |\alpha_k^n|^2 + a_k^2 & 1 \leq k \leq N \\ 0 & k = \ell \\ |\alpha_k^n|^2 & N < k \text{ and } k \neq \ell \end{cases} \quad (140)$$

where $\sum_{k=1}^N a_k^2 = |\alpha_\ell^n|^2$ (to ensure that $\sum_{k \in \mathbb{Z}} |\tilde{\alpha}_k^n|^2 = 1$). Such a set of coefficients $\{a_k\}_{k=1}^N$ can always be found since the N -term truncation of the remaining $N - 1$ sequences cannot span \mathbb{C}^N . With this sequence, the n th summand in the objective of (139) becomes

$$\begin{aligned} \sum_{k=1}^{\infty} |\tilde{\alpha}_k^n|^2 \frac{\lambda_k^2}{\lambda_k + \sigma_c^2} &= \sum_{k=1}^{\infty} |\alpha_k^n|^2 \frac{\lambda_k^2}{\lambda_k + \sigma_c^2} \\ &\quad + \left(\sum_{k=1}^N a_k^2 \frac{\lambda_k^2}{\lambda_k + \sigma_c^2} - |\alpha_\ell^n|^2 \frac{\lambda_\ell^2}{\lambda_\ell + \sigma_c^2} \right) \\ &\geq \sum_{k=1}^{\infty} |\alpha_k^n|^2 \frac{\lambda_k^2}{\lambda_k + \sigma_c^2} + \frac{\lambda_\ell^2}{\lambda_\ell + \sigma_c^2} \left(\sum_{k=1}^N a_k^2 - |\alpha_\ell^n|^2 \right) \\ &= \sum_{k=1}^{\infty} |\alpha_k^n|^2 \frac{\lambda_k^2}{\lambda_k + \sigma_c^2}, \end{aligned} \quad (141)$$

where we used the fact that $\lambda_k \geq \lambda_\ell$ for every $k < \ell$ and that $z^2/(a+z)$ is a monotone increasing function of z for all $z > 0$. This contradicts the optimality of $\{\alpha_k^n\}_{k \in \mathbb{Z}}$. Therefore, the set of sequences maximizing (139) satisfy $\alpha_k^n = 0$ for every $k > N$ and $n = 1, \dots, N$.

It remains to determine the optimal values of the first N elements of each of the N sequences $\{\alpha_k^n\}_{k \in \mathbb{Z}}$, $n = 1, \dots, N$. For this purpose, let \mathbf{A} denote the $N \times N$ matrix whose entries are $\mathbf{A}_{k,n} = \alpha_k^n$ and let $\mathbf{\Lambda}$ be a diagonal matrix with $\mathbf{\Lambda}_{k,k} = \lambda_k^2/(\lambda_k + \sigma_c^2)$. Then, the constraint (138) can be written as $\mathbf{A}^* \mathbf{A} = \mathbf{I}_N$, which is equivalent to $\mathbf{A} \mathbf{A}^* = \mathbf{I}_N$. Now, the objective in (139) can be expressed as

$$\begin{aligned} \sum_{n=1}^N \sum_{k=1}^{\infty} |\alpha_k^n|^2 \frac{\lambda_k^2}{\lambda_k + \sigma_c^2} &= \text{Tr}\{\mathbf{A}^* \mathbf{\Lambda} \mathbf{A}\} \\ &= \text{Tr}\{\mathbf{A} \mathbf{A}^* \mathbf{\Lambda}\} \\ &= \text{Tr}\{\mathbf{\Lambda}\}, \end{aligned} \quad (142)$$

which is independent of \mathbf{A} . Therefore, we conclude that any set of orthonormal sequences $\{\alpha_k^n\}_{k \in \mathbb{Z}}$, $n = 1, \dots, N$, whose elements vanish for every $k > N$ is optimal.

REFERENCES

- [1] C. Shannon, "A mathematical theory of communication," *Bell Syst. Tech. J.*, vol. 27, pp. 379423, 623656, 1948.
- [2] M. Vetterli, P. Marziliano, and T. Blu, "Sampling signals with finite rate of innovation," *IEEE Trans. Signal Process.*, vol. 50, no. 6, pp. 1417–1428, 2002.
- [3] I. Maravic and M. Vetterli, "Sampling and reconstruction of signals with finite rate of innovation in the presence of noise," *IEEE Trans. Signal Process.*, vol. 53, no. 8, pp. 2788–2805, Aug. 2005.
- [4] 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. 7, pp. 1741–1757, 2007.
- [5] K. Gedalyahu and Y. C. Eldar, "Time delay estimation from low rate samples: A union of subspaces approach," *IEEE Trans. Signal Process.*, vol. 58, no. 6, pp. 3017–3031, Jun. 2010.
- [6] K. Gedalyahu, R. Tur, and Y. C. Eldar, "Multichannel sampling of pulse streams at the rate of innovation," *IEEE Trans. Signal Process.*, submitted. [Online]. Available: <http://arxiv.org/abs/1004.5070>
- [7] R. Tur, Y. C. Eldar, and Z. Friedman, "Low rate sampling of pulse streams with application to ultrasound imaging," *IEEE Trans. Signal Process.*, submitted. [Online]. Available: <http://arxiv.org/abs/1003.2822>
- [8] Y. C. Eldar and M. Mishali, "Robust recovery of signals from a structured union of subspaces," *IEEE Trans. Inf. Theory*, vol. 55, no. 11, pp. 5302–5316, Nov. 2009.
- [9] 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, Feb. 2010.
- [10] P. Stoica and A. Nehorai, "MUSIC, maximum likelihood, and Cramer-Rao bound," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 37, no. 5, pp. 720–741, 1989.
- [11] H.-H. Kuo, "White noise theory," in *Handbook of Stochastic Analysis and Applications*, D. Kannan and V. Lakshmikantham, Eds. New York: Marcel Dekker, 2002, ch. 3.
- [12] I. A. Ibragimov and R. Z. Has'minskii, *Statistical Estimation: Asymptotic Theory*. New York: Springer, 1981.
- [13] M. Unser, "Sampling—50 years after Shannon," *Proc. IEEE*, vol. 88, pp. 569–587, Apr. 2000.
- [14] Y. C. Eldar and T. Michaeli, "Beyond bandlimited sampling," *IEEE Signal Process. Mag.*, vol. 26, no. 3, pp. 48–68, May 2009.
- [15] M. Unser and A. Aldroubi, "A general sampling theory for nonideal acquisition devices," *IEEE Trans. Signal Process.*, vol. 42, no. 11, pp. 2915–2925, Nov. 1994.
- [16] Y. C. Eldar and T. G. Dvorkind, "A minimum squared-error framework for generalized sampling," *IEEE Trans. Signal Process.*, vol. 54, no. 6, pp. 2155–2167, Jun. 2006.
- [17] M. Unser and T. Blu, "Generalized smoothing splines and the optimal discretization of the Wiener filter," *IEEE Trans. Signal Process.*, vol. 53, no. 6, pp. 2146–2159, 2005.
- [18] Y. C. Eldar and M. Unser, "Nonideal sampling and interpolation from noisy observations in shift-invariant spaces," *IEEE Trans. Signal Process.*, vol. 54, no. 7, pp. 2636–2651, Jul. 2006.
- [19] S. Ramani, D. Van De Ville, T. Blu, and M. Unser, "Nonideal Sampling and Regularization Theory," *IEEE Trans. Signal Process.*, vol. 56, no. 3, pp. 1055–1070, 2008.
- [20] H. J. Landau and W. L. Miranker, "The recovery of distorted bandlimited signals," *Journal of Math. Anal. and Appl.*, vol. 2, pp. 97–104, Feb. 1961.
- [21] I. W. Sandberg, "On the properties of some systems that distort signals I," *Bell Syst. Tech. J.*, vol. 42, pp. 2033–2047, September 1963.
- [22] T. G. Dvorkind, Y. C. Eldar, and E. Matusiak, "Nonlinear and nonideal sampling: theory and methods," *IEEE Trans. Signal Process.*, vol. 56, no. 12, pp. 5874–5890, Dec. 2008.
- [23] T. Faktor, T. Michaeli, and Y. C. Eldar, "Nonlinear and nonideal sampling revisited," *IEEE Signal Process. Lett.*, vol. 17, no. 2, pp. 205–208, February 2010.
- [24] M. Mishali and Y. C. Eldar, "From theory to practice: Sub-Nyquist sampling of sparse wideband analog signals," *IEEE J. Sel. Topics Signal Process.*, vol. 4, no. 2, pp. 375–391, Apr. 2010.
- [25] Y. C. Eldar, "Compressed sensing of analog signals in shift-invariant spaces," *IEEE Trans. Signal Process.*, vol. 57, no. 8, pp. 2986–2997, Aug. 2009.
- [26] M. Mishali and Y. C. Eldar, "Blind multi-band signal reconstruction: Compressed sensing for analog signals," *IEEE Trans. Signal Process.*, vol. 57, no. 3, pp. 993–1009, Mar. 2009.
- [27] J. G. Proakis, *Digital Communications*, 3rd ed. New York: McGraw-Hill, 1995.
- [28] A. Bruckstein, T. J. Shan, and T. Kailath, "The resolution of overlapping echos," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 33, no. 6, pp. 1357–1367, 1985.
- [29] W. U. Bajwa, K. Gedalyahu, and Y. C. Eldar, "Identification of underspread linear systems with application to super-resolution radar," *IEEE Trans. Signal Process.*, 2010, submitted. [Online]. Available: <http://arxiv.org/abs/1008.0851>
- [30] Y. M. Lu and M. N. Do, "A theory for sampling signals from a union of subspaces," *IEEE Trans. Signal Process.*, vol. 56, no. 6, pp. 2334–2345, Jun. 2008.
- [31] S. M. Kay, *Fundamentals of Statistical Signal Processing: Estimation Theory*. Englewood Cliffs, NJ: Prentice Hall, 1993.

- [32] E. L. Lehmann and G. Casella, *Theory of Point Estimation*, 2nd ed. New York: Springer, 1998.
- [33] J. Shao, *Mathematical Statistics*, 2nd ed. New York: Springer, 2003.
- [34] H. L. Van Trees, *Detection, Estimation, and Modulation Theory*. New York: Wiley, 1968, vol. 1.
- [35] Z. Ben-Haim and Y. C. Eldar, "The Cramér–Rao bound for estimating a sparse parameter vector," *IEEE Trans. Signal Process.*, vol. 58, no. 6, pp. 3384–3389, Jun. 2010.
- [36] Y. C. Eldar, "Sampling and reconstruction in arbitrary spaces and oblique dual frame vectors," *J. Fourier Analys.*, vol. 1, no. 9, pp. 77–96, Jan. 2003.
- [37] M. Loève, "Fonctions aléatoires du second ordre," in *Processus Stochastiques et Mouvement Brownien*, P. Lévy, Ed. Paris: Gauthier-Villars, 1948.
- [38] —, *Probability Theory*. Springer, 1978, vol. II.
- [39] O. P. Le Maître and O. M. Knio, *Spectral Methods for Uncertainty Quantification*. Springer, 2010.
- [40] Y. Yamashita and H. Ogawa, "Relative Karhunen–Loève transform," *IEEE Trans. Signal Process.*, vol. 44, no. 2, pp. 371–378, 1996.
- [41] Y. Hua and W. Liu, "Generalized Karhunen–Loève transform," *IEEE Signal Process. Lett.*, vol. 5, no. 6, pp. 141–142, 1998.
- [42] Y. Hua and T. K. Sarpar, "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, May 1990.

