FUNDAMENTAL PERFORMANCE LIMITS OF ANALOG-TO-DIGITAL COMPRESSION

A DISSERTATION
SUBMITTED TO THE DEPARTMENT OF ELECTRICAL ENGINEERING
AND THE COMMITTEE ON GRADUATE STUDIES
OF STANFORD UNIVERSITY
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

Alon Kipnis
September 2017

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

(Andrea J. Goldsmith)   Principal Adviser

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

(Tsachy Weissman)

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

(Abbas El Gamal)

Approved for the Stanford University Committee on Graduate Studies

Notational Conventions

Notation	Description
$\mathbb{R}$	real numbers
$\mathbb{N}$	natural numbers $1,2,\dots$
$\mathbb{Z}$	integers
$\mathbf{1}_A$	indicator function of the set $A \subset \mathbb{R}$
$\mu(A)$	Lebesgue measure of the set $A \subset \mathbb{R}$
$\text{supp}\, A$	support of the set A
$\mathbb{P}$	probability measure
$x(\cdot)$	real valued function over $\mathbb{R}$ representing a signal
$X(\cdot)$	stochastic process with index set $\mathbb{R}$
$X_T(\cdot)$	stochastic process $X(\cdot)$ limited to the interval $[-T/2, T/2]$
$Y[\cdot]$	stochastic process with index set $\mathbb{Z}$
$\mathbf{Y}[\cdot]$	vector-valued stochastic process with index set $\mathbb{Z}$
$\mathbb{E}[X]$	expectation of the random variable X
σ_X^2	variance of stationary process $X(\cdot)$
$\mathbf{L}_1(A)$	Banach space of absolutely Lebesgue integrable functions over domain A
$\mathbf{L}_2(A)$	Banach space of square Lebesgue integrable functions over domain A
$\lVert x(\cdot) \rVert$	$\mathbf{L}_2$ norm of the function $x(\cdot)$
$C_X(t,s)$	covariance function of the random process $X(\cdot)$
$S_X(f)$	power spectral density (PSD) of the continuous-time stationary process $X(\cdot)$
$S_Y\left(e^{2\pi i\phi}\right)$	power spectral density of the discrete-time stationary process $Y[\cdot]$
$\mathsf{mmse}(X\|Y)$	minimal mean squared error (MSE) in estimating X from Y
D	distortion in MSE per unit time
R	bitrate (bits per unit time)
f_s	sampling rate (samples per unit time)
$\bar{R} = R/f_s$	average rate of coding in bits
$K_H(t,\tau)$	bilinear kernel of a continuous linear operator between two spaces of signals
$H(f)$	frequency response of a linear time-invariant system
f_{Nyq}	Nyquist rate (the minimal length of the interval containing the support of $S_X(f)$).
f_{Lan}	Landau's rate (the spectral occupancy $\mu(\text{supp}\, S_X)$)

Preface

Analog-to-digital conversion is a fundamental operation in many electronic and communication systems. The principles describing the information loss as a result of transforming from continuous-time, continuous-amplitude signals to a sequence of bits lie at the intersection of sampling theory and quantization or lossy source coding theory. Classic results in sampling theory established the Nyquist rate of a signal, or, more precisely, its spectral occupancy, as the critical sampling rate above which the signal can be perfectly reconstructed from its samples. However, these results do not incorporate the quantization precision of the samples. Since it is impossible to obtain an exact digital representation of any continuous-amplitude sequence of samples, any digital representation of an analog signal will introduce some error, regardless of the sampling rate. This raises the question as to when sampling at the Nyquist rate is necessary. In other words, *is it possible to achieve the same or better performance by sampling at a rate lower than Nyquist without any further assumptions about the input signal other than a limited bitrate to describe the samples*?

When only quantization or limited bitrate is considered, the minimal distortion in encoding a random process subject to this bitrate limitation is described by Shannon's distortion-rate function (DRF). However, this DRF is given in terms of an optimization over a family of conditional probability distributions subject to a mutual information constraint and does not explicitly incorporate sampling and other inaccuracies arising from signal processing with a limited time-resolution. In fact, the standard achievability scheme in source coding assumes that the encoder can access the realization of the analog process, or expand it with respect to an analog basis function, without any restriction on the time resolution. Since this restriction often occurs in practice, the following question arises: *Given a finite number of samples per unit time from the analog process, what is the minimal distortion in recovering it from a digital encoding of these samples subject to a bitrate constraint ?*

The goal of this thesis is to address the two questions posed above by characterizing the minimal distortion that can be attained in recovering a random process using the most general form of quantization applied to its samples. By explicitly incorporating sampling into the minimum distortion optimization, the characterization of the minimal distortion bridges the missing theoretical gap between sampling theory and lossy data compression theory. By taking into account signal sampling and the associated distortion, the resulting distortion function generalizes and unifies sampling theory and source coding theory.

To Ima, Abba and Sarah – for shaping my dreams

Acknowledgments

First and foremost, I wish to express my great respect, admiration, and gratitude for my advisor, Andrea Goldsmith. On first meeting Andrea, her great intellect and clear vision for research were immediately apparent. Over the course of my PhD I also witnessed her skills as a mentor, in particular her commitment to me and others to creating a nurturing environment for future scientists. Despite her many professional endeavors, she always managed to make time available for her students. Her warm and energetic personality brought energy to the lab and made working with her fun and exciting.

Andrea contributed greatly to my individual development as a researcher. As a collaborator, she always brought out the best in me and helped round out my work. This thesis owes much to her contributions. On a personal level, I will forever remember Andrea's kindness and generosity over the years. She set a great example for what it means to be a researcher, advisor and mentor – one I will try to emulate throughout my future career.

During my PhD I have been also fortunate to work with Professor Yonina Eldar. Yonina's endless energy was the driving force behind many of the results in this thesis. I particularly admire Yonina's ability to identify important problems and engage multiple researchers, sometimes from very different fields, in solving them.

I am in debt to Professor Tsachy Weissman for multiple reasons. First, for influencing my decision to come to Stanford. Tsachy also mentored me in my first steps at Stanford. He introduced me to Andrea and to the research problem that later evolved into this thesis. Above all, Tsachy was a source for inspiring discussions and the first person to look for whenever I needed an advice.

I would like to thank Professor Robert Gray. Bob's endless knowledge in source coding literature was a valuable source of information, and his intellectual energy was an admired source of inspiration.

I would also like to thank to Professor Abbas El Gamal for his wonderful course and book on network information theory and for serving as a reader of my thesis and a committee member in my oral exam.

The Wireless System Lab at Stanford is a unique place where merit and friendship interact. I would like to thank the various friends I have been fortunate to interact and collaborate with: Stefano, Yair, Ernest, Javier, Yue, Neta, Jinyuan, Mahnoosh, Yoni, Yuxin, Alexandros, Nima, Mainak, Milind, Jeremy, Tom, and all other members and visitors of the WSL. I would also like to give a special thanks to Pat and Julia for handling administrative issues. It was a privilege to be your friend and colleague.

The final and most important gratitude goes to my family. I am grateful beyond words to my parents: for the loving home they built together, for showing me that hard work rewards itself, and for instilling in me the joy of discovery. Lastly, I would like to express my most heartfelt gratitude to my main collaborators in life: my soulmate Sarah – for her endless support and love, and my son Noam – for the joy and excitement he brings to our life. I am in debt to both for the meaning they give to my work.

Contents

List of Figures

Chapter 1

Introduction

1.1 Digital Representation of Analog Signals

Processing, storing, and communicating information that originates as an analog phenomenon involve the conversion of this information to bits. This conversion can be described by the combined effect of sampling and quantization, as illustrated in Figure 1.1. The digital representation in this procedure is achieved by first sampling the analog signal so as to represent it by a set of discrete-time samples and then quantizing these samples to a finite number of bits. Traditionally, these two operations are considered separately. The sampler is designed to minimize information loss due to sampling based on prior assumptions on the continuous-time input. The quantizer is designed to represent the samples as accurately as possible, subject to the constraint on the number of bits that can be used in the representation. The goal of this thesis is to revisit this paradigm by considering the joint effect of these two operations and to illuminate the dependency between them.

As motivation for exploring this dependency, consider the minimal information preserving sampling rate that arises in classical sampling theory due to Whittaker, Kotelnikov, Shannon, (WKS) and Landau [1, 2, 3]. These works establish the Nyquist rate or the *spectral occupancy* of the signal as the sampling rate above which a bandlimited signal can be perfectly reconstructed from its samples. This statement, however, focuses only on the minimal sampling rate required to perfectly reconstruct a bandlimited signal from its discrete

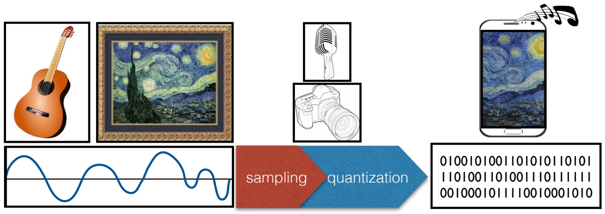

Figure 1.1: Analog-to-digital conversion is achieved by combining sampling and quantization.

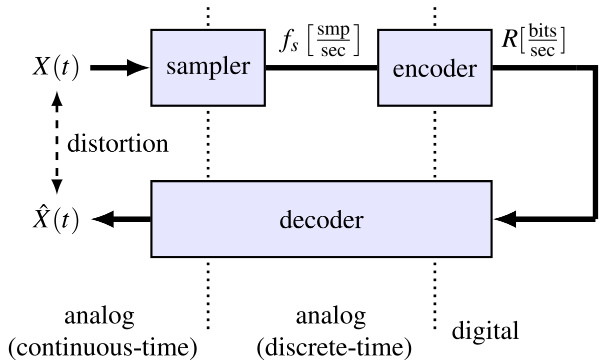

Figure 1.2: Analog-to-digital compression (ADX) and reconstruction setting. Our goal is to derive the minimal distortion between the signal and its reconstruction from any encoding at bitrate R of the samples of the signal taken at sampling rate f_s.

samples; it does not incorporate the quantization precision of the samples and does not apply to signals that are not bandlimited. It is in fact impossible to obtain an exact representation of any continuous-amplitude sequence of samples by a digital sequence of numbers due to finite quantization precision, and therefore any digital representation of an analog signal is prone to some amount of error. That is, no continuous amplitude signal can be reconstructed from its quantized samples with zero distortion regardless of the sampling rate, even when the signal is bandlimited. This limitation raises the following question: In converting a signal to bits via sampling and quantization at a given bit precision, can the signal be reconstructed from these samples with minimal distortion based on sub-Nyquist sampling? This thesis discusses this question by extending classical sampling theory to account for quantization and for non-bandlimited inputs. Namely, for an arbitrary random input signal and given a total budget of quantization bits, we consider the lowest sampling rate required to sample the signal such that reconstruction of the signal from its quantized samples results in minimal distortion.

1.2 Analog-to-Digital Compression (ADX)

The minimal distortion achievable in the presence of quantization depends on the particular way the signal is quantized or, more generally, encoded, into a sequence of bits. Since we are interested in the fundamental distortion limit in recovering an analog signal from its digital representation, we consider all possible encoding and reconstruction (decoding) techniques. As an example, in Figure 1.1 the smartphone display may be viewed as a reconstruction of the real world painting Starry Night from its digital representation. No matter how fine the smartphone screen, this recovery is not perfect since the digital representation of the analog image does not contain all information – due to loss of information during the transformation from analog to digital. Our goal is to analyze this loss as a function of limitations on the sampling mechanism and the number of bits used in the encoding. It is convenient to normalize this number of bits by the signal's free dimensions, that is, the dimensions along which new information is generated. For example, the free dimensions of a visual signal are usually the width and height of the frame, and the free dimension of an audio wave is time. For simplicity, we consider analog signals with a single free dimension, and we denote this dimension as *time*. Therefore, our restriction of the digital representation is given in terms of the *bitrate* and sampling rate – the number of bits per unit time and the number of samples per unit time, respectively. The random analog signal is assumed to be generated by some underlying source. Therefore, this signal is a continuous-time stochastic process and we refer to it as the *source signal*. The sampling and encoding operations are deterministic and can be designed based on the statistics of the process that are assumed to be known.

For an arbitrary continuous-time stochastic process with known statistics, the fundamental distortion limit due to the encoding of the signal using a limited bitrate is given by its Shannon distortion-rate function (DRF) [4, 5, 6], also referred to as the *information* DRF. This function provides the optimal tradeoff between the bitrate of the signal's digital representation and the distortion in recovering the original signal from this representation. The information DRF is described only in terms of the distortion criterion, the probability distribution on the continuous-time signal, and the maximal bitrate allowed in the digital representation. Consequently, the optimal encoding scheme that achieves the information DRF is a general mapping from continuous-time signal space to bits that does not consider any practical constraints in implementing such a mapping. In most applications, the encoding of analog signals into bits requires sampling followed by a digital mapping of these samples to bits. Therefore, in practice, the minimal distortion in recovering analog signals from their mapping to bits considers the digital encoding of the signal *samples*, with a constraint on both the *sampling rate* and the *bitrate* of the system. Here the sampling rate f_s is defined as the number of samples per unit time of the continuous-time signal; the bitrate R is the number of bits per unit time used in the representation of these samples. The resulting system describing our problem is illustrated in Figure 1.2, and is denoted as the *analog-to-digital compression* (ADX) setting.

The digital representation in this setting is obtained by transforming a continuous-time, continuous-amplitude random signal $X(\cdot)$ through a concatenated operation of a *sampler* and an *encoder*, resulting in a bit sequence. For instance, when the input signal $X(\cdot)$ is observed over a time interval of length T, then the sampler produces $\lfloor f_s T \rfloor$ samples, and the decoder maps these samples to $\lfloor TR \rfloor$ bits. The *decoder* estimates the original analog signal from this bit sequence. The *distortion* is defined to be the mean squared error (MSE) between the input signal $X(\cdot)$ and its reconstruction $\widehat{X}(\cdot)$. Since we are interested in the fundamental distortion limit subject to a sampling constraint, we allow optimization over the encoder, decoder, and the time horizon T. In addition, we also explore the optimal sampling mechanism but limit ourselves to the class of linear, continuous, and deterministic samplers.

The minimal distortion in ADX is bounded from below by two extreme cases of the sampling rate and the bitrate, as illustrated in Figure 1.3: (1) when the bitrate R is unlimited, the minimal ADX distortion reduces to the MSE in interpolating a signal from its samples at rate f_s. (2) When the sampling rate f_s is unlimited or above the Nyquist rate of the signal, the ADX distortion reduces to the information DRF, or simply the DRF,

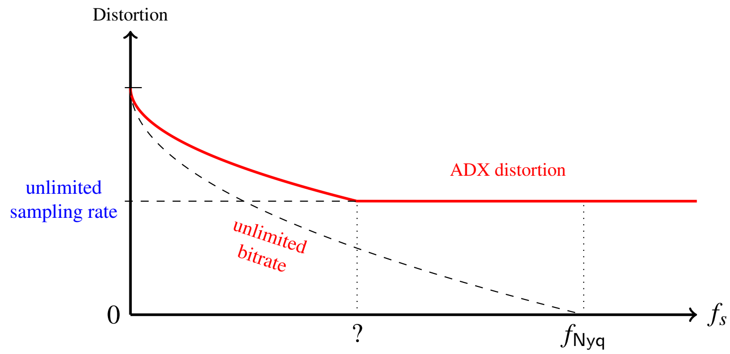

Figure 1.3: While Nyquist rate sampling is required to attain zero distortion without quantization constraints, the minimal sampling rate for attaining the minimal distortion achievable in the presence of quantization is usually below the Nyquist rate.

of the signal. Indeed, in this situation the optimal encoder can recover the original continuous-time source without distortion, and then encode this recovery in an optimal manner according to the optimal lossy source coding scheme that attains the DRF. In particular, we are interested in the minimal sampling rate for which the DRF, describing the minimal distortion subject to a bitrate constraint, is attained. As illustrated in Figure 1.3, and as will be explained in detail in Chapter 5, this sampling rate is usually below the Nyquist rate of the signal. We denote this minimal sampling rate as the *critical sampling rate* subject to a bitrate constraint, since it describes the minimal sampling rate required to attain the optimal performance in systems operating under quantization or bitrate restrictions. Therefore, the critical sampling rate extends the information-preserving or distortion-free sampling rate of WKS and Landau. It is only as the bitrate goes to infinity that sampling at the Nyquist rate or, equivalently, at the spectral occupancy, is necessary to attain zero distortion.

1.3 Applications

Figure 1.2 represents all systems that process information through sampling and are limited in the amount of memory they use, the number of states they can assume, or the number of bits they can transmit per unit time. These three sources of restrictions on system operations are explained in more detail below:

- **Memory** – Digital systems often operate under a constraint on their amount of memory or the states they can assume. Under such a restriction, the bitrate is the normalized amount of memory used over time (or dimension of the source signal). For example, consider a system of K states that analyzes information obtained by observing an analog signal for T seconds. The maximal bitrate of the system is $R = \log_2(K)/T$.

- **Power** – Emerging sensor network technologies, such as these for biomedical applications and "smart cities", use many low-cost sensors to collect data and transmit it to remote locations [7]. These sensors must operate under severe power restrictions, hence they are limited by the number of comparisons in their analog-to-digital operation. These comparisons are typically the most energy consuming part of the ADC unit, so that the overall power consumption in an ADC is proportional to their number [8, Sec 2.1]. In general, the number of comparisons is proportional to the bitrate, since any output of bitrate R is generated by at least R comparisons (although the exact number depends on the particular implementation of the ADC and may even grow exponentially in the bitrate [9]). Therefore, power restrictions lead to a bitrate constraint and to a non-zero MMSE distortion floor given by the DRF of the analog input signal.

An important scenario of power-restricted ADCs arises in wireless communication using millimeter waves [10]. Severe path-loss of electromagnetic waves in these frequencies is compensated by using a large number of receiver antennas. Each antenna is associated with an RF chain that includes an ADC unit. Due to the resulting large number of ADCs, power consumption is one of the major engineering challenges in mm-wave communication.

- **Communication** – Low-power sensors may also be limited by rates of communication to send their digital sensed information to a remote location. For example, consider a low-energy device collecting medical signals and transmitting its measurements wirelessly to a central processor (e.g. a smartphone). The communication rate from the sensor to the central processor depends on the capacity of the channel between them, which is a function of the available power for communication. When the power is limited, so is the capacity. As a result, the data rate associated with the digital representation of the sensed information cannot exceed this capacity limit since, without additional processing, there is no point in collecting more information than what can be communicated.

Specific examples for systems that operate under the above constraints are audio and video recorders, radio receivers, and digital cameras. Moreover, these restrictions also apply to signal processing techniques that use sampling and operate under bit-constraints, such as artificial neural networks [11], financial markets analyzers [12], and techniques to accelerate operations over large datasets by sampling [13].

For any of the above examples, the ADX setting provides the theoretical limits on estimation from sampled and quantized information. We note, however, that sampling and quantization restrictions may not only be the result of engineering limitations, but can also be inherent in the system model and the estimation problem. As an example, consider the estimation of an analog signal describing the price of a financial asset. Although we assume that the price follows some continuous-time behavior, the value of the asset is only "observed" when a transaction is reported. This limitation on the observation can be described by a sampling constraint. If transactions occur at intermittent times than this sampling is non-uniform. Moreover, it is often assumed that the instantaneous change of the price is given by a deterministic signal representing the *drift* plus an additive infinite bandwidth and stationary noise [12]. Therefore, the signal in question is of infinite bandwidth and sampling thus occurs below its Nyquist rate. In addition to the sampling constraint, it may be the case that the values of the transactions are hidden from us. The only information we receive is through

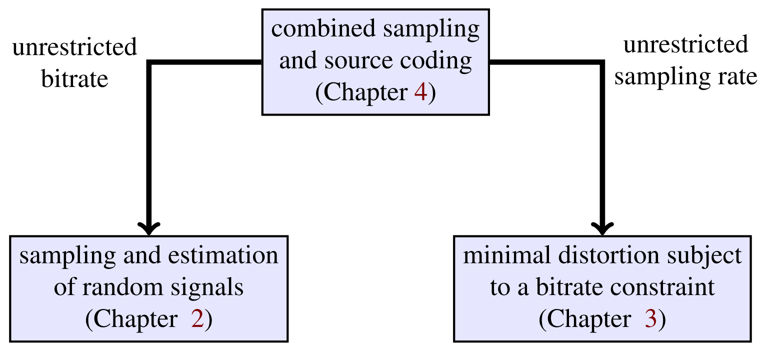

Figure 1.4: Relation between the background chapters and the main contribution (Chapter 4).

a sequence of actions taken by the agent controlling this asset. Assuming that the set of possible actions is finite, this last limitation corresponds to a quantization constraint. Therefore, the MMSE in estimating the continuous-time price based on the sequence of actions is described by the minimal distortion in the ADX. While in this case we have no control on the way the samples are encoded (into actions), the minimal distortion in the ADX setting provides a lower bound on the distortion in estimating the continuous-time price. This distortion can be expressed by an additional noise in a model that makes decisions based on the estimated price.

1.4 Organization and Contributions of the Thesis

The problem of characterizing the minimal distortion in ADX combines two classical problems in sampling theory and source coding theory. Hence, as preparation for studying this joint problem, we will consider the sampling part and the quantization part separately in Chapters 2 and 3, respectively. As illustrated in Figure 1.4, each of these problems corresponds to a special case of the ADX setting, obtained by relaxing the bitrate constraint or the sampling rate constraint. The full ADX setting is considered in Chapter 4, which comprises the main contribution of this thesis. The critical sampling rate under a bitrate constraint follows from this characterization as discussed in Chapter 5. Finally, concluding remarks and future research prospects are provided in Chapter 6. Detailed descriptions of the contribution of each chapter are given below.

1.4.1 Chapter 2

In Chapter 2 we focus on the sampling side of the ADX, as illustrated in Figure 1.4. We first provide a theoretical framework for sampling of random signals which we denote as bounded linear sampling. There are two main differences between our framework and standard frameworks in the sampling literature [14]: (1) we do not assume that the signals are bandlimited, and (2) we consider sampling operations that can be restricted to finite time horizons. The last property is required in order to consider the combined effect of sampling and encoding as in Figure 1.2, since encoding devices are usually restricted by their blocklength. Next, we consider the problem of estimating a Gaussian and stationary signal from its noisy samples under an MSE criterion. We focus in particular on the important class of linear time-invariant uniform samplers, which consists of samplers involving linear time-invariant operations followed by pointwise evaluations. The additional structure of samplers in this class allows us to derive closed-form expressions for this MSE. Finally, we consider the optimization over the sampling structure to minimize this MSE expression. This minimization is then used to recover classical results in sampling theory, showing that the conditions derived by Shannon and Landau for distortion-free sampling are obtained as special cases from the ADX setting, as indicated by the curved dashed line in Figure 1.3.

1.4.2 Chapter 3

In Chapter 3 we consider the general problem of encoding a random signal subject to a bitrate constraint. In the classical setting of Shannon, the encoder has full access to the analog signal it is required to encode, and the optimal tradeoff between bitrate and distortion is given by the DRF of the signal. In the ADX setting, however, the encoder can only observe a sampled and possibly noisy version of the signal, hence the resulting problem of encoding the samples with the goal of minimizing the MSE is denoted as *indirect* source coding.

In this chapter we review classical results in standard and indirect source coding. In particular, we show that the indirect distortion-rate function under MSE distortion is given as a sum of two terms: (1) the minimal MSE estimation error of the original source signal before encoding, and (2) the standard DRF of the optimal MSE estimator. Moreover, the optimal encoding procedure can be separated into a minimal MSE estimate followed by the encoding of this estimate. The aforementioned decomposition has been used before in various settings [15, 16, 6]. Our main new contribution is proving a new version of it that holds under time-varying operations that are commonly employed by sampling systems and, in particular, those defined in Chapter 2. We also demonstrate the usefulness of this decomposition to characterize the minimal distortion in various indirect source coding settings, as these characterizations will be used in subsequent chapters.

1.4.3 Chapter 4

In Chapter 4 we characterize the minimal MSE distortion in the ADX setting. This characterization is achieved by the following two steps:

(i) Given the output of the sampler, derive the optimal way to encode these samples subject to the bitrate R, so as to minimize the MSE distortion in reconstructing the original continuous-time signal.

(ii) Derive the optimal sampling scheme that minimizes the MSE in (i) subject to the sampling rate constraint.

In order to achieve (i), we formalize the ADX setting as a combined problem of sampling and source coding. We first derive a solution to this problem for the special case of time-invariant uniform samplers. This derivation is achieved by exploiting the special asymptotic properties of signals sampled by these samplers, in combination with the general decomposition of indirect DRFs derived in Chapter 3. Moreover, for this class of samplers, the optimization over the sampling scheme in (ii) coincides with the analogous optimization carried out in Chapter 2 without the bitrate constraint. The distortion function resulting from (i) and (ii) provides an achievable lower bound for the distortion in the ADX with the class of time-invariant uniform samplers. Finally, this lower bound is extended to the entire class of bounded linear samplers.

1.4.4 Chapter 5

In Chapter 5 we consider an interesting phenomenon that follows from our characterization of the minimal distortion in ADX: for most signal models, the minimal distortion subject to a bitrate constraint, described by the DRF of the signal, is attained by sampling at a rate lower than the Nyquist rate. Namely, for every bitrate R, there exists a sampling rate f_R that is usually smaller than the Nyquist rate, such that the optimal performance subject to a bitrate constraint is attained by sampling at this rate. Moreover, the new critical sampling rate f_R is finite even if the bandwidth of the signal is not.

Next, we provide interpretations of this phenomenon both in sampling theory and source coding theory. From a sampling theory perspective, the sampling rate f_R extends the classical Nyquist criterion in two aspects: (1) it describes the minimal sampling rate required to attain the minimal distortion subject to a bitrate constraint and (2) is also valid for signals that are not bandlimited. From a source coding theory perspective, the existence of a finite critical sampling rate for non-bandlimited signals provides conditions for encoding these signals with vanishing distortions. That is, for these signals, vanishing distortion in encoding is attained only in the asymptotic limit of infinite sampling rate and infinite bitrate. Since for each bitrate R the optimal sampling rate is f_R, the asymptotic ratio R/f_R provides the optimal number of bits per sample in this encoding.

1.5 Previously Published Material

Parts of this dissertation have appeared in the following publications:

- [17]: A. Kipnis, A. J. Goldsmith, Y. C. Eldar and T. Wiessman, “Distortion-Rate Function of Sub-Nyquist Sampled Gaussian Sources”, *IEEE Transactions on Information Theory*, vol. 62, no. 1, pp. 401–429, Jan 2016. Shorter version [18] appeared in *51st Annual Allerton Conference on Communication, Control, and Computing (Allerton)* IEEE, 2013, pp. 901-908.

- [19]: A. Kipnis, A. J. Goldsmith and Y. C. Eldar, “Rate-Distortion Function of Cyclostationary Gaussian Processes”, available https://arxiv.org/abs/1505.05586, to appear in *IEEE Transactions on Information Theory*. Shorter version [20] appeared in *Proceedings of the IEEE International Symposium on Information Theory*, pp. 2834-2838, 2014.

- [21]: A. Kipnis, Y. C. Eldar and A. J. Goldsmith, “Fundamental Distortion Limits of Analog-to-Digital

Compression", available https://arxiv.org/abs/1601.06421, to appear in *IEEE Transactions on Information Theory*.
Shorter versions appeared in:

  - [22]: –, "Gaussian Distortion-Rate Function under sub-Nyquist Nonuniform Sampling", *52nd Annual Allerton Conference on Communication, Control, and Computing (Allerton)* IEEE, 2014, pp. 874–880.
  - [23]: –, "Sub-Nyquist Sampling Achieves Optimal Rate-Distortion", *Information Theory Workshop (ITW)*, IEEE, 2015, pp. 1-5
  - [24]: –, "Optimal Trade-off Between Sampling Rate and Quantization Precision in A/D conversion", *54nd Annual Allerton Conference on Communication, Control, and Computing (Allerton)* IEEE, 2015, pp. 627-631.

- [25] A. Kipnis, A. J. Goldsmith and Y. C. Eldar, "The Distortion-Rate Function of Sampled Wiener Processes", available online, to appear in *IEEE Transactions on Information Theory*. Shorter version appeared in:

  - [26]: –"Information rates of sampled Wiener processes", *Proceedings of the IEEE International Symposium on Information Theory*, pp. 740-744, 2016.

Chapter 2

Sampling and Estimation of Random Signals

In this chapter we focus on the sampling part in the ADX setting of Figure 1.2. The main problem we consider is the estimation of a Gaussian stationary process from a discrete-time process representing the samples of the analog signal. A common assumption in sampling theory is that the analog process is bandlimited. This assumption is rooted in the celebrated Whittaker-Kotelnikov-Shannon (WKS) sampling theorem [27, 28, 2], which asserts that a bandlimited signal can be precisely represented by its uniform samples, as long as the sampling rate is larger than twice its bandwidth. Here the bandwidth is defined as the largest frequency in the Fourier transform of the signal, and the critical sampling rate, which is twice this frequency, was named the *Nyquist rate* by Shannon [29]. When the spectrum of the signal is a union of multiple disjoint intervals, sampling at the Nyquist rate is sufficient but not necessary to distortion-free reconstruction of the signal from its samples. An important extension of the WKS sampling theorem due to Landau [30, 3] asserts that reconstruction of a signal with known spectral support from its pointwise sampling is possible if and only if the sampling rate exceeds the signal's *spectral occupancy*. Here the spectral occupancy is the Lebesgue measure of the support of the spectrum of the signal, which is referred to as the *Landau rate* of the signal. Since the spectrum of real signals is symmetric, their Landau and Nyquist rates coincide whenever their spectrum is supported over a single interval [1]. We note that although Landau's result did not provide a recipe for sampling at the Landau rate, such procedures were later developed for many classes of sampling techniques, including bandpass sampling [31], multicoset sampling [32], and general nonuniform sampling [33].

The Nyquist and the Landau rates are central notions in sampling theory since they lead to necessary and sufficient conditions for optimal sampling, i.e., conditions for perfect representation of analog signals using their discrete-time samples. In many applications, however, it is necessary to sample signals below

[1] For this reason Landau referred to the spectral occupancy as the Nyquist rate in [3], although the former is clearly a generalization of the latter.

their Nyquist or Landau rates. For example, this need arises in sampling signals whose spectral support is larger than sampling rates of available hardware [34], when energy constraints preclude using Nyquist-rate sampling, or when fluctuations in the signal occur in any time resolution, such as in financial markets [12]. In these applications it is impossible to recover the original analog signal without error unless there exists some prior knowledge on the signal (aside from a known spectral distribution), such as sparsity [35, 36] or sparse (but otherwise unknown) spectral support [14]. It is important to stress that we will not pursue sampling under such priors in this thesis, although the ADX setting with a sparse signal prior was recently considered in [37]. Therefore, since in absence of such prior information, sampling below the Nyquist/Landau rates introduces non-zero reconstruction error, our goal in this chapter is to minimize this error and express it in a closed form.

In order to obtain a closed-form expression for the reconstruction distortion under sub Nyquist/Landau sampling, a statistical model for the analog signal being reconstructed is needed. The model of a Gaussian stationary process is chosen due to its usefulness in system theory [38], its maximal entropy and distortion properties [6], and the ability to derive error expressions in this model under a mean square error (MSE) criterion in closed forms. A review of the standard definitions for these processes is given in Section 2.1. Roughly speaking, the classical sampling theory, including WKS and Landua results, can be extended to the case of sampling Gaussian stationary signals using an isomorphism between two Hilbert spaces associated with the spectrum of the signal [39]. This connection is stressed in Section 2.2, where a framework for sampling stationary random signals under general conditions is developed. The most general form of sampling we consider is characterized by a linear pre-sampling operation and a discrete sampling set, and is denoted as *bounded linear sampling*. Later, in Section 2.3, we focus on a sub-class of bounded linear samplers in which the pre-sampling operation is time-invariant and the sampling set is uniform. As is shown is Section 2.4, the minimal MSE (MMSE) in recovering the original continuous-time analog signal from the output of samplers in this sub-class can be derived in a closed form. In addition to sampling, the estimation problem in Section 2.4 also assumes that the analog signal is corrupted by noise prior to sampling. Next, in Section 2.5, we consider the optimization of the linear pre-sampling operation so as to minimize the MSE under sub-optimal sampling. As we shall see, with enough uniform sampling branches, a particular choice of the pre-sampling operations and zero noise, a Gaussian stationary signal can be recovered with zero MSE by sampling at its Landau rate. For uniform sampling rates lower than the Landau rate and with non-zero noise, our characterization of the MSE in uniform sampling under an optimal pre-sampling operation leads to an achievable lower bound. Evidently, since the ADX adds another layer of lossy compression, this lower bound also bounds from below the distortion in our general ADX setting.

2.1 Notation

We now briefly describe the main notation involving continuous and discrete time random signals used throughout the thesis.

Throughout this thesis, the term *analog signal* refers to a continuous-time and continuous-amplitude stochastic process. Such a process is a set $X(\cdot) = \{X(t), t \in \mathbb{R}\}$ of real random variables indexed by the reals. For a time interval $I \subset \mathbb{R}$, we denote by $X_I(\cdot)$ the *restriction* of the process $X(\cdot)$ to the index set I. For either a random variable X or a stochastic process $X(\cdot)$, we use the lower case letters x and $x(\cdot)$ to denote a particular realization of the process. The lower case notation $x(\cdot)$ is also used to denote a deterministic analog signal.

The *probability law*, or simply the *law*, of the process $X(\cdot)$ is defined by the collection of all joint probability distribution functions of any finite number of elements from the process. Namely the law of $X(\cdot)$ is defined by the collection of all functions of the form:

$$F_{X(t_1),\ldots,X(t_n)}(\alpha_1,\ldots,\alpha_n) = \mathbb{P}\left(X(t_1) \le \alpha_1,\ldots,X(t_n) \le \alpha_n\right),$$

for $t_1,\ldots,t_n,\alpha_1,\ldots,\alpha_n \in \mathbb{R}$. We assume throughout that $X(\cdot)$ is *zero mean* in the sense that $\mathbb{E}X(t) = 0$ for all $t \in \mathbb{R}$.

The zero-mean process $X(\cdot)$ is said to be *Gaussian* if for any $(t_1,\ldots,t_n) \in \mathbb{R}^n$ there exists a positive semi-definite matrix $\Sigma \in \mathbb{R}^{n\times n}$ such that for any n dimensional vector $\mathbf{a} = (a_1,\ldots,a_n) \in \mathbb{R}^n$,

$$\mathbb{E}\exp\left\{ i\sum_{j=1}^{n} a_j X_j \right\} = \exp\left\{ -\frac{1}{2}\mathbf{a}^*\Sigma\mathbf{a} \right\}, \tag{2.1}$$

where $*$ denotes complex conjugation (or simply the transpose of a real vector). A zero mean process $X(\cdot)$ is said to be (wide-sense) *stationary* if for any $t,s \in \mathbb{R}$, the covariance function of $X(\cdot)$, defined as

$$C_X(t,s)\mathbb{E}\left[X(t)X(s)\right],$$

is only a function of $|t-s|$. For such a process, we set $C_X(\tau) = C_X(\tau,0)$. Since we only consider Gaussian random processes, stationarity in this wide sense implies stationarity in the strict sense of the joint law. The Wiener-Khinchin theorem implies that there exists a measure $dS_X(f)$ on $\mathbb{R}$, denoted as the spectral measure of $X(\cdot)$, such that

$$C_X(\tau) = \int_{-\infty}^{\infty} e^{2\pi i\tau f} dF(f).$$

In this thesis we assume that the spectral measure of $X(\cdot)$ is absolutely continuous with respect to the Lebesgue measure, so that there exists a function $S_X(f)$ such that $dS_X(f) = S_X(f)df$. $S_X(f)$ is denoted as the *power spectral density* of $X(\cdot)$.

An alternative way to describe a zero mean real Gaussian stationary process with absolutely continuous spectral measure is by defining its spectral properties first [40]. Specifically, we start with an $\mathbf{L}_1(\mathbb{R})$ non-negative symmetric function $S_X(f)$ which we call the PSD. We then define

$$C_X(\tau) = \int_{-\infty}^{\infty} e^{2\pi i \tau f} S_X(f) df.$$

The law of $X(\cdot)$ is now determined by its finite dimensional joint distributions through (2.1).

The discrete-time counterpart of the above notation with respect to discrete-time process $Y[\cdot] = \{Y[n], n \in \mathbb{Z}\}$ in defined in the obvious way. For example, the PSD of a stationary discrete-time process $Y[\cdot]$ is a non-negative function $S_Y\left(e^{2\pi i\phi}\right)$ symmetric around zero, periodic with period 1 and integrable over any period. This function defines a covariance function by

$$X_Y[k] = \mathbb{E}\left[X(n+k]X[n]\right] = \int_{-\frac{1}{2}}^{\frac{1}{2}} S_Y\left(e^{2\pi i\phi}\right) e^{2\pi i n\phi} d\phi.$$

2.2 Bounded Linear Sampling

Generally speaking, the sampling of a deterministic analog signal can be seen as a mapping from the space of signals to a real sequence. In particular, each sample can be seen as a mapping of the signal to a real number, and therefore defines a functional on the set of allowable signals. In addition, physical considerations of realizable systems dictate that this functional should be linear and bounded. In other words, each sample is obtained by an element of the dual space of the space of allowable signals. It is common to choose the space of signals such that the pointwise evaluation functional, i.e., the Dirac distribution, is continuous. When only bandlimited signals are considered, pointwise evaluations can be obtained, for example, by inner product in $\mathbf{L}_2$ with respect to the sinc function [14]. In our setting, however, we do not wish to restrict ourselves only to bandlimited signals. Instead, we assume that the space of signals $\mathscr{X}$ is a *nuclear* space of functions on $\mathbb{R}$ with topological dual $\mathscr{X}^\star$ [41]. Nuclear spaces are topological vector spaces with the additional property that for any bounded linear operator $H : \mathscr{X} \to \mathscr{X}^\star$ there exists a kernel K_A on $\mathbb{R}\times\mathbb{R}$ such that [42, Ch. 1]

$$Hx(t) = \int_{-\infty}^{\infty} K_A(\tau,t)dt, \quad t \in \mathbb{R}. \tag{2.2}$$

The pair $\mathscr{X}$ and $\mathscr{X}^\star$ are usually referred to as spaces of *test functions* and *distributions*, respectively, where in addition it is assumed that each element of any of these spaces has a well-defined Fourier transform. Hence, the bilinear operation between $x(\cdot) \in \mathscr{X}$ and $\phi \in \mathscr{X}^\star$ satisfies the Parseval identity:

$$\langle x(\cdot), \phi\rangle = \int_{-\infty}^{\infty} x(t)\phi^*(t)dt = \int_{-\infty}^{\infty} \mathscr{F}x(f)\left(\mathscr{F}\phi\right)^*(f)df, \tag{2.3}$$

where $\mathscr{F}$ is the Fourier transform and complex conjugation is represented by $*$. In (2.3) and henceforth, we use the integral notation to denote the bilinear operation between the set of signals and its dual space

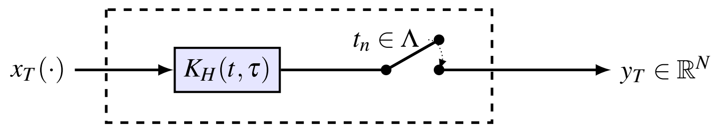

Figure 2.1: General structure of a bounded linear sampler.

(although standard integration may not be well-defined in some cases). We moreover assume that the Fourier transform of any $x(\cdot) \in \mathscr{X}$ is an $\mathbf{L}_1(\mathbb{R})$ function, which is a sufficient and necessary condition for $\mathscr{X}^\star$ to include the Dirac distribution δ_0.

For the aforementioned class of signals, a general *bounded linear sampler* is defined in terms of a bounded linear operator $H : \mathscr{X} \to \mathscr{X}$ with kernel $K_H(t,s)$, and a discrete ordered set $\Lambda \subset \mathbb{R}$ of sampling points, as illustrated in Figure 2.1. Note that here H is an operator from $\mathscr{X}$ to itself and not to its dual. Since $\mathscr{X} \subset \mathscr{X}^\star$, H still has the kernel representation (2.2). The set Λ is assumed to be *uniformly discrete* in the sense that there exists $\varepsilon > 0$ such that $|t_n - t_k| > \varepsilon$ for any $t_k, t_n \in \Lambda$. For an input signal $x(\cdot) \in \mathscr{X}$, the nth samples is given by

$$y[n] = Hx(t_n) = \int_{-\infty}^{\infty} K_H(t_n,s)x(s)ds. \tag{2.4}$$

We denote a sampler of this form as $S(H,\Lambda)$ and denote by $y(x)[\cdot] = S(H,\Lambda)(x)$ the discrete sequence defined by (2.4). The set Λ is referred to as the *sampling set*, and the operator H is referred to as the *pre-sampling operation*. The separation of the sampler into these two independent operations is in accordance with practical considerations in implementing sampling systems using linear components and sample-and-hold devices [14].

2.2.1 Finite time horizon sampling

In order to consider the concatenation of sampling and encoding as in the ADX setting, it is required to consider sampling over a finite time horizon $T > 0$. For this reason, for a sampler $S(H,\Lambda)$, we define $\Lambda_T = \Lambda \cap [-T/2, T/2]$ and

$$y_T(x) = S(H,\Lambda_T)(x)$$

as the finite vector of samples (y_T is finite since Λ is uniformly discrete) the sampler produces at time T. In addition to restricting the number of samples over the interval $[-T/2,T/2]$, it is sometimes useful to restrict the samples to be a function of inputs only over $[-T/2,T/2]$. In order to consider this restriction, we denote by H_T the restriction of H to the interval $[-T/2,T/2]$ which is defined by the kernel $K_{H_T} = \mathbf{1}_[-T/2,T/2]K_H(t,s)$, where $\mathbf{1}_A$ is the indicator function of a set A.

Example 1 *The most widely used example for a bounded linear sampler is uniform sampling. For this* *sampler, the pre-sampling operation is given in terms of the kernel $K(t,s) = \delta(t-s)$, where δ is the Dirac delta distribution at zero, and the sampling set $\Lambda = \mathbb{Z}/f_s$. Consequently, the nth sample equals*

$$y[n] = \int_{-\infty}^{\infty} \delta(n/f_s - s)x(s)ds, \quad = x(n/f_s), \quad n \in \mathbb{Z}.$$

The finite horizon version of this sampler produces the vector

$$y_T = \{x(n/f_s), n/f_s \in [-T/2, T/2]\}.$$

2.2.2 Sampling stationary random signals

In most of this thesis we are interested in the case where the signal $x(\cdot)$ is a realization of a zero-mean Gaussian stationary random process $X(\cdot)$. We assume that the part of the spectral measure of $X(\cdot)$ that is singular with respect to the Lebesgue measure is zero, so that its distribution is fully characterized by its power spectral density (PSD) function $S_X(f)$, defined by

$$\mathbb{E}[X(t)X(s)] = \mathbb{E}[X(t-s)X(0)] = \int_{-\infty}^{\infty} e^{2\pi i(t-s)f} S_X(f)df. \tag{2.5}$$

Equation (2.5) defines an isomorphism $X(t) \leftrightarrow e^{2\pi ift}$ between the Hilbert space generated by the closed linear span of the random process $X(\cdot) = \{X(t), t \in \mathbb{R}\}$ with norm $\|X(t)\|^2 = \mathbb{E}[X^2(t)]$ and the Hilbert space $\mathbf{L}_2(S_X)$ of complex valued functions generated by the closed linear span (CLS) of the exponentials $\mathscr{E} = \left\{e^{2\pi ift}, t \in \mathbb{R}\right\}$ with an $\mathbf{L}_2$ norm weighted by $S_X(f)$ [39]. This isomorphism allows us to define sampling of the random signal $X(\cdot)$ by describing its operation on the exponentials $\mathscr{E}$. Specifically, let $S(H,\Lambda)$ be a bounded linear sampler on the CLS of $\mathscr{E}$ and denote

$$\widehat{\phi}_n(f) = \int_{-\infty}^{\infty} e^{2\pi ifs} K_H(t_n, s)ds. \tag{2.6}$$

The nth sample of the sampler is given by the inverse map of $\widehat{\phi}_n$ under the isomorphism defined by (2.5). For example, when $K_H(t,s) = \delta_0(t-s)$, then $\widehat{\phi}_n = e^{2\pi if(n/f_s)}$, whose inverse image under (2.5) equals $X(n/f_s)$.

We note that when $X(\cdot)$ is bandlimited in the sense that $S_X(f)$ is supported within an interval $(-f_B, f_B)$, then the CLS of the exponentials $\mathscr{E}$ with an $\mathbf{L}_2$ norm weighted by $S_X(f)$ is isomorphic, through the Fourier transform operator on $\mathbf{L}_2$, to the Paley-Wiener space $\mathrm{PW}(-f_B, f_B)$ of $\mathbf{L}_2$ functions whose Fourier transform is supported on $(-f_B, f_B)$ [43]. Therefore, our definition of a bounded linear sampler for $S_X(f)$ reduces to the standard setting of sampling in the Paley-Wiener space $\mathrm{PW}(-f_B, f_B)$ of $\mathbf{L}_2$ functions whose Fourier transform is supported on $(-f_B, f_B)$, as considered for example by Landau [30].

So far we defined a framework for bounded linear sampling of stationary random processes whose PSD is any $\mathbf{L}_1(\mathbb{R})$ function. In the next section we consider two specific sub-classes of bounded linear samplers.

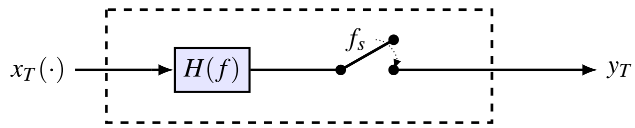

Figure 2.2: Single branch uniform sampler.

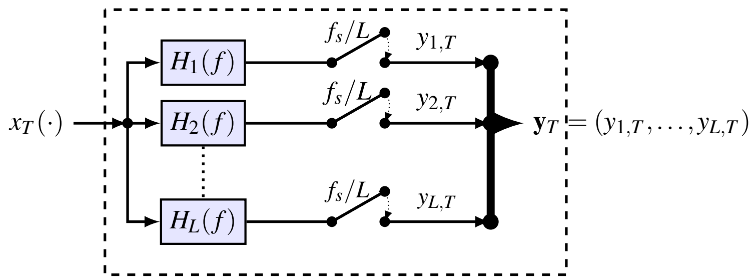

Figure 2.3: Multi-branch uniform sampler.

2.3 Time-Invariant Uniform Sampling

In this section we focus on two sub-classes from the class of bounded linear samplers defined in Section 2.2. The first sub-class consists of a sampler with a time-invariant pre-sampling operation and a uniform sampling set $\Lambda = \mathbb{Z}/f_s$, where f_s is denoted the *sampling rate*. For simplicity, we refer to each sampler in this class as a single-branch (SB) uniform sampler. The second class of samplers is called multi-branch (MB) uniform samplers, where each sampler in this class consists of multiple SB samplers. Our focus on SB and MB samplers is motivated by their simple structure, which leads to a relatively simple hardware implementation [14]. In the following sections we describe the sampling structures in more detail.

2.3.1 Single-branch uniform sampling

The general structure of a SB uniform sampler is illustrated in Figure 2.2. The pre-sampling operation H is a linear time-invariant system with frequency response $H(f)$ and impulse response $h(t)$. Namely, $K_H(t,s) = h(t-s)$ and $H(f) = \mathscr{F}h$. The sampling set of this filter is the grid $\mathbb{Z}/f_s$, so that the nth sample of the input $x(\cdot)$ is given by

$$y[n] = \int_{-\infty}^{\infty} h(n/f_s - s)x(s)dt.$$

For example, the uniform sampler of Example 1 is a special case of this sampler with $h(t) = \delta(t)$, i.e., the pre-sampling operator H is the identity.

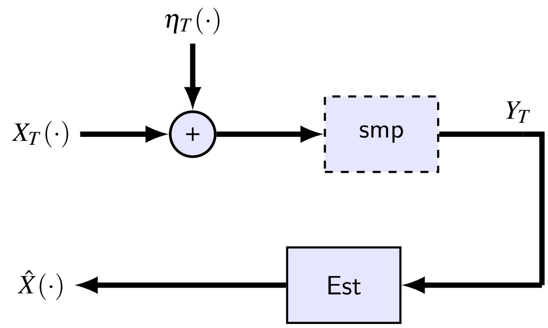

Figure 2.4: System model for MMSE estimation from noisy samples.

2.3.2 Multi-branch uniform sampling

We now generalize the SB uniform sampler by considering an array of such samplers, resulting in the MB uniform sampler illustrated in Figure 2.3. This sampler is characterized by a set of L linear time-invariant systems $H_1,\dots,H_L$ and a single sampling rate f_s, such that the output of each sampling branch is given by

$$y_l[n] = \int_{-\infty}^{\infty} x(s)h_l(nL/f_s - s)ds.$$

Note that the sampling rate of each branch is f_s/L, so that the overall sampling rate of the system is f_s.

We denote by $y_l[\cdot]$ the discrete-time signal at the output of the lth sampling branch. Henceforth, we treat the L discrete-time signals $Y_1[\cdot],\dots,Y_L[\cdot]$ resulting from the L branches, respectively, as a single vector-valued signal $\mathbf{y}[\cdot]$, defined by

$$\mathbf{y}[n] = (y_1[n],\dots,y_L[n]), \quad n \in \mathbb{Z}.$$

In order to show that a MB sampler is a special case of the bounded linear sampler illustrated in Figure 2.1, we consider the sampling set $\Lambda = \mathbb{Z}/f_s$ and define the operator H by the kernel

$$K_H(t,s) = \sum_{l=1}^{L} h_l(t-s)\mathbf{1}_{[(l-1)/f_s, l/f_s)}(t).$$

2.4 Optimal Estimation from Noisy Samples

In this section we consider the minimal MSE in estimating a continuous-time Gaussian stationary process $X(\cdot)$ from its samples obtained using a bounded linear sampler. We focus in particular on the time-invariant uniform samplers of Section 2.3, although we also comment on the conditions for zero MSE under any bounded linear sampler. In addition, in order to incorporate the most general conditions that arise in real systems, we assume that the signal $X(\cdot)$ is corrupted by an additive Gaussian stationary noise process $\eta(\cdot)$ prior to sampling, as illustrated in Figure 2.4. Special cases of this settings were considered in [44], [45, Prop. 3], [46] and [47]. Our setting generalizes these works since we allow for finite time-horizon sampling that is required in order to consider coding based on the samples in subsequent chapters.

In the setting of Figure 2.4, the noise $\eta(\cdot)$ and $X(\cdot)$ are independent and therefore their sum is a Gaussian stationary process with PSD $S_X(f)+S_\eta(f)$. Therefore, we sometimes abbreviate the notation by denoting $S_{X+\eta}(f)=S_X(f)+S_\eta(f)$. For a finite time horizon T, a bounded linear sampler $S=S(H,\Lambda)$ receives a realization of $X(\cdot)+\eta(\cdot)$ over $[-T/2,T/2]$ and produces the finite vector Y_T of dimension $N=|\Lambda\cap[-T/2,T/2]|$. The estimator produces an estimate of $X_T(\cdot)$, which, since we use the MSE criterion, can be assumed to be the conditional expectation of $X_T(\cdot)$ given Y_T. Namely, denote

$$\widetilde{X}_T(t)=\mathbb{E}\left[X(t)|Y_T\right]. \tag{2.7}$$

The finite horizon MMSE is defined as

$$\mathsf{mmse}(X_T|Y_T)=\frac{1}{T}\int_{-T/2}^{T/2}\mathbb{E}\left(X(t)-\widetilde{X}(t)_T\right)^2 dt. \tag{2.8}$$

Consequently, the asymptotic non-causal MMSE of a the bounded linear sampler S is defined as

$$\mathsf{mmse}_S=\liminf_{T\to\infty}\mathsf{mmse}(X_T|Y_T). \tag{2.9}$$

Note that since $X(\cdot)$ and $X(\cdot)+\eta(\cdot)$ are jointly Gaussian processes, for any bounded linear sampler $S(H,\Lambda)$ we have that

$$\mathsf{mmse}_{S(H,\Lambda)}\geq\mathsf{mmse}(X|X+\eta),$$

where $\mathsf{mmse}(X|X+\eta)$ is the error in the Wiener filter for estimating $X(\cdot)$ from $\eta(\cdot)$ [48], given by

$$\mathsf{mmse}(X|X+\eta)=\sigma_X^2-\int_{-\infty}^{\infty}\frac{S_X^2(f)}{S_{X+\eta}(f)}df=\int_{-\infty}^{\infty}\frac{S_X(f)S_\eta(f)}{S_{X+\eta}(f)}df. \tag{2.10}$$

In addition, for any time horizon T and time instances $t_1,\ldots,t_n\in[-T/2,T/2]$, the vector of samples Y_T concatenated with the vector $\mathbf{X}=(X(t_1),\ldots,X(t_n))$ is Gaussian. As a result, the conditional distribution of $\mathbf{X}$ given Y_T is Gaussian, and we have

$$\widetilde{X}_T(t_i)=\mathbb{E}\left[X(t_i)Y_T^*\right]\Sigma_{Y_T}^{-1}Y_T, \tag{2.11}$$

where $\Sigma_{Y_T}=\mathbb{E}\left[Y_TY_T^*\right]$ is the covariance matrix of the vector Y_T. If the covariance matrix of the vector $\left(\widetilde{X}(t_1),\ldots,\widetilde{X}(t_n)\right)$ converges to a constant, then we say that the law of $\widetilde{X}_T(\cdot)$ converges to an asymptotic Gaussian law. This asymptotic law defines a Gaussian process $\widetilde{X}(\cdot)$ by the joint probability distribution of

$\widetilde{X}(t_1),\dots,X(t_n)$ for any $t_1,\dots,t_n \in \mathbb{R}$.

In what follows we derive mmse_S in closed form when S is a SB or a MB uniform sampler. In particular, we will see that under these classes of samplers, the law of the estimator $\widetilde{X}_T(\cdot)$ converges to the law of a Gaussian process $\widetilde{X}(\cdot)$ with block-stationary or cyclostationary structure [49].

2.4.1 Single-branch uniform sampling

Consider the case where Y_T is the result of sampling the process $X(\cdot)+\eta(\cdot)$ under the SB uniform sampler of Figure 2.2 at time T. The covariance of $X(t)$ and the nth sample $Y_T[n]$ is given by

$$\mathbb{E}\left[X(t)Y_T[n]\right] = \int_{-T/2}^{T/2} \mathbb{E}\left[X(t)\left(X(s)+\eta(s)\right)\right] h(n/f_s-s)ds = \int_{-T/2}^{T/2} C_X(t-s)h(n/f_s-s)ds. \tag{2.12}$$

From the Parseval identity (2.3) it follows that the limit as $T\to\infty$ of (2.12) exists and is given by

$$\int_{-\infty}^{\infty} H^*(f)S_X(f)e^{2\pi i(t-n/f_s)f}df.$$

Similarly, the covariance of $Y_T[n]$ and $Y_T[k]$ is given by

$$\mathbb{E}\left[Y_T[n]Y_T[k]\right] = \int_{-T/2}^{T/2}\int_{-T/2}^{T/2} h(n/f_s-s)h^*(k/f_s-t)C_X(s-t)dsdt.$$

By basic properties of uniform sampling of functions and their Fourier transform (see, e.g., [50]), the above expression converges to

$$\int_{-f_s/2}^{f_s/2} \sum_{m\in\mathbb{Z}} S_X\left(f-mf_s\right)\left|H(f-f_sm)\right|^2 e^{2\pi i(n-k)f/f_s\phi} = \int_{-\frac{1}{2}}^{\frac{1}{2}} f_s \sum_{m\in\mathbb{Z}} S_X\left(f_s(\phi-m)\right)\left|H\left(f_s(\phi-m)\right)\right|^2 e^{2\pi i(n-k)\phi}. \tag{2.13}$$

It follows from (2.11) that the law of $\widetilde{X}_T(\cdot)$ converges to an asymptotic law $\widetilde{X}(\cdot)$. It also follows from (2.13) that as $T\to\infty$, the joint distribution of the vector Y_T defines a discrete-time process $Y[\cdot]$ with covariance function (2.13). In particular, $Y[\cdot]$ is stationary and its PSD is given by

$$S_Y\left(e^{2\pi i\phi}\right) = f_s \sum_{m\in\mathbb{Z}} S_X\left(f_s(\phi-m)\right)\left|H\left(f_s(\phi-m)\right)\right|^2.$$

The process $\widetilde{X}(\cdot)$ can be seen as the asymptotic non-causal MMSE estimator of $X(\cdot)$ from $Y[\cdot]$. Since $X(\cdot)$ and $Y[\cdot]$ are jointly Gaussian, the process $\widetilde{X}(\cdot)$ is a linear combination of $Y[\cdot]$. However, due to the fact that the estimation is from discrete-time to continuous-time, $\widetilde{X}(\cdot)$ may not be stationary. To illustrate this last fact, consider the case of a uniform sampler with no pre-sampling operation or noise. In this situation it can be checked that $Y[n]=X(n/f_s)$, for all $n\in\mathbb{Z}$. Therefore, regardless of f_s and the bandwidth of $X(\cdot)$, the variance of $\widetilde{X}(t)$ for $t=n/f_s$ is zero, while it is usually not zero for $t\notin\mathbb{Z}/f_s$. Nevertheless, since $X(\cdot)$ is stationary, the instantaneous estimation error $X(t)-\tilde{X}(t)$ is periodic and uniformly bounded in t. As a result, the limit in (2.9) exists and can be written as

$$\begin{aligned}\mathsf{mmse}_{\mathsf{SB}(H)}(f_s) &= \lim_{T\to\infty}\frac{1}{2T}\int_{-T}^{T}\mathbb{E}\left[\left(X(t)-\tilde{X}(t)\right)^2\right]dt\\ &= \int_0^1 \lim_{N\to\infty}\frac{1}{2N+1}\sum_{n=-N}^{N}\mathbb{E}\left[\left(X\left(\frac{n+\Delta}{f_s}\right)-\tilde{X}\left(\frac{n+\Delta}{f_s}\right)\right)^2\right]d\Delta\\ &= \int_0^1 \mathsf{mmse}(X_\Delta|Y)d\Delta, \end{aligned} \tag{2.14}$$

where $X_\Delta[\cdot]$ is a discrete-time process defined by

$$X_\Delta[n] \triangleq X\left(\frac{n+\Delta}{f_s}\right), \quad n\in\mathbb{Z}, \tag{2.15}$$

and is known as the Δ-*polyphase component* of $X(\cdot)$ [51]. The function $\mathsf{mmse}(X_\Delta|Y)$ is given by

$$\mathsf{mmse}(X_\Delta|Y) \triangleq \lim_{N\to\infty}\frac{1}{2N+1}\sum_{n=-N}^{N}\mathbb{E}\left(X_\Delta[n]-\widetilde{X}_\Delta[n]\right)^2,$$

where

$$\tilde{X}_\Delta[n] \triangleq \mathbb{E}\left[X_\Delta[n]|Y[\cdot]\right], \quad n\in\mathbb{Z}$$

is the MMSE estimator of $X_\Delta[\cdot]$ given $Y[\cdot]$. Since $X_\Delta[\cdot]$ and $Y[\cdot]$ are jointly Gaussian and stationary, $\mathsf{mmse}_{X_\Delta|Y}$ can be evaluated using the non-causal Wiener filter. By deriving an expression for the MSE in this estimation and integrating over Δ, we conclude the following:

Proposition 2.1 *The MMSE in estimating $X(\cdot)$ from the output of a SB uniform sampler at sampling rate f_s and pre-sampling filter $H(f)$ is given by*

$$\mathsf{mmse}_{\mathsf{SB}(H)}(f_s) = \sigma_X^2 - \int_{-\frac{f_s}{2}}^{\frac{f_s}{2}} \widetilde{S}_{X|Y}(f)df, \tag{2.16}$$

where $\sigma_X^2 = \mathbb{E}(X(t))^2$ and

$$\widetilde{S}_{X|Y}(f) \triangleq \frac{\sum_{k\in\mathbb{Z}} S_X^2(f-f_sk)\,|H(f-f_sk)|^2}{\sum_{k\in\mathbb{Z}} S_{X+\eta}(f-f_sk)\,|H(f-f_sk)|^2}. \tag{2.17}$$

Remark 1 *In* (2.17) *and henceforth, we interpret fractions of this form as zero whenever both numerator and denominator are zero.*

The proof of Proposition 2.1 is omitted since it follows as a special case of Proposition 2.2 in the next subsection.

We note that since the denominator in (2.17) is periodic in f with period f_s, (2.16) can be written as

$$\begin{aligned}\mathsf{mmse}_{\mathsf{SB}(H)}(f_s) &= \sigma_X^2 - \int_{-\infty}^{\infty} \frac{S_X^2(f)|H(f)|^2}{\sum_{k\in\mathbb{Z}} S_{X+\eta}(f-f_sk)|H(f-f_sk)|^2} df \\ &= \int_{-\infty}^{\infty} S_X(f)\left(1-\frac{S_X(f)|H(f)|^2}{\sum_{k\in\mathbb{Z}} S_{X+\eta}(f-f_sk)|H(f-f_sk)|^2}\right) df. \end{aligned} \tag{2.18}$$

This shows that the expression for $\mathsf{mmse}_{\mathsf{SB}(H)}(f_s)$ in Proposition 2.1 is equivalent to [44, Eq. 10]. The alternate derivation of this expression given here using (2.14) provides a new interpretation of the function $\widetilde{S}_{X|Y}(f)$ as the average of spectral densities of estimators of the stationary polyphase components of $X(\cdot)$, namely

$$\widetilde{S}_{X|Y}(f) = \int_0^1 f_s S_{X_\Delta|Y}(f/f_s) d\Delta. \tag{2.19}$$

It is interesting to note that, although the estimator $\tilde{X}(\cdot)$ is not a stationary process (it is in fact a *cyclostationary* process with period f_s^{-1} [52]), it still has a structure similar to the Wiener filter, namely

$$\mathbb{E}[X(t)|Y[\cdot]] = \sum_{n\in\mathbb{Z}} Y[n]w(t-n/f_s), \tag{2.20}$$

where the frequency response of the analog filter $w(t)$ is given by [47, Eq. 1 and 2] and [44, Eq. 10]

$$W(f) = \frac{S_X(f)H^*(f)}{\sum_{k\in\mathbb{Z}} (S_{X+\eta}(f-f_sk))|H(f-f_sk)|^2} \tag{2.21}$$

Consider the special case where $f_s > f_{\mathsf{Nyq}}$, i.e. the support of $S_X(f)$ is contained within $(-f_2/2, f_s/2)$, and where $H(f)$ does not block any part of this band. For this case, the denominator in (2.18) is non-zero only for $k=0$ and $H(f)$ can be eliminated. In this situation, the integral in the RHS of (2.18) is only over $\frac{S_X^2(f)}{S_{X+\eta}(f)}$, and the entire expression coincides with the error under non-causal Wiener filtering. Indeed, under these conditions the estimator (2.21) coincides with the Wiener filter [53], and there is no loss of information due to sampling. If, in addition, η equals zero, then the MMSE vanishes completely and the RHS of (2.20) takes the form of sinc interpolation as in the stochastic version of the WKS sampling theorem [54, 55].

2.4.2 Multi-branch uniform sampling

We now consider the MMSE in estimating the Gaussian stationary process $X(\cdot)$ from a sampled version of its noisy samples using the MB uniform sampler of Figure 2.3. The counterpart of Proposition 2.1 for the MB uniform sampler is as follows:

Proposition 2.2 *The MMSE in estimating $X(\cdot)$ from the process samples $\mathbf{Y}[\cdot]$ at the output of a MB uniform* *sampler of rate f_s with pre-sampling filters $H_1(f),\ldots,H_L(f)$, is given by*

$$\mathsf{mmse}_{\mathsf{MB}(H_1,\ldots,H_L)}(f_s) = \sigma_X^2 - \int_{-\frac{f_s}{2}}^{\frac{f_s}{2}} \mathrm{Tr}\left(\widetilde{\mathbf{S}}_{X|\mathbf{Y}}(f)\right) df. \tag{2.22}$$

Here, $\widetilde{\mathbf{S}}_{X|\mathbf{Y}}(f)$ is the $L\times L$ matrix defined by

$$\widetilde{\mathbf{S}}_{X|\mathbf{Y}}(f) \triangleq \tilde{\mathbf{S}}_{\mathbf{Y}}^{-\frac{1}{2}*}(f)\mathbf{K}(f)\tilde{\mathbf{S}}_{\mathbf{Y}}^{-\frac{1}{2}}(f), \tag{2.23}$$

where the matrices $\tilde{\mathbf{S}}_{\mathbf{Y}}(f),\mathbf{K}(f)\in\mathbb{C}^{L\times L}$ are given by

$$\left(\tilde{\mathbf{S}}_{\mathbf{Y}}(f)\right)_{i,j} = \sum_{k\in\mathbb{Z}} S_{X+\eta}\left(f-f_sk\right)H_i^*\left(f-f_sk\right)H_j\left(f-f_sk\right)\left(f-f_sk\right),$$

and

$$\left(\mathbf{K}(f)\right)_{i,j} = \sum_{k\in\mathbb{Z}} S_X^2\left(f-f_sk\right)H_i^*\left(f-f_sk\right)H_j\left(f-f_sk\right).$$

Proof We first write $\mathsf{mmse}_{\mathsf{MB}(H_1,\ldots,H_L)}(f_s)$ as in (2.14) where the output of the sampler is the vector valued process $\mathbf{Y}[\cdot]$. As in the proof of Proposition 2.1, we derive (2.22) by first evaluating

$$\mathsf{mmse}(X_\Delta|\mathbf{Y}) = \sigma_X^2 - \int_{-\frac{1}{2}}^{\frac{1}{2}} \mathbf{S}_{X_\Delta|\mathbf{Y}}\left(e^{2\pi i\phi}\right)d\phi \tag{2.24}$$

and then computing the integral in (2.14). The joint PSD of $X_\Delta[\cdot]$ and $\mathbf{Y}[\cdot]$ is the $1\times L$ matrix

$$\mathbf{S}_{X_\Delta\mathbf{Y}}\left(e^{2\pi i\phi}\right) = \left(S_{X_\Delta Y_1}\left(e^{2\pi i\phi}\right),\ldots,S_{X_\Delta Y_L}\left(e^{2\pi i\phi}\right)\right).$$

Using the fact that $Y_l[n] = \left\{H\star(X+\eta)\right\}(n/f_s)$ and properties of multi-rate signal processing (e.g. [56]), we have

$$S_{X_\Delta Y_l}\left(e^{2\pi i\phi}\right) = \sum_{m\in\mathbb{Z}}\mathbb{E}\left[X\left(\frac{n+m+\Delta}{f_s}\right)Y_l[n]\right]e^{-2\pi im\phi} = \sum_{k\in\mathbb{Z}} S_X\left(f_s(\phi-k)\right)H_l^*\left(f_s(\phi-k)\right)e^{2\pi ik\Delta}.$$

In addition, the (l,p)th entry of the $L\times L$ matrix $\mathbf{S}_{\mathbf{Y}}\left(e^{2\pi i\phi}\right)$ is given by

$$\left\{\mathbf{S}_{\mathbf{Y}}\left(e^{2\pi i\phi}\right)\right\}_{l,p} = \sum_{k\in\mathbb{Z}}\left\{S_{X+\eta}H_l^*H_p\right\}\left(f_s(\phi-k)\right),$$

where we have used the shortened notation $\left\{S_1S_2\right\}(x)\triangleq S_1(x)S_2(x)$ for two functions S_1 and S_2 with the same domain. It follows that

$$\mathbf{S}_{X_\Delta|\mathbf{Y}}\left(e^{2\pi i\phi}\right) = \left\{\mathbf{S}_{X_\Delta\mathbf{Y}}\mathbf{S}_{\mathbf{Y}}^{-1}\mathbf{S}_{X_\Delta\mathbf{Y}}^*\right\}\left(e^{2\pi i\phi}\right)$$

can also be written as

$$\mathbf{S}_{X_\Delta|\mathbf{Y}}\left(e^{2\pi i\phi}\right) = \mathrm{Tr}\left\{\mathbf{S}_{\mathbf{Y}}^{-\frac{1}{2}*}\mathbf{S}_{X_\Delta\mathbf{Y}}^{*}\mathbf{S}_{X_\Delta\mathbf{Y}}\mathbf{S}_{\mathbf{Y}}^{-\frac{1}{2}}\right\}\left(e^{2\pi i\phi}\right), \tag{2.25}$$

where $\mathbf{S}_{\mathbf{Y}}^{-\frac{1}{2}*}\left(e^{2\pi i\phi}\right)$ is the $L\times L$ matrix satisfying $\mathbf{S}_{\mathbf{Y}}^{-\frac{1}{2}*}\left(e^{2\pi i\phi}\right)\mathbf{S}_{\mathbf{Y}}^{-\frac{1}{2}}\left(e^{2\pi i\phi}\right) = \mathbf{S}_{\mathbf{Y}}^{-1}\left(e^{2\pi i\phi}\right)$.

The $(l,p)^{\text{th}}$ entry of $\mathbf{S}_{X_\Delta\mathbf{Y}}^{*}\left(e^{2\pi i\phi}\right)\mathbf{S}_{X_\Delta\mathbf{Y}}\left(e^{2\pi i\phi}\right)$ is given by

$$\begin{aligned}\left\{\mathbf{S}_{X_\Delta\mathbf{Y}}^{*}\mathbf{S}_{X_\Delta\mathbf{Y}}\right\}_{l,p}\left(e^{2\pi i\phi}\right) &= \sum_{k\in\mathbb{Z}}\left\{S_X H_l^*\right\}\left(f_s(\phi-k)\right)e^{2\pi ik\Delta}\sum_{m\in\mathbb{Z}}\left\{S_X H_p\right\}\left(f_s(\phi-m)\right)e^{-2\pi il\Delta} \\ = \sum_{k,m\in\mathbb{Z}}&\left[\left\{S_X H_l^*\right\}\left(f_s(\phi-k)\right)\left\{S_X H_l\right\}\left(f_s(\phi-m)\right)e^{2\pi i\Delta(k-m)}\right],\end{aligned}$$

which leads to

$$\int_0^1\left\{\mathbf{S}_{X_\Delta\mathbf{Y}}^{*}\mathbf{S}_{X_\Delta\mathbf{Y}}\right\}_{l,p}\left(e^{2\pi i\phi}\right)d\Delta = \sum_{k\in\mathbb{Z}}\left\{S_X^2 H_l^* H_p\right\}\left(f_s(\phi-k)\right).$$

By integrating (2.25) with respect to Δ from 0 to 1 leads to

$$\mathrm{Tr}\left\{\mathbf{S}_{\mathbf{Y}}^{-\frac{1}{2}*}\bar{\mathbf{K}}\mathbf{S}_{\mathbf{Y}}^{-\frac{1}{2}}\right\}\left(e^{2\pi i\phi}\right),$$

where $\bar{\mathbf{K}}$ is the $L\times L$ matrix given by

$$\bar{\mathbf{K}}_{l,p}\left(e^{2\pi i\phi}\right) = \sum_{k\in\mathbb{Z}}\left\{S_X^2 H_l^* H_p\right\}\left(f_s(\phi-k)\right).$$

The proof is completed by changing the integration variable in (2.24) from ϕ to $f=\phi f_s$, so $\mathbf{S}_{\mathbf{Y}}\left(e^{2\pi i\phi}\right)$ and $\bar{\mathbf{K}}\left(e^{2\pi i\phi}\right)$ are replaced by $\tilde{\mathbf{S}}_{\mathbf{Y}}(f)$ and $\mathbf{K}(f)$, respectively. □

2.4.3 Sufficient condition for optimal reconstruction under pointwise sampling

Let $S(H,\Lambda)$ be any bounded linear sampler and assume that H is the identity operator $\mathbf{I}$, so that the samples are the pointwise evaluations of the process $X(\cdot)+\eta(\cdot)$. In this subsection we derive conditions on the sampling set Λ such that $\mathsf{mmse}_{S(\mathbf{I},\Lambda)}$ equals zero. For this purpose, we introduce the following definition:

Definition 1 *The* lower Beurling density *of the sampling set* Λ *is*

$$d^-(\Lambda) \triangleq \lim_{T\to\infty}\inf_{r\in\mathbb{R}}\frac{|\Lambda\cap[T,T+r]|}{T}. \tag{2.26}$$

Similarly, one can define the *upper* Beurling density by replacing in (2.26) the inf with sup. The set Λ is said to have Beurling density $d^-(\Lambda)$ whenever the upper and lower Beurling densities coincide.

Using the above definition, we conclude the following:

Proposition 2.3 *Consider the MMSE estimation problem with sampler $S(\mathbf{I},\Lambda)$ and zero noise $\eta(\cdot)=0$. Then* $\mathsf{mmse}_{S(\mathbf{I},\Lambda)}=0$ *only if*

$$d^{+}(\Lambda)\geq \mu(\text{supp } S_X)=f_{\mathsf{Lan}}. \tag{2.27}$$

Proof The condition $\mathsf{mmse}_{S(\mathbf{I},\Lambda)}=0$ is equivalent to the following: for any $\varepsilon>0$ and T_0, there exists $T>T_0$ such that

$$\frac{1}{T}\int_{-T/2}^{T/2}\mathbb{E}\left(X(t)-\widetilde{X}_T(t)\right)^2 dt<\varepsilon. \tag{2.28}$$

Write

$$\widetilde{X}_T(t)=\mathbb{E}[X(t)|Y_T]=\sum_{n=1}^{N}\alpha_n(t)Y_T[n],$$

with $N=|\Lambda\cap[-T/2,T/2]|$. Then from the isomorphism (2.5), (2.28) can be written as

$$\frac{1}{T}\int_{-T/2}^{T/2}\int_{-\infty}^{\infty}\left|e^{2\pi itf}-\sum_{n=1}^{N}\alpha_n(t)e^{2\pi it_nf}\right|^2 S_X(f)df dt<\varepsilon. \tag{2.29}$$

Condition (2.29) implies that the set $\mathscr{E}(\Lambda)=\left\{e^{2\pi it_nf},t_n\in\Lambda\right\}$ is a *frame* [57] in the Hilbert space $\mathscr{H}(\mathscr{E})$ generated by the exponentials $\mathscr{E}(\mathbb{R})=\left\{e^{2\pi itf},t\in\mathbb{R}\right\}$ with an $\mathbf{L}_2$ measure weighted by $S_X(f)$. Since the measure $S_X(f)df$ and the Lebesgue measure df with support in suppS_X are mutually absolutely continuous, $\mathscr{E}(\Lambda)$ is a Reisz basis in $\mathscr{H}(\mathscr{E})$) if and only if it is a Reisz basis in the Paley-Wiener space PW(supp S_X) of functions with Fourier transform supported in supp S_X. The condition for $\mathscr{E}(\Lambda)$ to be a basis to PW(supp S_X) follows from Landau's work [30], and implies that the upper Beurling density of Λ must exceed the spectral occupancy $\mu(\text{supp } S_X)$, which is now denoted as the Landau rate f_{Lan}. □

We note that Proposition 2.3 holds even if the asymptotic law of the estimator $\widetilde{X}_T(\cdot)$ of (2.11) as well as the limit of $\mathsf{mmse}(X_T|Y_T)$ do not exist. Indeed, general conditions for the existence of such an asymptotic law for an arbitrary sampling set Λ seem hard to derive.

2.5 Optimal Pre-Sampling Operations

So far we considered the pre-sampling operation in SB and MB uniform sampling as part of the MMSE estimation problem. We now consider this pre-sampling operation as part of the system design and seek for an operation that minimizes the MMSE. That is, we look for the pre-sampling filter $H^\star$ that minimizes (2.16) in SB uniform sampling and the set of filters, or *filter-bank*, $H_1^\star,\ldots,H_L^\star$, in MB uniform sampling.

For the case of SB sampling, we shall see below that the optimal pre-sampling filter is fully characterized by the following two properties:

(i) **Aliasing-free** - the passband of $H^\star$ is such that no aliasing is present in $Y[\cdot]$ by sampling at rate f_s, i.e., all integer shifts of the support of the filtered signal by f_s are disjoint.

(ii) **SNR maximization** - the passband of $S^\star$ is chosen to maximize the SNR at the output of the filter, subject to the aliasing-free property (i).

In the case where $S_X(f)$ is unimodal and the noise spectrum $S_\eta(f)$ is flat, an ideal low-pass filter with cut-off frequency $f_s/2$ satisfies both the aliasing free and the SNR maximization properties, and is therefore the optimal pre-sampling filter that minimizes $D_{\mathsf{SB}(H)}(f_s,R)$. In general, however, the set that maximizes the SNR is not aliasing-free, so that the constraint in (i) implies that the optimal filter does not necessarily maximize the SNR of the signal at its output. As will be explained in more detail below, increasing the number of sampling branches as in MB uniform sampling allows for a relaxation the aliasing free constraint. This in turn allows for retaining spectral bands of higher SNR, thus reducing the overall MMSE.

2.5.1 Optimal pre-sampling filter in single-branch uniform sampling

As is apparent from (2.16), the problem of minimizing (2.16) over H is equivalent to finding the filter that maximizes $\widetilde{S}_{X|Y}(f)$ for every frequency $f \in (-f_s/2, f_s/2)$ independently. Thus, we are looking to determine the filter $H^\star$ that achieves the supremum in

$$\begin{aligned}\widetilde{S}^\star_{X|Y}(f) &\triangleq \sup_H \widetilde{S}_{X|Y}(f) \\ &= \sup_H \frac{\sum_{k\in\mathbb{Z}} S_X^2(f-f_sk)\,|H(f-f_sk)|^2}{\sum_{k\in\mathbb{Z}} S_{X+\eta}(f-f_sk)\,|H(f-f_sk)|^2}\end{aligned} \tag{2.30}$$

in the domain $(-f_s/2, f_s/2)$. In what follows, we will describe $H^\star$ by defining a set of frequencies $F^\star$ of minimal Lebesgue measure such that

$$\int_{F^\star} \frac{S_X^2(f)}{S_{X+\eta}(f)}\,df = \int_{-\frac{1}{2}}^{\frac{1}{2}} \sup_{k\in\mathbb{Z}} \frac{S_X^2(f-f_sk)}{S_{X+\eta}(f-f_sk)}\,df. \tag{2.31}$$

Since the integrand in the RHS of (2.31) is periodic in f with period f_s, excluding a set of Lebesgue measure zero, the set $F^\star$ does not contain two frequencies $f_1, f_2 \in \mathbb{R}$ that differ by an integer multiple of f_s due to its minimality. This property will be given the following name:

Definition 2 (aliasing-free set) *A measurable set $F \subset \mathbb{R}$ is said to be* aliasing-free *with respect to the sampling frequency f_s if, for almost*[2] *all pairs $f_1, f_2 \in F$, it holds that $f_1 - f_2 \notin f_s\mathbb{Z} = \{f_sk, k \in \mathbb{Z}\}$.*

The aliasing-free property imposes the following restriction on the Lebesgue measure of a bounded set:

Proposition 2.4 *Let F be an aliasing-free set with respect to f_s. If F is bounded, then the Lebesgue measure of F does not exceed f_s.*

[2] By *almost all* we mean *for all but a set of Lebesgue measure zero*.

Proof By the aliasing-free property, for any $n \in \mathbb{Z} \setminus \{0\}$ the intersection of F and $F + nf_s$ is empty. It follows that for all $N \in \mathbb{N}$, $\mu\left(\cup_{n=1}^{N} \{F + f_s n\}\right) = N\mu(F)$. Now assume F is bounded by the interval $(-M, M)$ for some $M > 0$. Then $\cup_{n=1}^{N} \{F + f_s n\}$ is bounded by the interval $(-M, M + Nf_s)$. It follows that

$$\frac{\mu(F)}{f_s} = \frac{N\mu(F)}{Nf_s} = \frac{\mu(\cup_{n=1}^{N} \{F + nf_s\})}{Nf_s} \le \frac{2M + Nf_s}{Nf_s}.$$

Letting $N \to \infty$ implies that $\mu(F) \le f_s$. □

We denote by $\mathsf{AF}(f_s)$ the collection of all bounded aliasing-free sets with respect to f_s. Note that a process with spectrum support in $\mathsf{AF}(f_s)$ admits no aliasing when uniformly sampled at frequency f_s, i.e., such a process can be reconstructed with probability one from its non-noisy uniform samples at frequency f_s [54]. As the following theorem shows, the optimal pre-sampling filter is characterized by an aliasing-free set with an additional maximality property.

Theorem 2.5 *For a fixed f_s, the optimal pre-sampling filter $H^\star$ that maximizes $\widetilde{S}_{X|Y}(f)$, $f \in (-f_s/2, f_s/2)$ and minimizes $\mathsf{mmse}_{\mathsf{SB}(H)}(f_s)$ is given by*

$$H^\star(f) = \begin{cases} 1 & f \in F^\star, \\ 0 & \textit{otherwise}, \end{cases} \tag{2.32}$$

where $F^\star \in AF(f_s)$ satisfies

$$\int_{F^\star} \frac{S_X^2(f)}{S_{X+\eta}(f)} df = \sup_{F \in \mathsf{AF}(f_s)} \int_F \frac{S_X^2(f)}{S_{X+\eta}(f)} df. \tag{2.33}$$

The optimal MMSE form sampling at rate f_s is

$$\mathsf{mmse}^\star_{\mathsf{SB}}(f_s) = \sigma_X^2 - \int_{F^\star} \frac{S_X^2(f)}{S_{X+\eta}(f)} df. \tag{2.34}$$

Proof Since $\widetilde{S}_{X|Y}(f) \ge 0$, we can maximize the integral over $\widetilde{S}_{X|Y}(f)$ by maximizing the latter for every f in $\left(-\frac{f_s}{2}, \frac{f_s}{2}\right)$. For a given f, denote $h_k = |H(f - f_s k)|^2$, $x_k = S_X^2(f - f_s k)$ and $y_k = S_{X+\eta}(f - f_s k) = S_X(f - f_s k) + S_\eta(f - f_s k)$. We arrive at the following optimization problem

$$\begin{aligned} &\text{maximize} \quad \frac{\sum_{k\in\mathbb{Z}} x_k h_k}{\sum_{k\in\mathbb{Z}} y_k h_k} \\ &\text{subject to} \quad h_k \ge 0, \quad k \in \mathbb{Z}. \end{aligned}$$

Because the objective function is homogeneous in $\mathbf{h} = (...,h_{-1},h_0,h_1,...)$, the last problem is equivalent to

$$\begin{array}{ll} \text{maximize} & \sum_{k\in\mathbb{Z}} x_k h_k \\ \text{subject to} & h_k \geq 0, \quad k\in\mathbb{Z}, \quad \sum_{k\in\mathbb{Z}} y_k h_k = 1. \end{array}$$

The optimal value of this problem is $\max_k \frac{x_k}{y_k}$, i.e. the maximal ratio over all pairs x_k and y_k. The optimal $\mathbf{h}$ is the indicator for the optimal ratio:

$$h_k^\star = \begin{cases} 1 & k \in \operatorname{argmax}_k \frac{x_k}{y_k}, \\ 0 & otherwise. \end{cases}$$

If there is more than one k that maximizes $\frac{x_k}{y_k}$, then we can arbitrarily decide on one of them. Going back to our original problem formulation, we see that for almost every $f \in \left(-\frac{f_s}{2},\frac{f_s}{2}\right)$, the optimal $\widetilde{S}_{X|Y}(f)$ is given by

$$\widetilde{S}^\star_{X|Y}(f) = \max_{k\in\mathbb{Z}} \frac{S_X^2(f-f_sk)}{S_{X+\eta}(f-f_sk)},$$

and the optimal $H(f)$ is such that $|H(f-f_sk)|^2$ is non-zero for the particular k that achieves this maximum. This last property also implies that $F^\star$, the support of $H^\star$, is an aliasing-free set that satisfies (2.33). □

Remarks

(i) The proof also shows that

$$\int_{F^\star(f_s)} \frac{S_X^2(f)}{S_{X+\eta}(f)} df = \int_{-\frac{f_s}{2}}^{\frac{f_s}{2}} \widetilde{S}^\star_{X|Y}(f) df,$$

where

$$\widetilde{S}^\star_{X|Y}(f) \triangleq \sup_k \frac{S_X^2(f-f_sk)}{S_{X+\eta}(f-f_sk)},$$

i.e.

$$\mathsf{mmse}^\star_{\mathsf{SB}}(f_s) = \sigma_X^2 - \int_{-\frac{f_s}{2}}^{\frac{f_s}{2}} \widetilde{S}^\star_{X|Y}(f) df.$$

(ii) Since the SNR at each spectral line f cannot be changed by $H(f)$, the filter $H^\star$ can be specified only in terms of its support, i.e. in (2.32) we may replace 1 by any non-zero value, which may even vary with f.

Theorem 2.5 motivates the following definition:

Definition 3 *For a given spectral density $S(f)$ and a sampling frequency f_s, an aliasing free set $F^\star \in \mathsf{AF}(f_s)$*

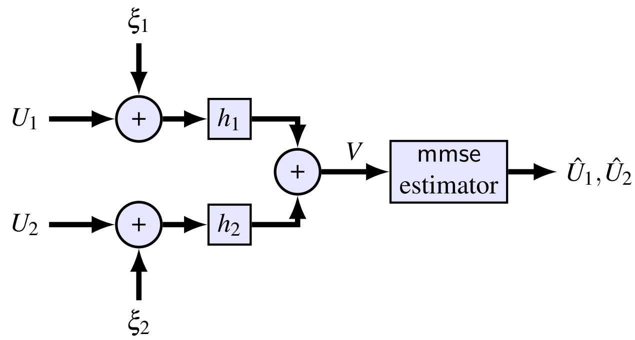

Figure 2.5: Joint MMSE estimation from a linear combination.

that satisfies

$$\int_{F^\star} S(f)df = \sup_{F \in \mathsf{AF}(f_s)} \int_F S(f)df$$

is called a maximal aliasing-free set *with respect to f_s and the spectral density $S(f)$. Such a set will be denoted by $F^\star(f_s, S)$.*

Roughly speaking, the maximal aliasing free set $F^\star(f_s, S)$ can be constructed by going over all frequencies $f \in (-f_s/2, f_s/2)$, and including in $F^\star(f_s, S)$ every frequency $f^\star \in \mathbb{R}$ such that $S(f^\star)$ is maximal among $\{S(f), f \in f^\star - f_s\mathbb{Z}\}$. The intuition behind Theorem 2.5 and the maximal aliasing-free property in Definition 3 can be seen as an interplay between the collection of information on all spectral components of the signal and elimination of interference between them, which is due to aliasing in sub-Nyquist sampling. It follows from Theorem 2.5 that the optimal pre-sampling filter eliminates aliasing at the price of completely suppressing information on weaker bands as these interfere with stronger ones. The following example provides an insight into this phenomena.

Example 2 (joint MMSE estimation) *Figure 2.5 illustrates two independent Gaussian random variables U_1 and U_2 with variances σ_1^2 and σ_2^2 respectively. We are interested in the MMSE estimate of the vector $\mathbf{U} = (U_1, U_2)$ from a noisy linear combination of their sum: $V = h_1(U_1 + \xi_1) + h_2(U_2 + \xi_2)$, where $h_1, h_2 \in \mathbb{R}$ and ξ_1, ξ_2 are another two Gaussian random variables with variances $\sigma_{\xi_1}^2$ and $\sigma_{\xi_2}^2$ respectively, representing independent additive noise. The MMSE in this estimation is given by*

$$\begin{aligned} \mathsf{mmse}(\mathbf{U}|V) &= \frac{1}{2}\left(\mathsf{mmse}(U_1|V) + \mathsf{mmse}(U_2|V)\right) \\ &= \frac{1}{2}\left(\sigma_1^2 + \sigma_2^2 - \frac{h_1^2\sigma_1^4 + h_2^2\sigma_2^4}{h_1^2(\sigma_1^2 + \sigma_{\xi_1}^2) + h_2^2(\sigma_2^2 + \sigma_{\xi_2}^2)}\right). \end{aligned} \tag{2.35}$$

The optimal choice of the coefficients vector $\mathbf{h} = (h_1, h_2)$ *that minimizes* (2.35) *is*

$$\mathbf{h} = \begin{cases} (c,0) & \frac{\sigma_1^4}{\sigma_1^2+\sigma_{\xi_1}^2} > \frac{\sigma_2^4}{\sigma_2^2+\sigma_{\xi_2}^2} \\ (0,c) & \frac{\sigma_1^4}{\sigma_1^2+\sigma_{\xi_1}^2} < \frac{\sigma_2^4}{\sigma_2^2+\sigma_{\xi_2}^2}, \end{cases}$$

where c is any constant different from zero. If

$$\frac{\sigma_1^4}{\sigma_1^2+\sigma_{\xi_1}^2} = \frac{\sigma_2^4}{\sigma_2^2+\sigma_{\xi_2}^2},$$

then any non-trivial linear combination results in the same estimation error.

The above example is easily generalized to any countable number of random variables $\mathbf{U} = (U_1, U_2, \ldots)$ and a respective noise sequence $\xi = (\xi_1, \xi_2, \ldots)$, such that $V = \sum_{i=1}^{\infty} h_i(U_i + \xi_i) < \infty$ with probability one. The optimal coefficient vector $\mathbf{h} = (h_1, h_2, \ldots)$ that minimizes $\mathsf{mmse}(\mathbf{U}|\mathbf{V})$ is the indicator for the maximum among

$$\left\{ \frac{\sigma_i^4}{\sigma_i^2+\sigma_{\xi_i}^2},\ i = 1,2,\ldots \right\}.$$

Since each frequency f in the support of $S_X(f)$ can be seen as an independent component of the process $X(\cdot)$ with an infinitely narrow band, the counterpart for the vector $\mathbf{U}$ in the MMSE estimation problem of $X(\cdot)$ from $Y[\cdot]$ are the spectral components of the source process that corresponds to the frequencies $f - f_s\mathbb{Z}$. These components are folded and summed together due to aliasing: each set of the form $f - f_s\mathbb{Z}$ corresponds to a linear combination of a countable number of independent Gaussian random variables attenuated by the coefficients $\{H(f - f_s k),\ k \in \mathbb{Z}\}$. The optimal choice of coefficients that minimizes the MMSE in joint estimation of all source components are those that pass only the spectral component with maximal $\frac{S_X^2(f')}{S_{X+\eta}(f')}$ among all $f' \in f - f_s\mathbb{Z}$, and suppress the remaining. This means that under the MSE criterion, the optimal choice is to eliminate aliasing at the price of losing all information contained in low-energy spectral components when they interfere with high-energy component.

An example of a maximal aliasing-free set for a specific PSD appears in Fig. 2.6. The MMSE with the optimal pre-sampling filter and with an all-pass filter is shown in Fig. 2.9. Interestingly, the two curves of Figure 2.9 have a similar shape although they correspond to different pre-sampling filters. Detailed analysis of this similarity is the subject of future work.

It also follows from Theorem 2.5 and Proposition 2.4 that a lower bound on $\mathsf{mmse}^\star_{\mathsf{SB}}(f_s)$ can be obtained by integrating over a set of Lebesgue measure f_s with maximal $\frac{S_X^2(f)}{S_{X+\eta}(f)}$ (that is, this set is not limited by the aliasing-free property), per the following corollary:

Corollary 2.6 *Let*

$$\underline{\mathsf{mmse}}(f_s) \triangleq \sigma_X^2 - \sup_{\mu(F) \le f_s} \int_F \frac{S_X^2(f)}{S_{X+\eta}(f)} df, \tag{2.36}$$

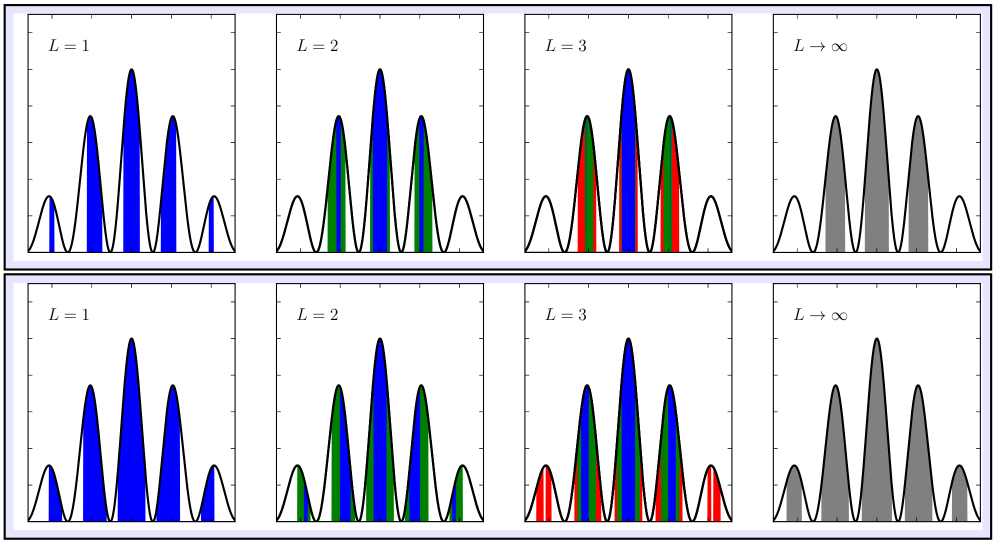

Figure 2.6: Maximal aliasing-free sets with respect to a multi-modal PSD and sampling frequencies $f_s = f_{\mathsf{Nyq}}/4$ (top) and $f_s = f_{\mathsf{Nyq}}/2$ (bottom), for $1,2$ and 3 sampling branches. The first, second and third maximal aliasing-free sets are given by the frequencies below the blue, green, and red areas, respectively. The Lebesgue measure of the colored area in each figure equals to f_s. Assuming $S_\eta(f) \equiv 0$, $\mathsf{mmse}^{\star}_{\mathsf{SB}}(f_s)$ equals to the area blocked by the filters, i.e., the white area bounded by the PSD. The ratio of this area to the total area bounded by the PSD is specified in each case.

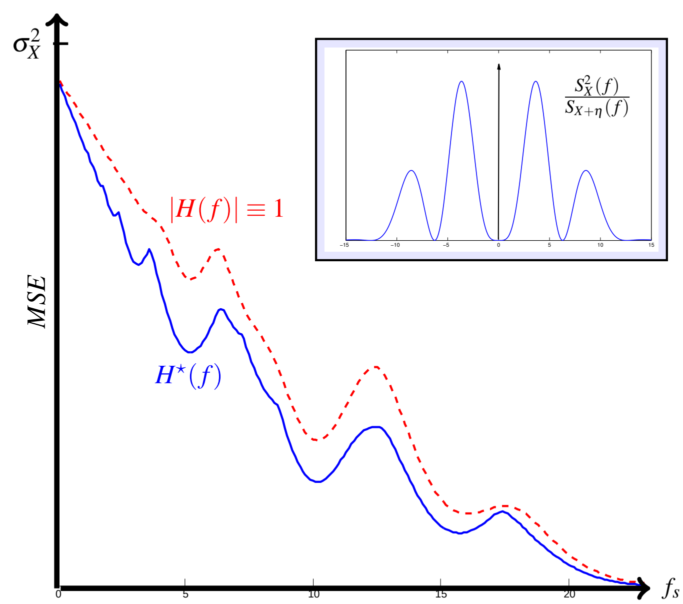

Figure 2.7: The MMSE as a function of the sampling frequency f_s in single branch sampling, with an optimal pre-sampling filter and an all-pass filter. The function $S_X^2(f)/S_{X+\eta}(f)$ is given in the small frame.

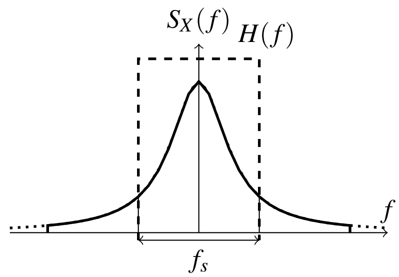

Figure 2.8: Unimodal PSD $S_\Omega(f)$ and its corresponding optimal pre-sampling filter.

where the supremum is taken over all measurable subsets of $\mathbb{R}$ *with Lebesgue measure not exceeding* f_s. *Then*

$$\underline{\mathsf{mmse}}(f_s) \leq \mathsf{mmse}^\star_{\mathsf{SB}}(f_s) \leq \mathsf{mmse}_{\mathsf{SB}(H)}(f_s)$$

for any SB uniform sampler.

When the noise process $\eta(\cdot)$ is taken to be zero we have

$$\underline{\mathsf{mmse}}(f_s) = \sigma_X^2 - \sup_{\mu(F)\leq f_s} \int_F S_X(f)df = \inf_{\mu(F)\leq f_s} \int_{\mathbb{R}\setminus F} S_X(f)df. \tag{2.37}$$

The RHS of (2.37) is zero if and only if the measure of the support of $S_X(f)$ does no exceed f_s. This last fact is the well-known condition of Landau for a stable set of sampling in Paley-Wiener spaces [30, 3]. Therefore, combined with the isomorphism defined by (2.5), Corollary 2.6 implies the necessity of sampling at the spectral occupancy in order to attain zero error in uniform sampling.

A special case in which the bound (2.36) is achieved is described in the following example.

Example 3 (unimodal PSD) *Consider the case where the function* $\Omega(f) = \frac{S_X^2(f)}{S_{X+\eta}(f)}$ *is unimodal in the sense that it is non-increasing for* $f > 0$, *as illustrated in Figure 2.8. It can be seen that the* $F^\star(\Omega, f_s)$ *is the interval* $(-f_s/2, f_s/2)$, *and the optimal pre-sampling filter is a lowpass filter with cutoff frequency* $f_s/2$. *Theorem 2.5 then implies that*

$$\mathsf{mmse}^\star_{\mathsf{SB}}(f_s) = \sigma_X^2 - \int_{-\frac{f_s}{2}}^{\frac{f_s}{2}} \frac{S_X^2(f)}{S_{X+\eta}(f)} df. \tag{2.38}$$

Since $S_X(f)$ *is symmetric and non-increasing for* $f > 0$, $\mathsf{mmse}^\star_{\mathsf{SB}}(f_s)$ *in* (2.14) *achieves the bound* (2.36).

In contrast to the case of a unimodal PSD described in Example 3, the bound in (2.36) cannot be achieved by a single sampling branch in general. It can, however, be approached by increasing the number of sampling branches, as will be discussed in the following two subsections.

2.5.2 Optimal filter-bank in multi-branch sampling

We now extend Theorem 2.5 to the MB uniform sampler of Figure 2.3.

Theorem 2.7 (optimal filter-bank in MB uniform sampling) *The optimal pre-sampling filters $H_1^\star(f),\ldots,H_L^\star(f)$ that minimize $\mathsf{mmse}_{\mathsf{MB}(H_1,\ldots,H_L)}(f_s)$ of (2.22) are given by*

$$H_l^\star(f) = \mathbf{1}_{F_l^\star}(f) \triangleq \begin{cases} 1 & f \in F_l^\star, \\ 0 & f \notin F_l^\star, \end{cases}, \quad l = 1,\ldots,L, \tag{2.39}$$

where the sets $F_1^\star,\ldots,F_L^\star \subset \mathbb{R}$ satisfy:

(i) $F_l^\star \in \mathsf{AF}(f_s/L)$ *for all* $l = 1,\ldots,L$.

(ii) *For* $l = 1$,

$$\int_{F_1^\star} \frac{S_X^2(f)}{S_{X+\eta}(f)} df = \int_{-\frac{f_s}{2}}^{\frac{f_s}{2}} \widetilde{S}_1^\star(f) df,$$

where

$$\widetilde{S}_1^\star(f) \triangleq \sup_{k\in\mathbb{Z}} \frac{S_X^2(f-kf_s/L)}{S_{X+\eta}(f-kf_s/L)},$$

and for $l = 2,\ldots,L$,

$$\int_{F_l^\star} \frac{S_X^2(f)}{S_{X+\eta}(f)} df = \int_{-\frac{f_s}{2}}^{\frac{f_s}{2}} \widetilde{S}_l^\star(f) df,$$

where

$$\widetilde{S}_l^\star(f) \triangleq \sup_{k\in\mathbb{Z}} \frac{S_X^2(f-kf_s/L)}{S_{X+\eta}(f-kf_s/L)} \mathbf{1}_{\mathbb{R}\setminus\{F_1^\star\cup\cdots\cup F_{l-1}^\star\}}.$$

The resulting MMSE equals

$$\mathsf{mmse}^\star_{\mathsf{MB}(L)}(f_s) \triangleq \sigma_X^2 - \sum_{l=1}^L \int_{F_l^\star} \frac{S_X^2(f)}{S_{X+\eta}(f)} df = \sigma_X^2 - \sum_{l=1}^L \int_{-\frac{f_s}{2}}^{\frac{f_s}{2}} \widetilde{S}_l^\star(f) df.$$

Remarks:

(i) The proof implies an even stronger statement than Theorem 2.7: the filters $H_1^\star,\ldots,H_L^\star$ yield a set of eigenvalues of $\widetilde{\mathbf{S}}_{X|\mathbf{Y}}(f)$ that are uniformly maximal, in the sense that the ith eigenvalue of $\widetilde{\mathbf{S}}_{X|\mathbf{Y}}(f)$ is always smaller than the ith eigenvalue of $\widetilde{\mathbf{S}}^\star_{X|\mathbf{Y}}(f)$.

(ii) As in the single-branch case in Theorem 2.5, the filters $H_1^\star,\ldots,H_L^\star$ are specified only in terms of their support, and in (2.39) we can replace 1 by any non-zero value which may vary with l and f.

(iii) Condition (ii) for the sets $F_1^\star,\ldots,F_L^\star$ can be relaxed in the following sense: if $F_1^\star,\ldots,F_L^\star$ satisfy condition (i) and (ii), then $\mathsf{mmse}^\star_{\mathsf{MB}(L)}(f_s)$ is achieved by any pre-sampling filters defined as the indicators of some sets $F_1',\ldots,F_L'$ in $\mathsf{AF}(f_s/L)$ that satisfy

$$\sum_{l=1}^{L}\int_{F_l'}\frac{S_X^2(f)}{S_{X+\eta}(f)}df=\sum_{l=1}^{L}\int_{F_l^\star}\frac{S_X^2(f)}{S_{X+\eta}(f)}df.$$

In particular, the set of optimal pre-sampling filters is usually not (essentially) unique.

(iv) One possible construction for $F_1^\star,\ldots,F_L^\star$ is as follows: For each $-\frac{f_s}{2L}\le f<\frac{f_s}{2L}$, denote by $f_1^\star(f),\ldots,f_L^\star(f)$ the L frequencies that correspond to the largest values among $\left\{\frac{S_X^2(f-f_sk)}{S_{X+\eta}(f-f_sk)},\,k\in\mathbb{Z}\right\}$. Then assign each $f_l^\star(f)$ to $F_l^\star$. Under this construction, the set $F_l^\star$ can be seen as the lth *maximal aliasing free set* with respect to f_s/L and $\frac{S_X^2(f)}{S_{X+\eta}(f)}$. This is the construction that was used in Figure 2.6.

Proof Let $\mathbf{H}(f)\in\mathbb{C}^{\mathbb{Z}\times L}$ be the matrix with infinite rows and L columns given by

$$\mathbf{H}(f)=\begin{pmatrix}\vdots & \vdots & \cdots & \vdots\\ H_1(f-2f_s) & H_2(f-2f_s) & \cdots & H_P(f-2f_s)\\ H_1(f-f_s) & H_2(f-f_s) & \cdots & H_P(f-f_s)\\ H_1(f) & H_2(f) & \cdots & H_P(f)\\ H_1(f+f_s) & H_2(f+f_s) & \cdots & H_P(f+f_s)\\ H_1(f+2f_s) & H_2(f+2f_s) & \cdots & H_P(f+2f_s)\\ \vdots & \vdots & \cdots & \vdots\end{pmatrix}.$$

In addition, denote by $\mathbf{S}(f)\in\mathbb{R}^{\mathbb{Z}\times\mathbb{Z}}$ and $\mathbf{S}_n(f)\in\mathbb{R}^{\mathbb{Z}\times\mathbb{Z}}$ the infinite diagonal matrices with diagonal elements $\{\mathbf{S}_X(f-f_sk),\ k\in\mathbb{Z}\}$ and $\left\{\mathbf{S}_{X+\eta}(f-f_sk),\ k\in\mathbb{Z}\right\}$, respectively. The infinite dimensional matrix notation is consistent with the theory of compact operators between Hilbert spaces [58]. For example, $\mathbf{H}(f)$ represents a compact operator between the Hilbert spaces $\mathbb{C}^L$ and $\ell_2(\mathbb{C})$.

With the notation above, we can write

$$\widetilde{\mathbf{S}}_{X|Y}(f)=\left(\mathbf{H}^*\mathbf{S}_n\mathbf{H}\right)^{-\frac{1}{2}*}\mathbf{H}^*\mathbf{S}^2\mathbf{H}\left(\mathbf{H}^*\mathbf{S}_n\mathbf{H}\right)^{-\frac{1}{2}},$$

where we suppressed the dependency on f in order to simplify notation. Denote by $\mathbf{H}^\star(f)$ the matrix $\mathbf{H}(f)$ that corresponds to the filters $H_1^\star(f),\ldots,H_L^\star(f)$ that satisfy conditions (i) and (ii) in the theorem. From the definition of the sets $F_1^\star,\ldots,F_L^\star$, the structure of $\mathbf{H}^\star(f)$ can be described as follows: each column has a single non-zero entry, such that the first column indicates the largest among $\left\{\frac{S_X^2(f-f_sk)}{S_{X+\eta}(f-f_sk)},\,k\in\mathbb{Z}\right\}$, which is the diagonal of $\mathbf{S}(f)\mathbf{S}_n^{-1}(f)\mathbf{S}(f)$. The second column corresponds to the second largest entry of $\left\{\frac{S_X^2(f-f_sk)}{S_{X+\eta}(f-f_sk)},\,k\in\mathbb{Z}\right\}$, and so on for all L columns of $\mathbf{H}^\star(f)$. It follows that $\widetilde{\mathbf{S}}^\star_{X|\mathbf{Y}}(f)$ is an $L\times L$ diagonal matrix whose non-zero entries are the L largest values among $\left\{\frac{S_X^2(f-f_sk)}{S_{X+\eta}(f-f_sk)},\,k\in\mathbb{Z}\right\}$, i.e, $\lambda_l\left(\widetilde{\mathbf{S}}^\star_{X|\mathbf{Y}}(f)\right)=J_l^\star(f)$, for all $p=1,\ldots,P$.

It is left to establish the optimality of this choice of pre-sampling filters. Since the rank of $\widetilde{\mathbf{S}}_{X|Y}(f)$ is at most L, in order to complete the proof it is enough to show that for any $\mathbf{H}(f)$, the L non-zeo eigenvalues of the corresponding $\widetilde{\mathbf{S}}_{X|\mathbf{Y}}(f)$ are smaller than the L largest eigenvalues of $\mathbf{S}(f)\mathbf{S}_n^{-1}(f)\mathbf{S}(f)$ compared by their respective order. Since the matrix entries of the diagonal matrices $\mathbf{S}(f)$ and $\mathbf{S}_n(f)$ are positive, the eigenvalues of $\widetilde{\mathbf{S}}_{X|\mathbf{Y}}(f)$ are identical to the L non-zero eigenvalues of the matrix

$$\mathbf{S}\mathbf{H}\left(\mathbf{H}^*\mathbf{S}_n\mathbf{H}\right)^{-1}\mathbf{H}^*\mathbf{S}.$$

It is enough to prove that the matrix

$$\mathbf{S}\mathbf{S}_n^{-1}\mathbf{S}-\mathbf{S}\mathbf{H}\left(\mathbf{H}^*\mathbf{S}_n\mathbf{H}\right)^{-1}\mathbf{H}^*\mathbf{S},$$

is positive, in the sense that it defines a positive linear operator on the Hilbert space $\ell_2(\mathbb{C})$. This last fact is equivalent to

$$\mathbf{a}^*\mathbf{S}\mathbf{S}_n^{-1}\mathbf{S}\mathbf{a}-\mathbf{a}^*\mathbf{S}\mathbf{H}\left(\mathbf{H}^*\mathbf{S}_n\mathbf{H}\right)^{-1}\mathbf{H}^*\mathbf{S}\mathbf{a}^*\geq 0, \tag{2.40}$$

for any sequence $\mathbf{a}\in\ell_2(\mathbb{C})$. By factoring out $\mathbf{S}\mathbf{S}_n^{-\frac{1}{2}}$ from both sides, (2.40) reduces to

$$\mathbf{a}^*\mathbf{a}-\mathbf{a}^*\mathbf{S}_n^{\frac{1}{2}}\mathbf{H}\left(\mathbf{H}^*\mathbf{S}_n\mathbf{H}\right)^{-1}\mathbf{H}^*\mathbf{S}_n^{\frac{1}{2}}\mathbf{a}\geq 0. \tag{2.41}$$

The Cauchy-Schwartz inequality implies that

$$\begin{aligned}&\left(\mathbf{a}^*\mathbf{S}_n^{\frac{1}{2}}\mathbf{H}\left(\mathbf{H}^*\mathbf{S}_n\mathbf{H}\right)^{-1}\mathbf{H}^*\mathbf{S}_n^{\frac{1}{2}}\mathbf{a}\right)^2\\ &\leq\mathbf{a}^*\mathbf{a}\times\mathbf{a}^*\mathbf{S}_n^{\frac{1}{2}}\mathbf{H}\left(\mathbf{H}^*\mathbf{S}_n\mathbf{H}\right)^{-1}\mathbf{H}^*\mathbf{S}_n^{\frac{1}{2}}\mathbf{S}_n^{\frac{1}{2}}\mathbf{H}\left(\mathbf{H}^*\mathbf{S}_n\mathbf{H}\right)^{-1}\mathbf{H}^*\mathbf{S}_n^{\frac{1}{2}}\mathbf{a}\\ &=\mathbf{a}^*\mathbf{a}\left(\mathbf{a}^*\mathbf{S}_n^{\frac{1}{2}}\mathbf{H}\left(\mathbf{H}^*\mathbf{S}_n\mathbf{H}\right)^{-1}\mathbf{H}^*\mathbf{S}_n^{\frac{1}{2}}\mathbf{a}\right).\end{aligned} \tag{2.42}$$

Dividing (2.42) by $\left(\mathbf{a}^*\mathbf{S}_n^{\frac{1}{2}}\mathbf{H}\left(\mathbf{H}^*\mathbf{S}_n\mathbf{H}\right)^{-1}\mathbf{H}^*\mathbf{S}_n^{\frac{1}{2}}\mathbf{a}\right)$ leads to (2.41). □

Figure 2.9 illustrates $\mathsf{mmse}^\star_{\mathsf{MB}(L)}(f_s)$ as a function of f_s for a specific PSD and $L=1,2$ and 3. As this figure shows, increasing the number of sampling branches does not necessarily decrease $\mathsf{mmse}^\star_{\mathsf{MB}(L)}(f_s)$ and may even increase it for some f_s. This perhaps-counter intuitive phenomenon arises since for L_2 and L_1 co-primes, any two MB sampling systems of orders L_1 and L_2 would lead to different samples at their output. Nevertheless, we will see below that $\mathsf{mmse}^\star_{\mathsf{MB}(L)}(f_s)$ converges to a fixed number as L increases.

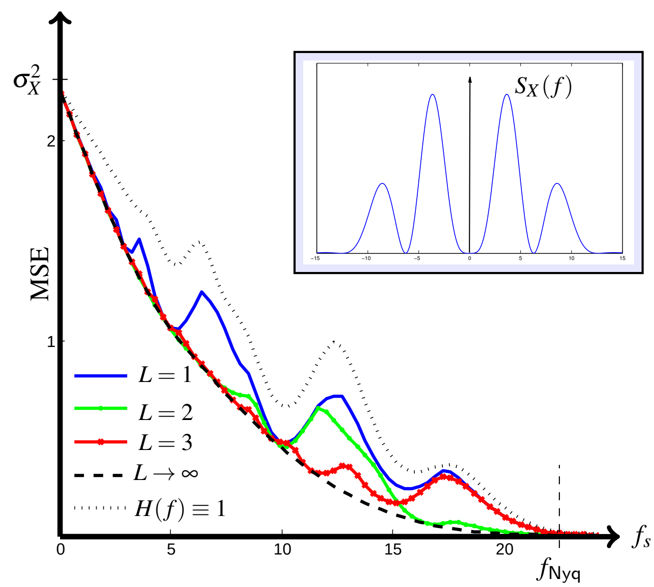

Figure 2.9: MMSE under multi-branch sampling and optimal pre-sampling filter-bank, for $L = 1,2$ and 3. The dashed line corresponds to the bound (2.36), which by Theorem 2.8 is attained as $L \to \infty$. The dotted line corresponds to $L = 1$ and an all-pass pre-sampling filter. The PSD $S_X(f)$ is given in the small frame, where we assumed $S_\eta(f) \equiv 0$.

2.5.3 Asymptotically many sampling branches

The set $F^\star$ defined in Theorem 2.5 to describe $\mathsf{mmse}^\star_{\mathsf{SB}}(f_s)$ was obtained by imposing two constraints: (1) a measure constraint $\mu(F^\star) \leq f_s$, associated only with the sampling frequency, and (2) an aliasing-free constraint, imposed by the sampling structure. Theorem 2.7 says that in the case of MB sampling, the aliasing-free constraint can be relaxed to aliasing-free in each branch, but with respect to f_s/L instead of f_s. Therefore, the overall passband of the system $\bigcup_{l=1}^L F_l^\star$ need not be aliasing-free, although its Lebesgue measure still does not exceed f_s since each $F_l^\star$ is aliasing free. As a result, it follows that

$$\mathsf{mmse}^\star_{\mathsf{MB}(L)}(f_s) \geq \sigma_X^2 - \inf_{\mu(F)\leq f_s} \int_F \frac{S_X^2(f)}{S_{X+\eta}(f)} df = \underline{\mathsf{mmse}}(f_s).$$

In other words, the statement in Corollary 2.6 can be extended to include MB uniform sampling.

While increasing the number of sampling branches does not improve the bound (2.36), it allows a MB sampler to approach it for any PSD (and not only for the unimodal PSD, where it is attained with a SB sampler): as the following theorem shows, and as suggested by Figure 2.6, the set $\bigcup_{l=1}^L F_l^\star$ eventually converges to some set $F^\dagger$ that attains (2.36). The set $F^\dagger$ corresponds to the frequencies below the gray area in the case $L\to\infty$ in Figure 2.6.

Theorem 2.8 *For any $f_s > 0$ and $\varepsilon > 0$, there exists $L \in \mathbb{N}$ and a set of L filters $H_1,\ldots,H_L$ with supports in $\mathsf{AF}(f_s/L)$, such that*

$$\mathsf{mmse}^\star_{\mathsf{MB}(H_1,\ldots,H_L)}(f_s) - \varepsilon < \underline{\mathsf{mmse}}(f_s).$$

Proof Since any interval of length at most f_s/L is aliasing free for sampling at rate f_s/L, it is enough to show that any set $F^\dagger$ of measure f_s that satisfies

$$\int_{F^\dagger} \frac{S_X^2(f)}{S_{X+\eta}(f)} df = \sup_{\mu(F)\leq f_s} \int_F \frac{S_X^2(f)}{S_{X+\eta}(f)} df$$

can be approximated by L intervals of measure at most f_s/L. Consider the measure μ_{S_X} defined by

$$\mu_S(A) = \int_A \frac{S_X^2(f)}{S_{X+\eta}(f)} df$$

for a Lebesgue measurable set A. Since $S_X(f)$ is $\mathbf{L}_1(\mathbb{R})$, we can choose a set $G \subset F^\dagger$ such that $S_X(f)$ is bounded on G and such that $\mu_S(G) > \mu_S(F^\star) - \varepsilon/3$. Denote by S_M the maximal value of $S_X(f)$ on G and note that $S_X^2(f)/S_{X+\eta}(f) \leq S_M$ on G as well. The measure μ_S is absolutely continuous with respect to the Lebesgue measure and hence is a regular measure [59]. Therefore, there exists M intervals $I_1,\ldots,I_M$ such that $\cup_{i=1}^M I_i \subset G$ and $\mu_S(\cup_{i=1}^M I_i) > \mu_S(G) - \varepsilon/3 > \mu_S(F^\dagger) - 2\varepsilon/3$. We can assume that $I_1,\ldots,I_M$ are disjoint, for otherwise we use $I_1' = I_1$, $I_2' = I_2 \setminus I_1'$, $I_3' = I_3 \setminus (I_1' \cup I_2')$, and so forth. Therefore, $\sum_{i=1}^M \mu(I_i) \leq f_s$. For $\delta > 0$, let $L_i = \lfloor M\mu(I_i)/\delta \rfloor$ and $L = \sum_{i=1}^M L_i$. We now define L pre-sampling filters as follows: for each $i = 1,\ldots,M$, consider L_i disjoint intervals $I_{i,1},\dots,I_{i,j}$ of length $r=\delta/M$ that are sub-intervals of I_i. Since $\mu(I_i)\ge L_i r$, such L_i intervals exist and we set $A_i = I_i \setminus \cup_{j=1^{L_i}} I_{i,j}$. That is, A_i is the part of the interval I_i that cannot be covered by L_i intervals of length r. In particular, $\mu(A_i)\le r$. Finally, set the support of each filter $H_{i,j}$ to be $I_{i,j}$. Note that

$$\mu(\sum_{i,j} \mathrm{supp} H_{i,j}) = \sum_{i=1}^{M} L_i r \le r \sum_{i=1}^{M} M\mu(I_i)/\delta \le f_s.$$

This way we have defined $L=L_1+\ldots+L_M$ filters, each of passband of width $r\le f_s/L$. It is left to show that

$$\sum_{i,j} \mu_S(\mathrm{supp} H_{i,j}) = \sum_{i,j} \int_{\mathrm{supp} H_{i,j}} \frac{S_X^2(f)}{S_{X+\eta}(f)} df > \int_{F^\dagger} \frac{S_X^2(f)}{S_{X+\eta}(f)} df - \varepsilon.$$

Let m_S be the essential supremum of $S_X/S_{X+\eta}$ on G. We have

$$\mu_S(A_i) = \int_{A_i} \frac{S_X^2(f)}{S_{X+\eta}(f)} df \le m_s \mu(A_i) \le m_s r.$$

It follows that

$$\begin{aligned}\mu_S(\sum_{i,j} \mathrm{supp} H_{i,j}) &= \sum_{i=1}^{M}\sum_{j=1}^{L_i} \mu_S(I_{i,j}) = \sum_{i=1}^{M} \mu_S(I_i) - \sum_{i=1}^{M} \mu_S(A_i) \\ &\ge \mu_S(G) - \varepsilon/3 - M m_s r \ge \mu_S(F^\dagger) - 2\varepsilon/3 - M m_s \delta.\end{aligned}$$

Taking $\delta = \varepsilon/(3Mm_s)$ leads to the desired result. □

An immediate conclusion from Theorem 2.8 is that $\mathsf{mmse}^\star_{\mathsf{MB}(L)}(f_s)$ is guaranteed to converge to $\underline{\mathsf{mmse}}(f_s)$ as the number of sampling branches L goes to infinity, as illustrated in Figure 2.9. In other words, with enough sampling branches and a careful selection of the pre-sampling filters, MB uniform sampling attains the lower bound (2.36). Recall that with zero noise, this lower bound vanishes if and only if f_s exceeds the spectral occupancy of $X(\cdot)$. Hence, MB uniform sampling can be used to attain zero MMSE for sampling rates at or above the theoretical limit of Landau [30].

2.6 Chapter Summary

This chapter presented a general sampling framework for random stationary processes that are not necessarily bandlimited. We call this framework linear bounded sampling. Our framework also allows for sampling and pre-sampling operations limited to a finite time horizon, a feature that is required in order to consider the concatenation of sampling and source coding as in the ADX setting. In addition, we considered the problem of estimating a Gaussian stationary process from its noisy samples obtained through a bounded linear sampler under an MSE criterion. For the special case of time-invariant uniform sampling, we derived closed form expressions for this MSE as a function of the uniform sampling rate f_s and the pre-sampling operation. Further optimization over the pre-sampling operation yielded the function $\underline{\mathsf{mmse}}(f_s)$. This function provides a lower bound for the MMSE in any time-invariant uniform sampling system, and depends only on the PSD of the source, the noise, and the sampling rate. Moreover, we showed that $\underline{\mathsf{mmse}}(f_s)$ can be attained using a multi-branch uniform sampler with a sufficiently large number of sampling branches. When the noise is zero, the function $\underline{\mathsf{mmse}}(f_s)$ vanishes if and only if the sampling rate f_s exceeds the Landau rate of $X(\cdot)$. Therefore, $\underline{\mathsf{mmse}}(f_s)$ can be seen as the stochastic version of Landau's characterization of stable sampling in Paley Wiener spaces.

Since the ADX setting of Figure 1.2 includes, in addition to sampling, encoding of the samples with bitrate R, the function $\underline{\mathsf{mmse}}(f_s)$ provides a lower bound on the distortion in ADX. Since the precise description of any sample of a Gaussian process requires an infinite number of bits, this lower bound is attained in the ADX setting only as the bitrate R goes to infinity.

Chapter 3

Minimal Distortion subject to a Bitrate Constraint

In this chapter we consider the problem of representing a random signal under a constraint on the number of bits per unit time available for this representation. The random signal can be seen as the realization of an information source that generates information over time. The problem of encoding the signal subject to a bit per time unit constraint is denoted as the source coding problem, as introduced and solved by Shannon in [4]. The most simple source coding setting describes an encoder that maps the signal at its input to a finite set and a decoder that recovers the original signal from this set. This setting follows from our general ADX setting in Figure 1.2 if we ignore the sampling constraint, or otherwise assume that the encoder can recover the original signal $X(\cdot)$ before encoding it. Therefore, the minimal distortion in this recovery, denoted as the distortion-rate function (DRF), provides a lower bound for the distortion in ADX.

In general, however, sampling or noise may prevent the encoder from obtaining an exact description of $X(\cdot)$. Therefore, the ADX setting gives rise to a variant of the source coding problem in which the encoder has no direct access to the signal it is required to encode. This variant, first considered in [15], is known as the *remote*, *noisy*, or *indirect* source coding problem. Consequently, the minimal distortion in recovering the original signal is denoted as the indirect DRF. In this chapter we study general properties of the indirect DRF. It is well known that under a quadratic distortion measure, the indirect DRF can be decomposed into a minimal MSE term plus a standard source coding problem with respect to the optimal MSE estimator [16] [Ch. 3.5][6]. In this chapter we generalize this decomposition to allow for the encoding of a realization of a signal over a finite time interval T, where the stochastic relation between the observable signal and the original one may vary with T. This generalization is required in order to consider the concatenation of sampling and encoding as in the ADX setting. In addition, we demonstrate the usefulness of the aforementioned decomposition in various scenarios involving Gaussian signals.

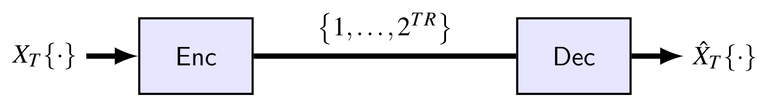

Figure 3.1: The standard source coding problem: encoding with full source information.

3.1 Classical Source Coding

Since the ADX system of Figure 1.2 is subject to a bitrate constraint, a lower bound for the distortion in ADX is given by the minimal distortion in the source coding problem of encoding an analog signal using a finite number of bits per unit time. In this section we review classical results from source coding theory on the characterization of this minimal distortion. In addition to providing a lower bound for the ADX distortion, these results are useful as references as we consider their extensions in subsequent sections.

Figure 3.1 illustrates the classical source coding problem of Shannon [4]: the source signal is the stochastic process $X(\cdot)$. The *encoder* observes a realization $x_T(\cdot)$ of $X(\cdot)$ over the time interval $[-T/2, T/2]$ and is required to describe its observation to the decoder using a single index out of a set of 2^{TR} possible indices. The *decoder* is required to recover $x_T(\cdot)$ by observing the index produced by the encoder. Since there are only 2^{TR} possible indices, each index i is associated with a single reconstruction waveform $\widehat{x}_i^T(\cdot)$. We assume that the set of reconstruction waveforms and index assignments is revealed beforehand to the encoder and decoder.

Specialized to the case of a quadratic distortion measure, the operational distortion-rate function (DRF) of the process $X(\cdot)$, or the *optimal performance theoretically achievable* (OPTA), is defined to be the minimal expected $\mathbf{L}_2$ norm distance between $X_T(\cdot)$ and its reconstruction $\widehat{X}_T$, where the optimization is over all encoder and decoder mappings of TR bits and over the time-horizon T. Namely, the OPTA distortion function is defined as

$$\delta_X(R) = \liminf_{T\to\infty} \inf_{\mathsf{Enc-Dec}} \frac{1}{T}\int_{-T/2}^{T/2} \mathbb{E}\left(X(t)-\widehat{X}(t)\right)^2 dt. \tag{3.1}$$

Shannon's classical *source coding theorem* [4, 5] and its extensions to continuous alphabets [60, 61] imply that, for a stationary and block-ergodic process $X(\cdot)$, the function $\delta_X(R)$ equals its *information* DRF, defined as

$$D_X(R) \triangleq \liminf_{T\to\infty} \inf_{I(X_T(\cdot);\widehat{X}_T(\cdot))\le RT} \frac{1}{T}\int_{-T/2}^{T/2} \mathbb{E}\left(X(t)-\widehat{X}(t)\right)^2 dt, \tag{3.2}$$

where the infimum is over joint probability distributions $P_{X_T,\widehat{X}_T}$ whose mutual information is constrained to RT bits, and their marginals over $X_T(\cdot)$ agrees with the distribution of $X(\cdot)$.

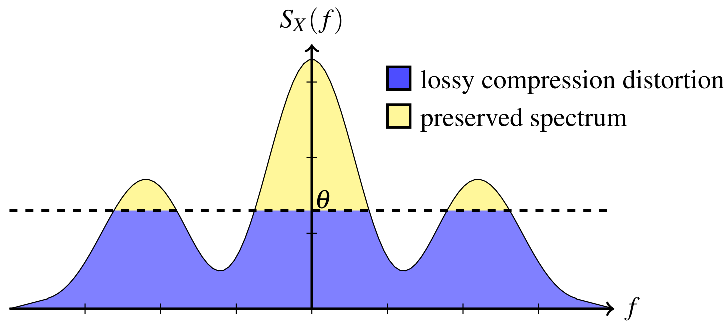

Figure 3.2: Water-filling interpretation of (3.3): water is poured into the area bounded by the graph of $S_X(f)$ up to level θ. The bitrate R is tied to the water level θ through the preserved part of the spectrum (3.3b). The distortion D is the error due to lossy compression given by (3.3a).

3.1.1 Gaussian stationary processes

For the special case of a Gaussian stationary process $X(t)$ with PSD $S_X(f)$, Pinsker [62] showed that (3.2) has the following parametric representation:

$$D_X(R_\theta) = \int_{-\infty}^{\infty} \min\{S_X(f), \theta\}\, df \tag{3.3a}$$

$$R_\theta = \frac{1}{2}\int_{-\infty}^{\infty} \log^+ [S_X(f)/\theta]\, df. \tag{3.3b}$$

Expression (3.3) has the water-filling interpretation given in Figure 3.2: the distortion in (3.3a) is the area bounded by the PSD and the *water-level* parameter θ, where the latter is determined to satisfy (3.3b). The joint probability distribution $P_{\widehat{X}_T(\cdot),X_T(\cdot)}$ which attains the infimum in the RHS of (3.2) converges, as $T \to \infty$, to the joint distribution of two jointly Gaussian and stationary processes $X(\cdot)$ and $\widehat{X}(\cdot)$, defined by the Gaussian channel

$$X(t) = \widehat{X}(t) + \varepsilon(t), \tag{3.4}$$

where $\widehat{X}(\cdot)$ and $\varepsilon(\cdot)$ are two Gaussian stationary processes independent of each other with PSDs given by the preserved and the lossy compression distortion parts of the spectrum in Figure 3.2, respectively [6, 63].

The source coding theorem also implies that the asymptotic distribution (3.4) can be used to derive an achievable coding scheme that attains $D_X(R)$. Roughly speaking, this scheme is described as follows [6]: fix $\varepsilon > 0$ and $T > 0$; generate $2^{\lfloor T(R+\varepsilon)\rfloor}$ waveforms $\left\{\widehat{x}_i(\cdot),\, i = 1,\ldots,2^{\lfloor T(R+\varepsilon)\rfloor}\right\}$, where each waveform is drawn randomly from the distribution of $\widehat{X}(\cdot)$ defined by (3.4). The encoder, upon observing a realization $x(\cdot)$ of $X(\cdot)$ over $[-T/2, T/2]$, sends the index i for which the $\mathbf{L}_2$ norm between $x(\cdot)$ and $\widehat{x}_i(\cdot)$ is minimal. The decoder uses $\widehat{x}_i(\cdot)$ as its estimate for $X(\cdot)$.

Going back to our general ADX setting of Figure 1.2, the distortion-rate function of the continuous-time Gaussian process $X(\cdot)$ describes the optimal trade-off between bitrate and distortion regardless of any other parameters, and thus provides a lower bound for the minimal distortion in ADX. In the general ADX setting,

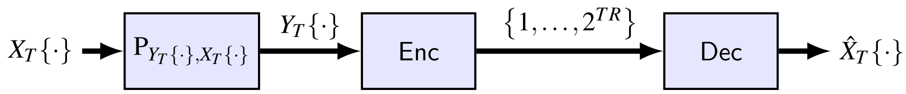

Figure 3.3: Generic indirect source coding setting with information source X and observation process Y.

however, the encoder is required to encode the source $X(\cdot)$ using R bits per unit time by observing only samples of $X(\cdot)$, rather than $X(\cdot)$ itself. In what follows, we explore the general source coding setting in which an encoder is required to describe an information source which it cannot observe directly.

3.2 Indirect Source Coding

An important extension to Shannon's source coding problem arises when the encoder has no direct access to the signal it tries to describe to the decoder. This extension is referred to as the *indirect* source coding problem. In this section we review general known properties of the indirect source coding setting. We then extend many of these properties to the kind of processes that arise from sampling and estimation in bounded linear sampling. Throughout this section the curly brackets notation $\{\cdot\}$ is used instead of $(\cdot)$ or $[\cdot]$ for expressions that apply whenever each of the processes is indexed by either discrete or continuous time. Depending on the index set, integration is either with respect to the Lebesgue measure or the counting measure.

The general indirect source coding problem setting is illustrated in Figure 3.3. In this problem, the encoder observes the signal $Y_T\{\cdot\}$ and is required to encode it so that the distortion between $X_T\{\cdot\}$ and the decoder output $\widehat{X}_T\{\cdot\}$ is minimized. The relation between $X_T\{\cdot\}$ and $Y_T\{\cdot\}$ is described by a family of joint probability distributions $\left\{P_{Y_T\{\cdot\},X_T\{\cdot\}}, T>0\right\}$ whose marginal distribution over $X_T\{\cdot\}$ agrees with the distribution of the process $X(\cdot)$. By working with joint distributions rather than conditional distributions, we avoid many of the complications occurring in conditional distributions of continuous alphabets [64]. The OPTA or the operational indirect DRF in this setting is denoted as $\delta_{X|Y}(R)$.

3.2.1 Indirect source coding theorem under quadratic distortion

Under the special case of a quadratic distortion measure $d(x,\widehat{x}) = (x-\widehat{x})^2$ the OPTA can be represented using a simplified expression. This expression highlights conditions on the posterior distribution of $X_T\{\cdot\}$ given $Y_T\{\cdot\}$ that are necessary for the characterizations of the OPTA using an information expression.

Theorem 3.1 *(i) For $T>0$, let $X\{\cdot\}$ be a stochastic process such that*

$$\sigma_X^2 \triangleq \lim_{T\to\infty} \frac{1}{T} \int_{-T/2}^{T/2} \mathbb{E}\left[X\{t\}^2\right] dt$$

exists and finite. Denote by $X_T\{\cdot\}$ the restriction of $X\{\cdot\}$ to the interval $[-T/2, T/2]$ and let $Y_T\{\cdot\}$ be another process (on the same probability space). The OPTA in recovering $X\{\cdot\}$ from an encoded version of $Y_T\{\cdot\}$ using TR bits satisfies

$$\delta_{X|Y}(R) = \liminf_{T\to\infty} \mathsf{mmse}(X_T|Y_T) + \liminf_{T\to\infty} \delta_{\widetilde{X}_T}(R), \tag{3.5}$$

where $\delta_{\widetilde{X}_T}(R)$ is the OPTA in encoding the process

$$\widetilde{X}_T\{t\} = \mathbb{E}[X\{t\}|Y_T\{\cdot\}] \tag{3.6}$$

using TR bits.

(ii) Assume that the conditions in (i) hold, and that $\widetilde{X}_T\{\cdot\}$ has continuous amplitude and converges in law to a continuous-time continuous-amplitude process $\widetilde{X}\{\cdot\}$ that is asymptotic mean stationary [65] in the following sense: for every $\alpha_1, \ldots, \alpha_n, t_1, \ldots, t_n \in \mathbb{R}$,

$$\frac{1}{T}\int_{-T/2}^{T/2} \mathbb{P}(\widetilde{X}_T\{t_1+\tau\} \le \alpha_1, \ldots . \widetilde{X}_T\{t_n+\tau\} \le \alpha_n) d\tau \overset{T\to\infty}{\longrightarrow} \mathbb{P}(\widetilde{X}\{0\} \le \alpha_1, \ldots \widetilde{X}\{t_n - t_1\} \le \alpha_n).$$

Then

$$\delta_{X|Y}(R) = \mathsf{mmse}(X|Y) + D_{\widetilde{X}}(R), \tag{3.7}$$

where

$$\mathsf{mmse}(X|Y) = \sigma_X^2 - \lim_{T\to\infty} \frac{1}{T}\int_{-T/2}^{T/2} \mathbb{E}\left[\widetilde{X}^2\{t\}\right] dt,$$

and $D_{\widetilde{X}}(R)$ is the information DRF of the process $\widetilde{X}\{\cdot\}$ given by (3.2).

Proof From the orthogonality principle, it follows that for any estimator $\widehat{X}\{t\}$ of $X\{t\}$ based on an encoded version of $Y_T\{\cdot\}$ we have

$$\frac{1}{T}\int_{-T/2}^{T/2} \mathbb{E}\left(X\{t\} - \widetilde{X}\{t\}\right)^2 dt + \frac{1}{T}\int_{-T/2}^{T/2} \mathbb{E}\left(\widetilde{X}\{t\} - \widehat{X}\{t\}\right)^2 dt. \tag{3.8}$$

Therefore, (i) follows from (3.8) and from the definition of the operational DRF (3.1). In order to prove (ii), first note that

$$\begin{aligned}\mathsf{mmse}(X_T|Y_T) &= \frac{1}{T}\int_{-T/2}^{T/2} \mathbb{E}\left(X\{t\} - \widetilde{X}_T\{t\}\right)^2 dt = \sigma_X^2 - \frac{1}{T}\int_{-T/2}^{T/2} \mathbb{E}\left[\widetilde{X}_T^2\{t\}\right] dt \\ &= \sigma_X^2 - \frac{1}{T}\int_{-T/2}^{T/2} \mathbb{E}\left[\widetilde{X}^2\{t\}\right] dt + \frac{1}{T}\int_{-T/2}^{T/2} \mathbb{E}\left[\widetilde{X}^2\{t\} - \widetilde{X}_T\{t\}\right] dt. \end{aligned} \tag{3.9}$$

Since $\widetilde{X}_T\{\cdot\}$ converges in law to $\widetilde{X}\{\cdot\}$ and since the second moment of $\widetilde{X}_T\{t\}$ is bounded by that of $X\{\cdot\}$, the second moment of $\widetilde{X}\{\cdot\}$ is also bounded. As a result, the last term in the RHS of (3.9) goes to zero as T goes to infinity. In addition, asymptotic mean stationarity of $\widetilde{X}\{\cdot\}$ ensures the existence of the limit

$$\lim_{T\to\infty}\frac{1}{T}\int_{-T/2}^{T/2}\mathbb{E}\widetilde{X}^2\{t\}dt,$$

so that $\mathsf{mmse}(X|Y)$ is well defined and $\mathsf{mmse}(X_T|Y_T)\to\mathsf{mmse}(X|Y)$. For the second term in (3.7), we write

$$\frac{1}{T}\int_{-T/2}^{T/2}\mathbb{E}\left(\widetilde{X}_T\{t\}-\widehat{X}\{t\}\right)^2dt=\frac{1}{T}\int_{-T/2}^{T/2}\mathbb{E}\left(\widetilde{X}_T\{t\}-\widetilde{X}\{t\}+\widetilde{X}\{t\}-\widehat{X}\{t\}\right)^2dt$$
$$=\frac{1}{T}\int_{-T/2}^{T/2}\mathbb{E}\left(\widetilde{X}\{t\}-\widehat{X}\{t\}\right)^2dt+\frac{1}{T}\int_{-T/2}^{T/2}\mathbb{E}\left(\widetilde{X}_T\{t\}-\widetilde{X}\{t\}\right)^2dt$$
$$+\frac{1}{T}\int_{-T/2}^{T/2}\mathbb{E}\left(\widetilde{X}_T\{t\}-\widetilde{X}\{t\}\right)\left(\widetilde{X}\{t\}-\widehat{X}\{t\}\right)dt.$$

Convergence in law of $X_T\{\cdot\}$ to $\widetilde{X}(\cdot)$ and bounded second moments of both implies that the last two terms go to zero as $T\to\infty$. Thus, it is enough to consider the encoding of the limiting process $\widehat{X}\{\cdot\}$. Since the latter is assumed to be asymptotic mean stationary, the convergence of the OPTA to an information DRF representation analogous to (3.2) follows from [66, Ch. 11]. □

We note that the simplification in evaluating $\delta_{X|Y}(R)$ provided by Theorem 3.1 is twofold. First, it decomposes the indirect source coding problem into a minimal MSE estimation problem and a standard source coding problem. Second, it provides general conditions for the expression in the RHS of (3.5) to admit an information DRF representation. In this last sense, Theorem 3.1 is an extension of [15] and [16]. In the next section we consider various examples for using Theorem 3.1 where the processes $X\{\cdot\}$ and $Y\{\cdot\}$ are jointly Gaussian.

3.3 Indirect DRF in Gaussian Settings

In this section we consider the indirect source coding setting of Figure 3.3 under various special cases in which the distortion is quadratic and the processes $X\{\cdot\}$ and $Y\{\cdot\}$ are jointly Gaussian. In addition to demonstrating the usefulness of Theorem 3.1, these examples are later used in characterizing the minimal distortion in our general ADX setting of Figure 1.2.

3.3.1 Gaussian stationary source corrupted by noise

Consider the scenario where the source process $X(\cdot)$ is a Gaussian stationary process and the encoder observes the process $Y(t)=X(t)+\eta(t)$, where $\eta(\cdot)$ is a Gaussian stationary noise independent of $X(\cdot)$. In this scenario, the MMSE estimator of $X(\cdot)$ from $Y(\cdot)$ is given by the Wiener filter, which is a stationary Gaussian

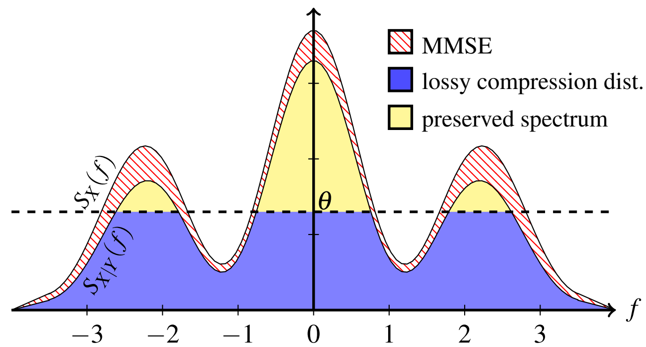

Figure 3.4: Water-filling interpretation of (3.10). The distortion is the sum of the MMSE and the lossy compression error.

process [53] with PSD

$$S_{\widetilde{X}}(f) = S_{X|Y}(f) = \frac{S_X^2(f)}{S_X(f)+S_\eta(f)}.$$

In particular, the asymptotic MMSE is given by

$$\mathsf{mmse}(X|Y) = \sigma_X^2 - \int_{-\infty}^{\infty} S_{X|Y}(f)df = \int_{-\infty}^{\infty} \frac{S_X(f)S_\eta(f)}{S_X(f)+S_\eta(f)}df,$$

and the DRF of $\widetilde{X}(\cdot)$ is obtained by (3.3). In the setting of Theorem 3.1, the asymptotic law of (3.6) converges to the output of the Wiener filter, which is stationary Gaussian, and hence asymptotic mean stationary [65]. It therefore follows from (3.7) that the indirect DRF of $X(\cdot)$ given $Y(\cdot)$ is given by the following expression, first derived in [15]:

$$D_{X|Y}(R_\theta) = \int_{-\infty}^{\infty} S_{X|Y}(f)df + \int_{-\infty}^{\infty} \min\left\{\theta, S_{X|Y}(f)\right\}df \tag{3.10a}$$

$$= \sigma_X^2 - \int_{-\infty}^{\infty} \left[S_{X|Y}(f) - \theta\right]^+ df$$

$$R_\theta = \int_{-\infty}^{\infty} \log^+ \left[S_{X|Y}(f)/\theta\right] df, \tag{3.10b}$$

where $\log^+[x] = \log(\max\{1,x\})$, and $S_X(f)$, $S_\eta(f)$ are the PSD of the processes $X(\cdot)$ and $\eta(\cdot)$, respectively. A graphic interpretation of the expression (3.10) as the combination of a reverse water-filling term and an MMSE term is illustrated in Figure 3.4.

3.3.2 Discrete-time jointly Gaussian vector processes

Consider a discrete-time and M-dimensional vector valued stationary process $\mathbf{X}[\cdot]$ with PSD matrix $\mathbf{S_X}\left(e^{2\pi i\phi}\right)$. The encoder observes the process $Y[\cdot]$, which is a discrete-time and P-dimensional vector valued stationary process with PSD matrix $\mathbf{S_Y}\left(e^{2\pi i\phi}\right)$. Assume moreover that $\mathbf{X}[\cdot]$ and $\mathbf{Y}[\cdot]$ are jointly Gaussian and stationary with cross PSD matrix $\mathbf{S}_{\mathbf{XY}}\left(e^{2\pi i\phi}\right)$ defined by

$$\left(\mathbf{S}_{\mathbf{XY}}\left(e^{2\pi i\phi}\right)\right)_{m,p} \triangleq \sum_{n\in\mathbb{Z}} \mathbb{E}\left[X_i[n]Y_j[0]\right] e^{-2\pi i\phi n}, \quad m=1,\ldots,M,\ p=1,\ldots P.$$

Then from linear estimation theory [53], it follows that the MMSE estimate of $\mathbf{X}[\cdot]$ from $\mathbf{Y}[\cdot]$ is given by the discrete-time vector Wiener filter:

$$\widetilde{\mathbf{X}}[n] = \sum_{k\in\mathbb{Z}} \mathbf{G}[n-k]\mathbf{Y}[k],$$

where the Fourier transform of the process defining the $M\times P$ matrix $\mathbf{G}$ is given by

$$\hat{\mathbf{G}}\left(e^{2\pi i\phi}\right) = \mathbf{S}_{\mathbf{XY}}\left(e^{2\pi i\phi}\right)\mathbf{S}_{\mathbf{Y}}^{-1}\left(e^{2\pi i\phi}\right).$$

Therefore, $\widetilde{X}[\cdot]$ is a discrete-time Gaussian stationary process with PSD matrix

$$\mathbf{S}_{\mathbf{X}|\mathbf{Y}} = \mathbf{S}_{\mathbf{XY}}\left(e^{2\pi i\phi}\right)\mathbf{S}_{\mathbf{Y}}^{-1}\left(e^{2\pi i\phi}\right)\mathbf{S}_{\mathbf{YX}}\left(e^{2\pi i\phi}\right),$$

and the MMSE in Wiener filtering is given by

$$\mathsf{mmse}(\mathbf{X}|\mathbf{Y}) = \frac{1}{M}\int_{-\frac{1}{2}}^{\frac{1}{2}} \mathrm{Tr}\,\mathbf{S}_{\mathbf{X}}\left(e^{2\pi i\phi}\right)d\phi - \frac{1}{M}\int_{-\frac{1}{2}}^{\frac{1}{2}} \mathrm{Tr}\,\mathbf{S}_{\mathbf{X}|\mathbf{Y}}\left(e^{2\pi i\phi}\right)d\phi.$$

The DRF of the process $\widetilde{\mathbf{X}}[\cdot]$ is obtained by the counterpart of Pinsker's expression (3.3) for Gaussian stationary vector sources as derived in [67, Eq. (20) and (21)]:

$$D_{\widetilde{\mathbf{X}}}(\bar{R}_\theta) = \frac{1}{M}\sum_{i=1}^{M}\int_{-\frac{1}{2}}^{\frac{1}{2}} \min\left\{\lambda_i\left(e^{2\pi i\phi}\right),\theta\right\}d\phi \tag{3.11a}$$

$$\bar{R}_\theta = \sum_{i=1}^{M}\int_{-\frac{1}{2}}^{\frac{1}{2}} \frac{1}{2}\log^+\left[\lambda_i\left(e^{2\pi i\phi}\right)/\theta\right]d\phi, \tag{3.11b}$$

where $\lambda_1\left(e^{2\pi i\phi}\right),\ldots,\lambda_M\left(e^{2\pi i\phi}\right)$ are the eigenvalues of the matrix $\mathbf{S}_{\mathbf{X}|\mathbf{Y}}\left(e^{2\pi i\phi}\right)$ and where $\bar{R}$ is measured in bits per source symbol (unlike R, which denotes bits per unit time).

In the setting of Theorem 3.1, the law of the process $\widetilde{X}_T[\cdot]$ of (3.6) converges to the law of the stationary Gaussian process $\widetilde{\mathbf{X}}[\cdot]$ at the output of the Wiener filter. Therefore, the indirect DRF of $\mathbf{X}[\cdot]$ given $\mathbf{Y}[\cdot]$ follows from (3.7):

Theorem 3.2 *Let $\mathbf{X}[\cdot]$ be an M-dimensional, vector-valued Gaussian stationary process and let $\mathbf{Y}[\cdot]$ be another vector valued process, such that $\mathbf{X}[\cdot]$ and $\mathbf{Y}[\cdot]$ are jointly Gaussian and stationary. The indirect* *distortion-rate function of* $\mathbf{X}[\cdot]$ *given* $\mathbf{Y}[\cdot]$ *under quadratic distortion is given by*

$$\begin{aligned} D_{\mathbf{X}|\mathbf{Y}}(\bar{R}_\theta) =& \mathsf{mmse}(\mathbf{X}|\mathbf{Y}) + \frac{1}{M}\sum_{i=1}^{M}\int_{-\frac{1}{2}}^{\frac{1}{2}} \min\left\{\lambda_i\left(e^{2\pi i\phi}\right), \theta\right\} d\phi \\ =& \frac{1}{M}\int_{-\frac{1}{2}}^{\frac{1}{2}} \operatorname{Tr} \mathbf{S_X}\left(e^{2\pi i\phi}\right) d\phi - \frac{1}{M}\sum_{i=1}^{M}\int_{-\frac{1}{2}}^{\frac{1}{2}} \left[\lambda_i\left(e^{2\pi i\phi}\right) - \theta\right]^+ d\phi, \end{aligned}$$

$$\bar{R}_\theta = \sum_{i=1}^{M}\int_{-\frac{1}{2}}^{\frac{1}{2}} \frac{1}{2}\log^+\left[\lambda_i\left(e^{2\pi i\phi}\right)/\theta\right] d\phi,$$

where $\lambda_1\left(e^{2\pi i\phi}\right), \ldots, \lambda_M\left(e^{2\pi i\phi}\right)$ *are the eigenvalues of the matrix* $\mathbf{S}_{\mathbf{X}|\mathbf{Y}}\left(e^{2\pi i\phi}\right)$.

Proof Since the MMSE estimate of $\mathbf{X}(\cdot)$ given $\mathbf{Y}_T(\cdot)$ converges to $\mathsf{mmse}(\mathbf{X}|\mathbf{Y})$, the theorem is an immediate consequence of (3.11) and (3.7). □

3.3.3 Jointly Gaussian i.i.d. vector sources

Another example for an indirect source coding problem involving a pair of Gaussian sources is as follows: each symbol from the sources $X\{\cdot\}$ and $Y\{\cdot\}$ is an i.i.d. vector $\mathbf{X}$ and $\mathbf{Y}$ with covariance matrices Σ_X, Σ_Y, respectively. The cross-covariance between $\mathbf{X}$ and $\mathbf{Y}$ is Σ_{XY}. Although this example is relatively simple, it provides intuition for the kind of optimization required in subsequent chapters to derive the optimal sampling structure in the ADX. The indirect DRF in this scenario is obtained as a special case of Theorem 3.2 by setting $\mathbf{S_X}\left(e^{2\pi i\phi}\right) \equiv \Sigma_X$, $\mathbf{S_Y}\left(e^{2\pi i\phi}\right) \equiv \Sigma_Y$, and $\Sigma_{XY}\left(e^{2\pi i\phi}\right) \equiv \Sigma_{XY}$. The result is as follows:

$$\begin{aligned} D_{\mathbf{X}|\mathbf{Y}}(\bar{R}_\theta) &= \sum_{i=1}^{n}\left(\lambda_i\left(\Sigma_X\right) - \lambda_i\left(\Sigma_{X|Y}\right)\right) + \sum_{i=1}^{n}\min\left\{\theta, \lambda_i\left(\Sigma_{X|Y}\right)\right\} \\ &= \sum_{i=1}^{n}\lambda_i\left(\Sigma_X\right) - \sum_{i=1}^{n}\left[\lambda_i\left(\Sigma_{X|Y}\right) - \theta\right]^+, \end{aligned} \tag{3.13a}$$

$$\bar{R}_\theta = \frac{1}{2}\sum_{i=1}^{n}\log^+\left[\lambda_i\left(\Sigma_{X|Y}\right)/\theta\right]. \tag{3.13b}$$

3.3.4 Minimizing a combined MMSE and water-filling expression

Before concluding this section, we consider the minimization of expressions of the form (3.10) and (3.13) over the PSD or the eigenvalues of the covariance matrix of the estimator.

Proposition 3.3 *Fix $R > 0$ and consider the function $D : \mathbb{R}^n \to \mathbb{R}$ defined by*

$$D(x_1,\dots,x_n) = -\sum_{i=1}^{n} [x_i - \theta]^+$$
$$R = \frac{1}{2}\sum_{i=1}^{n} \log^+ [x_i/\theta].$$

Then D is non-increasing in any of its arguments.

Proof For any i such that $x_i > \theta$, the partial derivative of θ with respect to x_i equals $\theta/x_i < 1$. Therefore, the derivative of D with respect to x_i equals $-1 + \theta/x_i < 0$. If $X_i < \theta$ then the partial derivative of D with respect to x_i is zero. □

When an expression of the form (3.10) is considered, Proposition 3.3 takes the following form:

Proposition 3.4 *Fix $R > 0$ and set $A \subset \mathbb{R}$. For an integrable function f over A, define*

$$D(f) = -\int_A [f(x) - \theta]^+ \, dx$$
$$R = \frac{1}{2}\int_A \log^+ [f(x)/\theta] \, dx.$$

Let f and g be two integrable functions such that

$$\int_A f(x)dx \le \int_A g(x)dx.$$

Then $D(g) \le D(f)$.

Proof Since f and g are integrable, we approximate their integrals by sums over a finite number of elements. Proposition 3.3 applied to this sum leads to the desired result. □

3.4 Chapter Summary

We considered the minimal MSE distortion in recovering a random process from an encoded version of another process, statistically correlated with the first one. This problem is known as the indirect source coding problem, for which the optimal tradeoff between code rate and distortion is characterized by the indirect DRF. Under quadratic distortion, the indirect DRF can be decomposed into an estimation problem under a MSE criterion and a standard source coding problem with respect to the estimator. In this chapter we provided new conditions on the posterior law of the process given the observations such that the above decomposition leads to a computable expression of the indirect DRF. These conditions permit the statistical relation between the source and the observations at the encoder to depend on the time horizon, as long as this relation converges in law to an asymptotic mean stationary distribution. We used this decomposition to evaluate the indirect DRF in closed form for recovering a Gaussian processes from jointly Gaussian observations that may be noisy. Since the input signal in the ADX setting in Figure. 1.2 is stationary and Gaussian, this last derivation provides a lower bound for the distortion in ADX when noise is added to the source signal before sampling.

In the following chapter we combine the source coding problem discussed in this chapter with the MMSE problem of Chapter 2 to characterize the fundamental distortion limit in ADX.

Chapter 4

Combined Sampling and Source Coding

In this chapter we provide the fundamental distortion limit in ADX by combining the problem of estimating signals from their noisy sampled version, considered in Chapter 2, with the source coding problem of encoding signals subject to a bitrate constraint considered in Chapter 3.2. Specifically, in Section 4.1, we formalize the ADX setting as a combined sampling and source coding problem with an additive noise. Therefore, the minimal distortion in recovering the analog signal is due to the joint effect of sampling, lossy compression, and independent noise. In Section 4.2, we solve this combined problem in the simple case of a single-branch uniform sampler and explore properties of this solution through various examples. In Section 4.3, we then generalize our solution to multi-branch sampling. We also consider an optimization over the pre-sampling operations in both cases and derive an achievable lower bound on the distortion in ADX under time-invariant uniform sampling. In Section 4.4 we show that this lower bound on the distortion in ADX holds even under the most general case of an arbitrary bounded linear sampler.

4.1 Combined Sampling and Source Coding

In this section we formalize the ADX setting as a combined problem of sampling and source coding. Our formalization combines the bounded linear sampling system of Chapter 2 with the indirect source coding setting of Chapter 3.

4.1.1 Combined sampling and source coding setting

The combined sampling and source coding setting is described in Figure 4.1. The source signal $X(\cdot)$ is a zero-mean Gaussian stationary process with an $\mathbf{L}_1(\mathbb{R})$ PSD $S_X(f)$ and variance σ_X^2. The process $X(\cdot)$ is corrupted by a Gaussian stationary noise process $\eta(\cdot)$ with PSD $S_\eta(f)$. The *sampler* $S = S(H,\Lambda)$ is a bounded linear sampler with a pre-sampling operation H and a sampling set Λ, as defined in Section 2.2. The input to the sampler is the process $X(\cdot)+\eta(\cdot)$, and the output of this sampler at time T is an N dimensional vector Y_T,

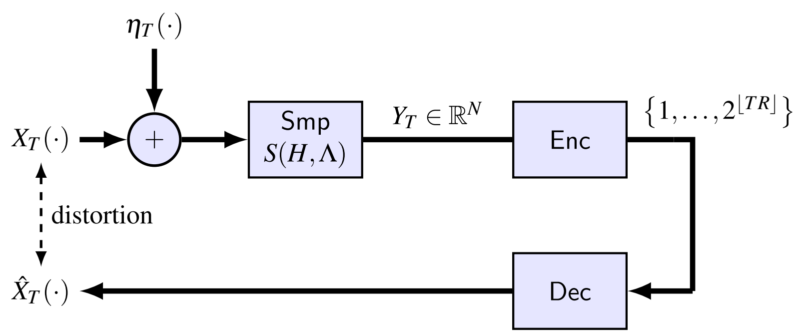

Figure 4.1: Combined sampling and source coding setting.

where $N = |\Lambda \cap [-T/2, T/2]|$. The *encoder* maps the vector Y_T to an index i in the set $\{1,\ldots,2^{\lfloor RT \rfloor}\}$. The decoder, upon receiving the index i, produces an estimate $\widehat{X}_T^i(\cdot)$ of $X_T(\cdot)$. The *distortion* between the source signal and its estimate is measured according to a MSE criterion. Namely, the distortion between a source waveform $x_T(\cdot)$ and a reconstruction waveform $\widehat{x}_T(\cdot)$ is given by

$$\frac{1}{T}\int_{-T/2}^{T/2} \left(x_T(t) - \widehat{x}_T(t)\right)^2 dt.$$

For the bounded linear sampler S limited to a time horizon T, our goal is to minimize the expected distortion

$$D_T = \inf_{\mathsf{Enc-Dec}} \frac{1}{T}\int_{-T/2}^{T/2} \mathbb{E}\left(X_T(t) - \widehat{X}_T(t)\right)^2 dt,$$

where the infimum is with respect to all encoders and decoders operating to and from a set of $2^{\lfloor TR \rfloor}$ elements. The minimal distortion in ADX associated with S is the indirect DRF of $X(\cdot)$ given the vector of samples Y_T, defined as

$$D_S(R) \triangleq \liminf_{T\to\infty} D_T. \tag{4.1}$$

4.1.2 Basic properties of $D_S(R)$

Below is a list of various properties of the function $D_S(R)$ that follow from its definition:

Proposition 4.1 *For any bitrate R and a bounded linear sampler S:*

(i) $D_S(R) \le \sigma_X^2$ *and* $D_S(R) \to \sigma_X^2$ *as* $R \to 0$.

(ii) $D_S(R) \ge \mathsf{mmse}_S$ *and* $D_S(R) \to \mathsf{mmse}_S$ *as* $R \to \infty$, *where* mmse_S *is defined in* (2.9).

(iii) $D_S(R) \ge D_{X|X+\eta}(R)$, *where* $D_{X|X+\eta}(R)$ *is the indirect DRF of* $X(\cdot)$ *given* $X(\cdot)+\eta(\cdot)$ *of* (3.10).

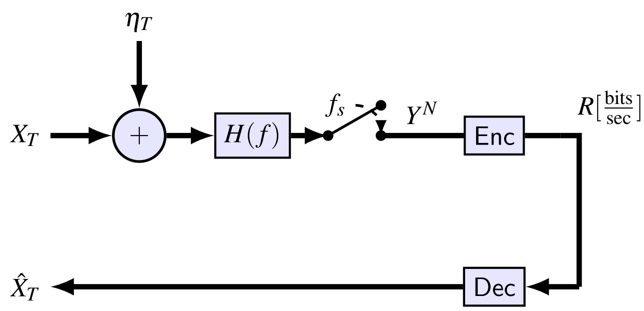

Figure 4.2: Combined sampling and source coding with a single-branch uniform sampler.

Proof (i) follows by taking $\widehat{X}_T(\cdot) \equiv 0$ in the definition of $D_S(R)$ in (4.1). As $R \to 0$, no information on $X_T(\cdot)$ is provided and the mean, which is assumed to be zero, is the best estimate for $X_T(\cdot)$. (ii) As $R \to \infty$, the encoder can encode the samples Y_T with arbitrarily small distortion, hence the problem is reduced to the minimal MSE estimation problem of Section 2.4. (iii) follows since any optimal encoding of Y_T for describing $X(\cdot)$ is no better than the optimal encoding of $X(\cdot)+\eta(\cdot)$ to describe $X(\cdot)$. $\square$

For the special case of a SB sampler with sampling rate f_s, the function $D_S(R)$ defines a manifold in the three-dimensional space (D, f_s, R). It follows from Proposition 4.1 that this manifold is bounded from above by the plane $D = \sigma_X^2$, and from below by the two cylinders $D = \mathsf{mmse}_{\mathsf{SB}(H)}(f_s)$ and $D = D_{X|X+\eta}(R)$. Next, we derive precise expressions for $D_S(R)$ for SB and MB uniform samplers.

4.2 Single Branch Uniform Sampling

We now consider the case where the sampler S is the SB uniform sampler of the form described in Figure 2.2 and in Section 2.3.1, leading to the combined sampling and source coding system illustrated in Figure 4.2. For this special class of samplers, we denote the function $D_S(R)$ by $D_{\mathsf{SB}(H)}(f_s)$.

Consider the MMSE estimator $\widetilde{X}_T(\cdot)$ of $X_T(\cdot)$ given $Y_T(\cdot)$. It was shown in Section 2.4.1 that the asymptotic law of this estimator converges to the law of the process $\widetilde{X}(\cdot)$ given by (2.20). Since $\widetilde{X}(\cdot)$ is a cyclostationary process, it is also asymptotic mean stationary [65, Exc. 6.3.1]. Hence, the conditions of (ii) in Theorem 3.1 hold and

$$D_{\mathsf{SB}(H)}(f_s, R) = \mathsf{mmse}_{\mathsf{SB}(H)}(f_s) + D_{\widetilde{X}}(R). \tag{4.2}$$

A closed form expression for $\mathsf{mmse}_{\mathsf{SB}(H)}(f_s)$ is given by Proposition 2.1 so it is only left to evaluate the second term in (4.2), which is the DRF of the process $\widetilde{X}(\cdot)$. Note that for $f_s > f_{\mathsf{Nyq}}$, as explained in Section 2.3.1, the MMSE estimator of $X(\cdot)$ from $Y[\cdot]$ coincides with the estimator resulting from the Wiener filter for estimating $X(\cdot)$ from the output of the filter $H(f)$. This last estimator is stationary with PSD $\frac{S_X^2(f)}{S_X(f)+S_\eta(f)}$. Therefore, we conclude that for $f_s > f_{\mathsf{Nyq}}$,

$$D_{\mathsf{SB}(H)}(f_s, R) = D_{X|X+\eta}(R),$$

where $D_{X|X+\eta}(R)$ is given by (3.10).

Next, we consider the DRF of $\widetilde{X}(\cdot)$ when f_s is below the Nyquist rate of $X(\cdot)$. In this case, $\widetilde{X}(\cdot)$ is not a stationary process but rather cyclostationary with a pulse-amplitude structure [52]. A general expression for the DRF of Gaussian cyclostationary processes is given in Appendix A. For the special case of cyclostationary processes with pulse-amplitude structure, the following theorem is obtained from the results of Appendix A:

Theorem 4.2 (DRF of PAM-modulated signals) *Let $X_{\mathsf{PAM}}(\cdot)$ be defined by*

$$X_{\mathsf{PAM}}(t) = \sum_{n \in \mathbb{Z}} U(n/f_0) p(t - n/f_0), \quad t \in \mathbb{R}, \tag{4.3}$$

where $U(\cdot)$ is a Gaussian stationary process with PSD $S_U(f)$ and $p(t)$ is an analog deterministic signal with $\int_{-\infty}^{\infty} |p(t)|^2 dt < \infty$ and Fourier transform $P(f)$. Assume, moreover, that the covariance function $\mathbb{E}[X_{\mathsf{PAM}}(t+\tau) X_{\mathsf{PAM}}(t)]$ of $X_{\mathsf{PAM}}(\cdot)$ is Lipschitz continuous in τ. The distortion-rate function of $X_{\mathsf{PAM}}(\cdot)$ is given by

$$D(\theta) = f_0 \int_{-\frac{f_0}{2}}^{\frac{f_0}{2}} \min\left\{\widetilde{S}(f), \theta\right\} df \tag{4.4a}$$

$$R(\theta) = \frac{1}{2} \int_{-\frac{f_0}{2}}^{\frac{f_0}{2}} \log^+ \left[\widetilde{S}(f)/\theta\right] df, \tag{4.4b}$$

where

$$\widetilde{S}(f) \triangleq \sum_{k \in \mathbb{Z}} |P(f - k f_s)|^2 S_U(f - k f_s). \tag{4.5}$$

Proof See Appendix A. □

Using Theorem 4.2 we obtain the following:

Theorem 4.3 *The minimal distortion in the combined sampling and source coding problem with a SB uniform sampler of sampling rate f_s and pre-sampling operation $H(f)$ is given by*

$$D_{\mathsf{SB}(H)}(f_s, R_\theta) = \mathsf{mmse}_{\mathsf{SB}(H)}(f_s) + \int_{-\frac{f_s}{2}}^{\frac{f_s}{2}} \min\left\{\widetilde{S}_{X|Y}(f), \theta\right\} df \tag{4.6a}$$

$$= \sigma_X^2 - \int_{-\frac{f_s}{2}}^{\frac{f_s}{2}} \left[\widetilde{S}_{X|Y}(f) - \theta\right]^+ df$$

$$R_\theta = \frac{1}{2} \int_{-\frac{f_s}{2}}^{\frac{f_s}{2}} \log^+ \left[\widetilde{S}_{X|Y}(f), \theta\right] df, \tag{4.6b}$$

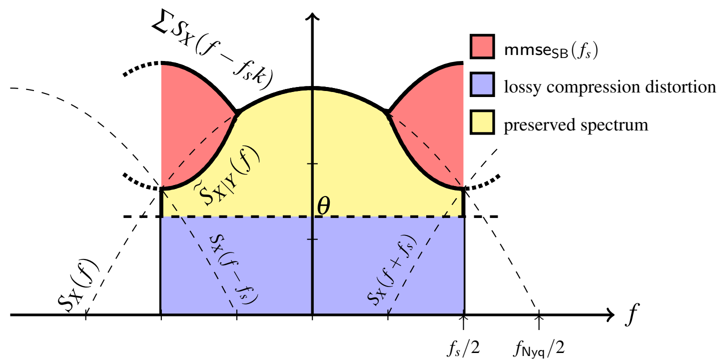

Figure 4.3: Water-filling interpretation of (4.6) with an all pass pre-sampling filter ($|H(f)| \equiv 1$): The function $\sum_{k\in\mathbb{Z}} S_X(f-f_sk)$ is the aliased PSD that represents the full energy of the original signal within the discrete-time spectrum interval $(-f_s/2, f_s/2)$. The part of the energy recovered by the MMSE estimator of (2.20) is $\widetilde{S}_{X|Y}(f)$. The distortion due to lossy compression is obtained by reverse water-filling over the recovered energy according to (4.6a). The overall distortion $D_{\mathsf{SB}(H)}(f_s,R)$ is the sum of the MMSE due to sampling and the lossy compression distortion.

where

$$\widetilde{S}_{X|Y}(f) = \frac{\sum_{k\in\mathbb{Z}} S_X(f-f_sk)^2 \left|H(f-f_sk)\right|^2}{\sum_{n\in\mathbb{Z}} \left|H(f-f_sn)\right|^2 S_{X+\eta}(f-f_sn)}. \tag{4.7}$$

Proof We use Theorem 4.2 with $U(t) = \int_{-\infty}^{\infty} (X(\tau)+\eta(\tau))h(t-\tau)d\tau$. The PSD of the process $U(\cdot)$ is given by

$$S_U(f) = (S_X(f)+S_\eta(f))\left|H(f)\right|^2.$$

From Theorem 4.2 it follows that the DRF of $\widetilde{X}(\cdot)$ of (2.20) is given by

$$D_\theta = \int_{-\frac{f_s}{2}}^{\frac{f_s}{2}} \min\left\{\widetilde{S}_{X|Y}(f),\theta\right\} df$$

$$R_\theta = \frac{1}{2}\int_{-\frac{f_s}{2}}^{\frac{f_s}{2}} \log^+\left[\widetilde{S}_{X|Y}(f)/\theta\right] df,$$

hence (4.6) follows from (4.2). □

4.2.1 Discussion

A graphical interpretation of (4.6) is given in Figure 4.3, under the assumption that $\eta(\cdot)$ equals zero and the pre-sampling filter $H(f)$ does not block any part of the signal energy (for example, if $H(f)$ is an all pass filter). The band $(-f_s/2, f_s/2)$ represents the spectrum of the discrete-time process $Y[\cdot]$ in the analog domain, and the entire signal's energy σ_X^2 is obtained by integrating the aliased version of $S_X(f)$ in this band. Consequently, the function $\widetilde{S}_{X|Y}(f)$ represents the part of the energy that is recovered by the MMSE estimator $\widetilde{X}(\cdot)$. The parameter θ is the water level that couples the distortion expression (4.6a) with the bitrate expression (4.6b) in a similar manner as in Pinsker's water-filling expression (3.3). Theorem 4.3 shows that the function $D_{\mathsf{SB}(H)}(f_s,R)$ is given by the sum of two terms, both of which depend on $\widetilde{S}_{X|Y}(f)$: (1) the MMSE in this estimation, given by integrating the difference

$$\sum_{k\in\mathbb{Z}}\left(S_X(f-f_sk)-\widetilde{S}_{X|Y}(f-f_sk)\right)=\sum_{k\in\mathbb{Z}}S_X(f-f_sk)-\widetilde{S}_{X|Y}(f)$$

over $(-f_s/2,f_s/2)$, and (2) water-filling over $\widetilde{S}_{X|Y}(f)$. By comparing (4.6) with (3.10), we conclude that the function $\widetilde{S}_{X|Y}(f)$ plays the role of the PSD of the MMSE estimator of $X(t)$ from $Y[\cdot]$. Recall, however, that this estimator is not stationary in general, and therefore its PSD is not well-defined. In fact, as follows from Section 2.4, the function $\widetilde{S}_{X|Y}(f)$ can be seen as the average of the polyphase components of $\widetilde{X}(\cdot)$. This last interpretation for $\widetilde{S}_{X|Y}(f)$ is explained in Appendix A and generalized to derive the DRF of an arbitrary Gaussian cyclostationary process.

The function $\widetilde{S}_{X|Y}(f)$ depends on the sampling rate, the pre-sampling filter $H(f)$, and the spectral densities $S_X(f)$ and $S_\eta(f)$ – but is independent of R. By fixing R and considering a small change in $\widetilde{S}_{X|Y}(f)$ such that the overall estimated energy

$$\int_{-\frac{f_s}{2}}^{\frac{f_s}{2}}\widetilde{S}_{X|Y}(f)df \tag{4.8}$$

is increased, then it follows from (4.6b) that θ is also increased to maintain the same fixed source coding rate R. On the other hand, the expression for $D_{\mathsf{SB}(H)}(f_s,R)$ in (4.6a) exhibits a negative linear dependency on (4.8). In this interplay between the two terms in (4.6), the negative linear dependency in $\widetilde{S}_{X|Y}(f)$ is stronger then a logarithmic dependency of θ in (4.8), and the distortion reduces with an increment in (4.8), as implied by Proposition 3.4. We therefore conclude:

Corollary 4.4 *For a fixed $R\geq 0$, minimizing $D_{\mathsf{SB}(H)}(f_s,R)$ is equivalent to maximizing*

$$\int_{-\frac{f_s}{2}}^{\frac{f_s}{2}}\widetilde{S}_{X|Y}(f)df.$$

This last corollary will be useful in determining the pre-sampling filter that minimizes $D_{\mathsf{SB}(H)}(f_s,R)$ in Section 4.2.3.

4.2.2 Examples

As an example for using Theorem 4.3, consider the case where the the spectrum of $X(\cdot)$ equals

$$S_X(f) = S_\Pi(f) = \begin{cases} \frac{\sigma_X^2}{2W} & |f| \leq W, \\ 0 & \text{otherwise}, \end{cases} \tag{4.9}$$

for some $W > 0$. In addition, assume that the noise is constant over the band $|f| \leq W$ with intensity $\sigma_\eta^2 = \gamma^{-1}\sigma_X^2/(2W)$, where $\gamma > 0$ is the signal-to-noise ratio (SNR) in this model. For all frequencies $f \in [-f_s/2, f_s/2]$, we have

$$\widetilde{S}_{X|Y}(f) = \frac{\sum_{k\in\mathbb{Z}} S_X^2\left(f - f_s k\right)}{\sum_{k\in\mathbb{Z}}\left(S_X\left(f - f_s k\right) + S_\eta\left(f - f_s k\right)\right)} = \frac{\sigma_X^2}{2W}\begin{cases} \frac{\gamma}{1+\gamma} & |f| < W, \\ 0 & |f| \geq W. \end{cases}$$

For this case, (4.6) is given by

$$D_{\mathsf{SB}(H)}\left(f_s, R_\theta\right) = \sigma_X^2 \begin{cases} \left[1 - \frac{f_s}{2W}\right]^+ + \frac{\theta f_s}{\sigma_X^2} & \frac{\theta}{\sigma_X^2} \leq \min\left\{\frac{f_s \gamma}{2W(1+\gamma)}, 1\right\}, \\ 1 & \text{otherwise}, \end{cases}$$

and

$$R_\theta = \begin{cases} \frac{f_s}{2}\log\left(\frac{\sigma_X^2\gamma}{2W\theta(1+\gamma)}\right) & 0 \leq \frac{\theta}{\sigma^2}(1+\gamma^{-1}) < \frac{f_s}{2W} < 1, \\ W\log\left(\frac{\sigma_X^2\gamma}{2W\theta(1+\gamma)}\right) & 0 \leq \frac{\theta}{\sigma_X^2}(1+\gamma^{-1}) < 1 \leq \frac{f_s}{2W}, \\ 0 & \text{otherwise}. \end{cases}$$

The parametric expressions above can be written in a single equation as

$$D_{\mathsf{SB}(H)}\left(f_s, R\right) = \sigma_X^2 \begin{cases} 1 - \frac{f_s}{2W} + \frac{f_s}{2W}\frac{\gamma}{1+\gamma}2^{\frac{-2R}{f_s}} & \frac{f_s}{2W} < 1, \\ \frac{1}{1+\gamma} + \frac{\gamma}{1+\gamma}2^{-\frac{R}{W}} & \frac{f_s}{2W} \geq 1. \end{cases} \tag{4.10}$$

Expression (4.10) has a very intuitive structure: for sampling rates below the Nyquist frequency of the signal, the distortion as a function of the bitrate increases by a constant factor due to the error as a result of non-optimal sampling. This factor completely vanishes for f_s greater than the Nyquist frequency of the signal, in which case $D_{\mathsf{SB}}\left(f_s, R\right)$ equals the indirect DRF of the process $X(\cdot)$ given $X(\cdot) + \eta(\cdot)$, which, by (3.10), equals

$$D_{X|X+\eta}(R) = \sigma_X^2\left(\frac{1}{1+\gamma} + \frac{\gamma}{1+\gamma}2^{-R/W}\right).$$

Expression (4.10) is depicted in Figure 4.4 for $\gamma = 5$ and $\gamma \to \infty$.

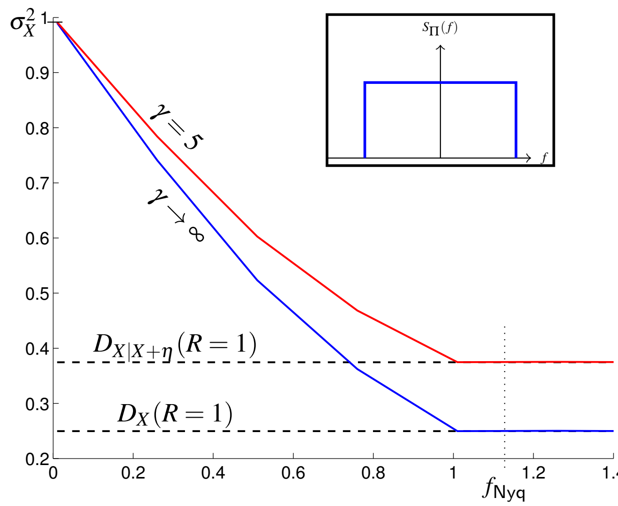

Figure 4.4: The function $D_{\mathsf{SB}(H)}(f_s, R)$ versus f_s for the process with rectangular PSD of (4.10), bandwidth 0.5, and an all-pass filter $|H(f)| \equiv 1$. The lower curve corresponds to zero noise, and the upper curve corresponds to signal-to-noise ration $\gamma = 5$. The dashed lines represent the indirect DRF of the source given the noisy process before sampling and when this noise is zero. In this example, $f_{\mathsf{Nyq}} = f_{\mathsf{Lan}} = 1$.

As another example, consider the case where $S_X(f)$ has the band-pass structure

$$S_X(f) = \begin{cases} \frac{\sigma_X^2}{2} & 1 \le |f| \le 2, \\ 0 & \text{otherwise}, \end{cases} \tag{4.11}$$

and zero noise, i.e. $S_\eta(f) \equiv 0$. We also assume that the pre-sampling filter $H(f)$ does not block any spectral line of $S_X(f)$. As in the example of the PSD (4.9), we obtain that for any $f \in (-f_s/2, f_s/2)$, the function $\widetilde{S}_{X|Y}(f)$ either equals $\frac{\sigma^2}{2}$ or 0. Thus, in order to find $D_{\mathsf{SB}(H)}(f_s, R)$, we only need to know for which values of $f \in (-f_s/2, f_s/2)$ the function $\widetilde{S}_{X|Y}(f)$ vanishes. This leads to

$$D_{\mathsf{SB}(H)}(f_s, R) = \sigma_X^2 \begin{cases} 2^{-R} & 4 \le f_s, \\ 1 - \frac{f_s - 2}{2}\left(1 - 2^{-\frac{2R}{f_s-2}}\right) & 3 \le f_s < 4, \\ 1 - \frac{4-f_s}{2}\left(1 - 2^{-\frac{2R}{4-f_s}}\right) & 2 \le f_s < 3, \\ 1 - (f_s - 1)\left(1 - 2^{-\frac{R}{f_s-1}}\right) & 1.5 \le f_s < 2, \\ 1 - (2 - f_s)\left(1 - 2^{-\frac{R}{2-f_s}}\right) & 4/3 \le f_s < 1.5, \\ 1 - \frac{f_s}{2}\left(1 - 2^{-\frac{2R}{f_s}}\right) & 0 \le f_s < 4/3, \end{cases} \tag{4.12}$$

which is depicted in Figure 4.5 for two different values of R. It can be seen from Figure 4.5 that the function

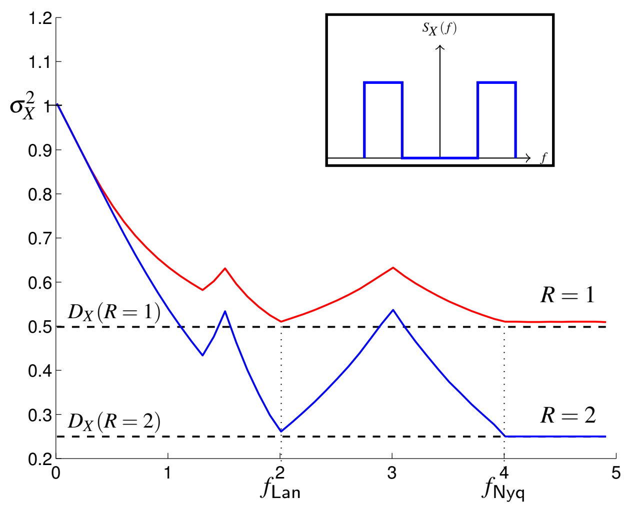

Figure 4.5: Illustration of (4.12): the function $D_{\mathsf{SB}(H)}(f_s,R)$ for the process with spectrum given in the small frame and two values of the source coding rate R.

$D_{\mathsf{SB}(H)}(f_s,R)$ is not necessarily decreasing in the sampling rate f_s. In fact, it follows from this figure that the function $D_{\mathsf{SB}}(f_s,R)$ coincides with $D_X(R)$ for $f_s=2$, indicating that sampling at this particular sub-Nyquist rate does not lead to additional distortion due to sampling compared to sampling at any super-Nyquist rate. This reduction in sampling rate compared to the Nyquist rate is the result of the particular bandpass structure of the process that allows sampling at the Landau rate with a uniform sampler. The technique to design sub-Nyquist samplers for sampling bandpass signals is known as *bandpass sampling* [68]. Note that since the Landau rate (i.e. the support of $S_X(f)$) in this example equals 2, Proposition 2.3 implies that it is impossible to sample $X(\cdot)$ at a sampling rate smaller than 2 without introducing additional distortion.

In the two examples above with the PSDs (4.9) and (4.11), the pre-sampling filter $H(f)$ has no effect on the function $D_{\mathsf{SB}(H)}(f_s,R)$ as long as the support of $S_X(f)$ is included in the support of $H(f)$. Intuitively, the reason for this indifference to the pre-sampling filter is the uniform energy distribution of the signal over the support of the PSD. This uniformity implies that all spectral lines of the signal contribute the same amount to the overall distortion, and hence all lines are estimated with equal importance. In particular, there is no benefit in attenuating or amplifying any part of the spectrum using a pre-sampling filter, since this does not change the SNR. As it turns out, the picture is different when the energy distortion is not uniform over the spectral support: for a general PSD $S_X(f)$ and each sampling rate f_s, a particular choice of $H(f)$ minimizes the function $D_{\mathsf{SB}(H)}(f_s,R)$.

4.2.3 Optimal pre-sampling filter

We now consider the pre-sampling filter $H(f)$ that minimizes (4.6) subject to a fixed bitrate R and sampling rate f_s. It follows from Corollary 4.4 that this minimization is equivalent to maximizing $\widetilde{S}_{X|Y}(f)$ for any f in the interval $(-f_s/2, f_s/2)$. In particular, since $\widetilde{S}_{X|Y}(f)$ is independent of R, the optimal pre-sampling filter is only a function of f_s and the spectrum of the signal and the noise. Recall that the optimal pre-sampling filter $H^\star(f)$ that maximizes $\widetilde{S}_{X|Y}(f)$ was already given in Theorem 2.5 in terms of the maximal aliasing free set associated with $\frac{S_X^2(f)}{S_X(f)+S_\eta(f)}$. This leads us to the following conclusion:

Theorem 4.5 *Given $f_s > 0$, the optimal pre-sampling filter $H^\star(f)$ that minimizes $D_{\mathsf{SB}(H)}(f_s,R)$, for all $R \geq 0$, is given by*

$$H^\star(f) = \begin{cases} 1 & f \in F_1^\star, \\ 0 & \text{otherwise}, \end{cases}$$

where $F_1^\star \in \mathsf{AF}(f_s)$ and satisfies

$$\int_{F_1^\star} \frac{S_X^2(f)}{S_{X+\eta}(f)} df = \int_{-\frac{1}{2}}^{\frac{1}{2}} \sup_k \frac{S_X^2(f - f_s k)}{S_{X+\eta}(f - f_s k)} df.$$

Moreover, the minimal distortion obtained using $H^\star$ is given by

$$\begin{aligned} D_{\mathsf{SB}}^\star(f_s, R_\theta) &\triangleq D_{\mathsf{SB}(H^\star)}(f_s, R_\theta)\, \mathsf{mmse}_{\mathsf{SB}}^\star(f_s) + \int_{F_1^\star} \min\left\{\frac{S_X^2(f)}{S_{X+\eta}(f)}\right\} df \qquad (4.13a) \\ &= \sigma_X^2 - \int_{F_1^\star} \left[\frac{S_X^2(f)}{S_{X+\eta}(f)} - \theta\right]^+ df \end{aligned}$$

$$R_\theta = \frac{1}{2}\int_{F^\star} \log^+ \left[\frac{S_X^2(f)}{S_{X+\eta}(f)}/\theta\right] df,$$

where $\mathsf{mmse}_{\mathsf{SB}}^\star(f_s)$ is given by (2.34).

Proof From Theorem 2.5 we conclude that the filter $H^\star(f)$ that maximizes $\widetilde{S}_{X|Y}(f)$ is given by the indicator function of the maximal aliasing free set $F_1^\star$. Moreover, with this optimal filter, (4.6) reduces to (4.13). □

We emphasize that even in the absence of noise, the optimal pre-sampling filter $H^\star(f)$ plays a crucial role in reducing distortion by preventing aliasing as described in Section 2.5. The effect of using the optimal pre-sampling filter on the function $D_{\mathsf{SB}(H)}(f_s,R)$ is illustrated in Figure 4.6.

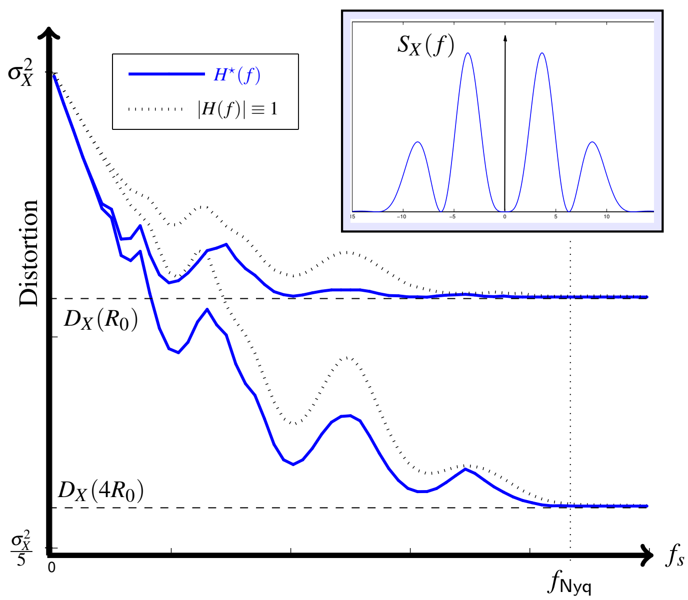

Figure 4.6: The functions $D^\star_{\mathsf{SB}}(f_s,R)$ and $D_{\mathsf{SB}(H)}(f_s,R)$, corresponding to using an optimal pre-sampling filter $H^\star(f)$ and an all-pass filter ($|H(f)| \equiv 1$), respectively, at two fixed values of R. Here $\eta(\cdot)$ is zero, and $S_X(f)$ is given in the small frame. The dashed lines represent the distortion-rate functions of the source corresponding to direct encoding or sampling above the Nyquist rate.

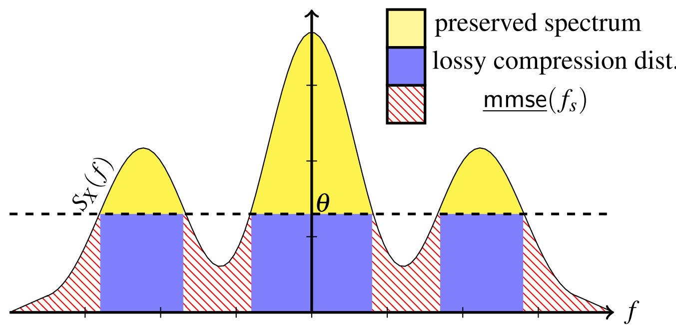

Figure 4.7: Water-filling interpretation of the function $\underline{D}(f_s,R)$. The overall distortion is the sum of the MMSE due to sampling and the lossy compression distortion. The set $F^\star$ is the support of the preserved spectrum.

4.2.4 The fundamental distortion limit

Analogous to the lower bound for the MMSE of (2.36), we now derive a lower bound on $D_{\mathsf{SB}(H)}(f_s,R)$ using a set $F^\star$ of Lebesgue measure f_s that satisfies

$$\int_{\mathscr{F}^\star} \frac{S_X^2(f)}{S_{X+\eta}(f)} df = \sup_{\mu(F)\leq f_s} \int_F \frac{S_X^2(f)}{S_{X+\eta}(f)} df. \tag{4.14}$$

Consider the function:

$$\underline{D}(f_s,R_\theta) = \underline{\mathsf{mmse}}(f_s,R) + \int_{F^\star} \min\left\{ \frac{S_X^2(f)}{S_{X+\eta}(f)}, \theta \right\} df \tag{4.15a}$$

$$= \sigma_X^2 - \int_{F^\star} \left[\frac{S_X^2(f)}{S_{X+\eta}(f)} - \theta \right]^+ df,$$

$$R_\theta = \frac{1}{2}\int_{F^\star} \log^+ \left[\frac{S_X^2(f)}{S_{X+\eta}} / \theta \right] df. \tag{4.15b}$$

Since $F^\star$ is not restricted to be aliasing free, it follows from Corollary 4.5 that for any f_s, R, and H,

$$\underline{D}(f_s,R_\theta) \leq D^\star_{\mathsf{SB}}(f_s,R) \leq D_{\mathsf{SB}(H)}(f_s,R).$$

A water-filling interpretation of $\underline{D}(f_s,R_\theta)$ is given in Figure 4.7. As in the case of the MMSE, the lower bound $\underline{D}(f_s,R_\theta)$ is attained with equality whenever the PSD is unimodal and the pre-sampling filter H is an ideal low-pass filter with pass band $(-f_s/2, f_s/2)$. In general, however, it is impossible to attain $\underline{D}(f_s,R_\theta)$ with equality using a single branch. As we will see in the next section, this lower bound can be attained by increasing the number of sampling branches.

4.3 Multi-Branch Uniform Sampling

We now consider the minimal distortion in the combined sampling and source coding setting of Figure 4.1 where the sampler is a MB uniform sampler, as defined in Section 2.3. We denote the function $D_S(R)$ under this sampler by $D_{\mathsf{MB}(H_1,\ldots,H_L)}(f_s,R)$, where $H_1,\ldots,H_L$ are the L LTI systems that comprise the pre-sampling operation of the sampler. Clearly, the function $D_{\mathsf{SB}(H)}(f_s,R)$ of the previous section is obtained as a special case of $D_{\mathsf{MB}(H_1,\ldots,H_L)}(f_s,R)$ for $L=1$.

In principle, the characterization of $D_{\mathsf{MB}(H_1,\ldots,H_L)}(f_s,R)$ can be obtained using a similar technique to the one that was used in Section 4.2 to derive $D_{\mathsf{SB}(H)}(f_s,R)$. Namely, first use the decomposition in (3.7) to separate the distortion into an MMSE term and a standard information DRF with respect to the asymptotic distribution $X(\cdot)$ of the estimator $X_T(\cdot)$. While the MMSE term is readily available from Proposition 2.2, no analytic expression is known for the DRF of the estimator of $X(\cdot)$ from the output of a MB filter in terms of the PSD of $X(\cdot)$ and $\eta(\cdot)$. We therefore proceed along a different path to evaluate the indirect DRF in the combined sampling and source coding problem under MB sampling: we first consider a discrete-time version of this problem and derive a closed form expression for the indirect DRF in this case. We then increase the time-resolution in the discrete-time problem and show that the indirect DRF there converges to $D_{\mathsf{MB}(H_1,\ldots,H_L)}(f_s,R)$ under mild conditions.

In addition to deriving the expression for $D_{\mathsf{MB}(H_1,\ldots,H_L)}(f_s,R)$, we consider the counterpart of Corollary 4.5, i.e., the optimal set of pre-sampling filters $H_1^\star,\ldots,H_L^\star$ that minimizes $D_{\mathsf{MB}(H_1,\ldots,H_L)}(f_s,R)$.

4.3.1 MB uniform sampling in the discrete-time setting

In this subsection we consider the discrete-time counterpart of the combined sampling and source coding problem of Figure 4.1. While only the expression for the indirect DRF is required in order to evaluate the continuous-time setting, we describe the full source coding setting since its characterization is interesting on its own.

In this discrete-time counterpart, the source $\bar{X}[\cdot]$ is a Gaussian stationary discrete-time process with PSD $S_X\left(e^{2\pi i\phi}\right)$, and the noise $\bar{\eta}[\cdot]$ is another Gaussian stationary process with PSD $S_\eta\left(e^{2\pi i\phi}\right)$ independent of $\bar{X}[\cdot]$. The lth sampling branch consists of a discrete-time LTI pre-sampling operation followed by a factor $M\in\mathbb{N}$ decimator, i.e., the nth sample is given by

$$\bar{Y}_l[n]=\sum_{k=-\infty}^{\infty}h[Mn-k]\left(\bar{X}[k]+\bar{\eta}[k]\right).$$

The encoder receives the sequence $\bar{\mathbf{Y}}[\cdot]=(\bar{Y}_1[\cdot],\ldots,\bar{Y}_L[\cdot])$. After observing $N\in\mathbb{N}$ samples from $\bar{\mathbf{Y}}[\cdot]$, the encoder produces an index out of $2^{MN\bar{R}}$ possible indices (the code-rate is the number of bits per symbol of $\bar{X}[\cdot]$). The decoder produces a reconstruction sequence $\widehat{X}[\cdot]$, and the distortion is evaluated according to a

MSE criterion

$$d_N(\bar{x}[\cdot],\widehat{x}[\cdot]) = \frac{1}{2N+1}\sum_{n=-N}^{N}(\bar{x}[n]-\widehat{x}[n])^2.$$

The minimal expected MSE $\lim_{N\to\infty}\mathbb{E}\left[d_N(\bar{X}[\cdot],\widehat{X}[\cdot]\right]$ taken over all encoders and decoders limited to $\bar{R}$ bits per source symbol of $\bar{X}[\cdot]$ is denoted by $\bar{D}_{MB(H_1,\ldots,H_L)}(M,\bar{R})$. For this setting we have the following result:

Theorem 4.6 (discrete-time multi-branch sampling) *Consider a discrete-time Gaussian stationary process $\bar{X}(\cdot)$ corrupted by a Gaussian stationary noise $\bar{\eta}[\cdot]$. Let $\bar{\mathbf{Y}}[\cdot]$ be the result of sampling $\bar{X}[\cdot]+\bar{\eta}[\cdot]$ using L sampling branches, each with a pre-sampling filter $H\left(e^{2\pi i\phi}\right)$ and decimation factor M. The indirect distortion-rate function of $\bar{X}[\cdot]$ given $\bar{\mathbf{Y}}[\cdot]$ is given by*

$$\bar{D}_{\mathsf{MB}(H_1,\ldots,H_L)}(L,R_\theta) = \mathsf{mmse}_{\mathsf{MB}(H_1,\ldots,H_L)}(M) + \sum_{l=1}^{L}\int_{-\frac{1}{2}}^{\frac{1}{2}}\min\left\{\lambda_l\left(e^{2\pi i\phi}\right),\theta\right\}d\phi \tag{4.16a}$$

$$= \sigma_{\bar{X}}^2 - \sum_{l=1}^{L}\int_{-\frac{1}{2}}^{\frac{1}{2}}\left[\lambda_l\left(e^{2\pi i\phi}\right)-\theta\right]^+ d\phi$$

$$R_\theta = \frac{1}{2}\sum_{l=1}^{L}\int_{-\frac{1}{2}}^{\frac{1}{2}}\log^+\left[\lambda_l\left(e^{2\pi i\phi}\right)/\theta\right]d\phi \tag{4.16b}$$

where

$$\mathsf{mmse}_{\mathsf{MB}(H_1,\ldots,H_L)}(M) = \sigma_{\bar{X}}^2 - \sum_{l=1}^{L}\lambda_l\left(e^{2\pi i\phi}\right)$$

is the MMSE in estimating $\bar{X}[\cdot]$ from $\bar{\mathbf{Y}}[\cdot]$, and $\lambda_1\left(e^{2\pi i\phi}\right) = \ldots\lambda_L\left(e^{2\pi i\phi}\right)$ are the eigenvalues of the $L\times L$ matrix

$$\mathbf{J}_M\left(e^{2\pi i\phi}\right) \triangleq \mathbf{S}_{\bar{\mathbf{Y}}}^{-\frac{1}{2}*}\left(e^{2\pi i\phi}\right)\mathbf{K}_M\left(e^{2\pi i\phi}\right)\mathbf{S}_{\bar{\mathbf{Y}}}^{-\frac{1}{2}}\left(e^{2\pi i\phi}\right). \tag{4.17}$$

Here $\mathbf{S}_{\bar{\mathbf{Y}}}\left(e^{2\pi i\phi}\right)$ is the PSD matrix of the process $\bar{\mathbf{Y}}[\cdot]$ and is given by

$$\left(\mathbf{S}_{\bar{\mathbf{Y}}}\left(e^{2\pi i\phi}\right)\right)_{i,j} \triangleq \frac{1}{ML}\sum_{r=0}^{ML-1}\left\{S_{\bar{X}+\bar{\eta}}H_i^*H_j\right\}\left(e^{2\pi i\frac{\phi-r}{ML}}\right),$$

and $\mathbf{S}_{\bar{\mathbf{Y}}}^{\frac{1}{2}}\left(e^{2\pi i\phi}\right)$ is such that $\mathbf{S}_{\bar{\mathbf{Y}}}\left(e^{2\pi i\phi}\right) = \mathbf{S}_{\bar{\mathbf{Y}}}^{\frac{1}{2}}\left(e^{2\pi i\phi}\right)\mathbf{S}_{\bar{\mathbf{Y}}}^{\frac{1}{2}*}\left(e^{2\pi i\phi}\right)$. The (i,j)th entry of the $L\times L$ matrix $\mathbf{K}_M\left(e^{2\pi i\phi}\right)$ is given by

$$\left(\mathbf{K}_M\right)_{i,j}\left(e^{2\pi i\phi}\right) \triangleq \frac{1}{(ML)^2}\sum_{r=0}^{ML-1}\left\{S_X^2H_i^*H_j\right\}\left(e^{2\pi i\frac{\phi-r}{ML}}\right).$$

Remark 2 *The case where the matrix $\mathbf{S}_{\bar{\mathbf{Y}}}\left(e^{2\pi i\phi}\right)$ is not invertible for some $\phi\in\left(-\frac{1}{2},\frac{1}{2}\right)$ corresponds to a linear dependency between the spectral components of the vector $\bar{\mathbf{Y}}[\cdot]$ in this frequency. In this case we can apply the theorem to the process that is obtained from $\bar{\mathbf{Y}}[\cdot]$ by removing linearly dependent components.*

Proof By stacking the process $X[\cdot]$ into length LM blocks, we see that this indirect DRF coincides with the indirect DRF of the vector-valued process $\mathbf{X}^{PM}[\cdot]$ defined by

$$\mathbf{X}^{ML}[n] = \left(\bar{X}[LMn], \bar{X}[LMn+1], \ldots \bar{X}[LMn+LM-1]\right).$$

For a given $M \in \mathbb{N}$, $\mathbf{X}^{ML}[\cdot]$ is a stationary Gaussian process with PSD matrix

$$\left(\mathbf{S_X}\right)_{r,s}\left(e^{2\pi i\phi}\right) = \frac{1}{ML}\sum_{m=0}^{ML-1} e^{2\pi i(r-s)\frac{\phi-m}{ML}} S_X\left(e^{2\pi i\frac{\phi-m}{ML}}\right).$$

For $l = 1,\ldots,L$, we defined the process $\bar{Z}_l[\cdot]$ as

$$\bar{Z}_l[n] = \sum_{k\in\mathbb{Z}} \left(\bar{X}[k]+\bar{\eta}[k]\right) h_l[n-k],$$

so

$$S_{\bar{Z}_l}\left(e^{2\pi i\phi}\right) = S_{\bar{X}+\bar{\eta}}\left(e^{2\pi i\phi}\right) H^*\left(e^{2\pi i\phi}\right).$$

The processes $\mathbf{Y}[\cdot]$ and $\mathbf{X}^{ML}[\cdot]$ are jointly Gaussian and stationary with a $ML\times L$ cross PSD whose $(m+1,p)^{\text{th}}$ entry is given by

$$\begin{aligned}\left(\mathbf{S}_{\mathbf{X}^{ML}\bar{\mathbf{Y}}}\right)_{m,p}\left(e^{2\pi i\phi}\right) = S_{\mathbf{X}^{ML}_m\bar{\mathbf{Y}}_p}\left(e^{2\pi i\phi}\right) &= \sum_{k\in\mathbb{Z}} \mathbb{E}\left[\bar{X}[MLk+m]\bar{Z}_p[0]\right] e^{-2\pi i\phi k} \\ &= \frac{1}{ML}\sum_{r=0}^{ML-1} e^{2\pi i m\frac{\phi-r}{ML}} S_{\bar{X}\bar{Z}_p}\left(e^{2\pi i\frac{\phi-r}{ML}}\right),\end{aligned}$$

where we denote by $X^{ML}_m[\cdot]$ the mth coordinate of $\mathbf{X}^{ML}[\cdot]$. The PSD of the MMSE estimator of $\mathbf{X}^{ML}[\cdot]$ from $\bar{\mathbf{Y}}[\cdot]$ is given by

$$\mathbf{S}_{\mathbf{X}^{ML}|\bar{\mathbf{Y}}}\left(e^{2\pi i\phi}\right) = \left\{\mathbf{S}_{\mathbf{X}^{ML}\bar{\mathbf{Y}}}\mathbf{S}_{\bar{\mathbf{Y}}}^{-1}\mathbf{S}^*_{\mathbf{X}^{ML}\bar{\mathbf{Y}}}\right\}\left(e^{2\pi i\phi}\right), \tag{4.18}$$

Since only the non-zero eigenvalues of $\mathbf{S}_{\mathbf{X}^{ML}|\bar{\mathbf{Y}}}\left(e^{2\pi i\phi}\right)$ contribute to the distortion in (4.16), we are interested in the non-zero eigenvalues of (4.18). These eigenvalues are identical to the non-zero eigenvalues of

$$\left\{\mathbf{S}_{\bar{\mathbf{Y}}}^{-\frac{1}{2}*}\mathbf{S}^*_{\mathbf{X}^{ML}\bar{\mathbf{Y}}}\mathbf{S}_{\mathbf{X}^{ML}\mathbf{Y}}\mathbf{S}_{\bar{\mathbf{Y}}}^{-\frac{1}{2}}\right\}\left(e^{2\pi i\phi}\right), \tag{4.19}$$

where $\left\{\mathbf{S}_{\bar{\mathbf{Y}}}^{-\frac{1}{2}*}\mathbf{S}_{\bar{\mathbf{Y}}}^{-\frac{1}{2}}\right\}\left(e^{2\pi i\phi}\right)=\mathbf{S}_{\bar{\mathbf{Y}}}^{-1}\left(e^{2\pi i\phi}\right)$. The (p,q)th entry of the $L\times L$ matrix $\left\{\mathbf{S}_{\mathbf{X}^{PM}\bar{\mathbf{Y}}}^{*}\mathbf{S}_{\mathbf{X}^{ML}\bar{\mathbf{Y}}}\right\}\left(e^{2\pi i\phi}\right)$ is given by

$$
\begin{aligned}
&\frac{1}{(ML)^2}\sum_{l=0}^{ML-1}\sum_{r=0}^{ML-1}e^{-2\pi i l\frac{\phi-r}{ML}}S_{\bar{X}\bar{Z}_p}\left(e^{2\pi i\frac{\phi-r}{ML}}\right)\sum_{k=0}^{ML-1}e^{2\pi i l\frac{\phi-k}{ML}}S_{\bar{X}\bar{Z}_q}\left(e^{2\pi i\frac{\phi-k}{ML}}\right)\\
=&\frac{1}{(ML)^2}\sum_{r=0}^{ML-1}S_{\bar{X}\bar{Z}_p}\left(e^{2\pi i\frac{\phi-r}{ML}}\right)S_{\bar{X}\bar{Z}_q}\left(e^{2\pi i\frac{\phi-r}{ML}}\right)\\
=&\frac{1}{(ML)^2}\sum_{r=0}^{ML-1}\left\{S_{\bar{X}}^2H_p^*H_q\right\}\left(e^{2\pi i\frac{\phi-r}{ML}}\right),
\end{aligned}
$$

which equals the entries of the matrix $\mathbf{K}_M\left(e^{2\pi i\phi}\right)$ defined in the theorem. Applying Theorem 3.2 with the eigenvalues of (4.19) completes the proof. □

In the next section we use Theorem 4.6 to evaluate the indirect DRF under MB uniform sampling in continuous time.

4.3.2 Multi-branch uniform sampling in continuous-time

We now return to the continuous-time setting of the combined sampling and source coding problem with MB uniform sampling. The following theorem provides a characterization of the function $D_{\mathsf{MB}(H_1,\ldots,H_L)}(f_s,R)$, describing the minimal distortion in this setting.

Theorem 4.7 *Consider a Gaussian stationary process $X(\cdot)$ corrupted by a Gaussian stationary noise $\eta(\cdot)$. Let the discrete-time vector valued process $\mathbf{Y}[\cdot]$ be the result of sampling $X(\cdot)+\eta(\cdot)$ using L sampling branches, each with a pre-sampling filter $H(f)$ at sampling rate f_s/L. The indirect distortion-rate function of $X(\cdot)$ given $\mathbf{Y}[\cdot]$ is given by*

$$
\begin{aligned}
D_{\mathsf{MB}(H_1,\ldots,H_L)}(f_s,R_\theta)&=\mathsf{mmse}_{\mathsf{MB}(H_1,\ldots,H_L)}(f_s)+\sum_{l=1}^{L}\int_{-\frac{f_s}{2}}^{\frac{f_s}{2}}\min\left\{\lambda_l\left(\widetilde{\mathbf{S}}_{X|\mathbf{Y}}(f)\right),\theta\right\}df,\\
&=\sigma_X^2-\sum_{l=1}^{L}\int_{-\frac{f_s}{2}}^{\frac{f_s}{2}}\left[\lambda_l\left(\widetilde{\mathbf{S}}_{X|\mathbf{Y}}(f)\right)-\theta\right]^+df, \qquad (4.20a)
\end{aligned}
$$

$$
R_\theta=\frac{1}{2}\sum_{l=1}^{L}\int_{-\frac{f_s}{2}}^{\frac{f_s}{2}}\log^+\left[\lambda_l\left(\widetilde{\mathbf{S}}_{X|\mathbf{Y}}(f)\right)/\theta\right]df \qquad (4.20b)
$$

where $\lambda_1\left(\widetilde{\mathbf{S}}_{X|\mathbf{Y}}(f)\right)=\ldots=\lambda_L\left(\widetilde{\mathbf{S}}_{X|\mathbf{Y}}(f)\right)$ are the eigenvalues of the $L\times L$ matrix

$$
\widetilde{\mathbf{S}}_{X|\mathbf{Y}}(f)=\tilde{\mathbf{S}}_{\mathbf{Y}}^{-\frac{1}{2}*}(f)\mathbf{K}(f)\tilde{\mathbf{S}}_{\mathbf{Y}}^{-\frac{1}{2}}(f),
$$

and the (i,j)th entry of the matrices $\tilde{\mathbf{S}}_{\mathbf{Y}}(f), \mathbf{K}(f) \in \mathbb{C}^{L\times L}$ are given by

$$\left(\tilde{\mathbf{S}}_{\mathbf{Y}}\right)_{i,j}(f) = \sum_{k\in\mathbb{Z}} \left\{S_{X+\eta} H_i H_j^*\right\}(f - f_s k),$$

and

$$\mathbf{K}_{i,j}(f) = \sum_{k\in\mathbb{Z}} \left\{S_X^2 H_i H_j^*\right\}(f - f_s k).$$

Proof As explained in Section 2.4, the finite horizon MMSE estimator of $X(\cdot)$ from $\mathbf{Y}_T$ under time-invariant uniform sampling converges to an asymptotic law with periodic covariance function. In particular, the asymptotic process is asymptotic mean stationary and therefore Theorem 3.1 holds. We do not fully utilize the decomposition (3.7), but only use it to connect the indirect DRF for the MB sampler to an information expression as appears in (3.7). Indeed, we now evaluate both expressions in (3.7) simultaneously by approximating the process $X(\cdot)$ by a discrete-time process and take the limit in this approximation. Specifically, for $M \in \mathbb{N}$, define $X^M[\cdot]$ to be the processes obtained by uniformly sampling $X(\cdot)$ at rate $f_s M$, i.e. $X^M[n] = X(n/(f_s M))$. Similarly, for $l = 1,\ldots,L$, let $Z_l^M[\cdot]$ be the process obtained by sampling the continuous-time process at the output of the filter H_l at rate $f_s M$. We have

$$S_{X^M}\left(e^{2\pi i\phi}\right) = M f_s \sum_{k\in\mathbb{Z}} S_X\left(M f_s(\phi - k)\right)$$

and

$$S_{Z_l^M}\left(e^{2\pi i\phi}\right) = M f_s \sum_{k\in\mathbb{Z}} \left\{S_X |H_l|^2\right\}\left(M f_s(\phi - k)\right).$$

For $l, r = 1,\ldots,L$, $Z_l^M[\cdot]$ and $Z_r^M[\cdot]$ are jointly stationary with cross PSD

$$S_{Z_l^M Z_r^M}\left(e^{2\pi i\phi}\right) = M f_s \sum_{k\in\mathbb{Z}} S_{Z_l Z_r}\left(M f_s(\phi - k)\right).$$

In addition, $X^M[\cdot]$ and $Z_l^M[\cdot]$ are jointly stationary processes with cross PSD

$$S_{X^M Z_l^M}\left(e^{2\pi i\phi}\right) = M f_s \sum_{k\in\mathbb{Z}} S_{X Z_p}\left(M f_s(\phi - k)\right).$$

Since the sampling rate of each branch is f_s/L, $Y_l[\cdot]$ is a factor ML decimated version of $Z_l^M[\cdot]$. Therefore, the indirect DRF of $X^M[\cdot]$ given $\mathbf{Y}[\cdot]$ is given by Theorem 4.6, where the (p,r)th entry of the PSD matrix

$\mathbf{S_Y}\left(e^{2\pi i\phi}\right)$is given by

$$\begin{aligned}
\left(\mathbf{S_Y}\left(e^{2\pi i\phi}\right)\right)_{p,r} &= \frac{1}{ML}\sum_{m=0}^{ML-1} S_{Z_p^M Z_r^M}\left(e^{2\pi i\frac{\phi-m}{ML}}\right) \\
&= f_s \sum_{m=0}^{ML-1}\sum_{k\in\mathbb{Z}} S_{Z_pZ_r}\left(f_s\left(\phi-m-MLk\right)\right) \\
&= f_s \sum_{n\in\mathbb{Z}} S_{Z_pZ_r}\left(f_s\left(\phi-n\right)\right) \\
&= f_s \sum_{n\in\mathbb{Z}} \left\{S_{X+\eta}H_pH_r^*\right\}\left(f_s\phi - f_sn\right) \\
&= \left(\tilde{\mathbf{S}}_{\mathbf{Y}}\right)_{p,r}\left(\phi f_s\right) f_s,.
\end{aligned} \tag{4.21}$$

Similarly, the matrix $\mathbf{K}_M$ is given by

$$\begin{aligned}
\left(\mathbf{K}_M\right)_{p,r} &= \frac{1}{(ML)^2}\sum_{m=0}^{ML-1}\left\{S_{X^M}^2 H_pH_r\right\}\left(e^{2\pi i\frac{\phi-m}{ML}}\right) \\
&= \frac{1}{(ML)^2}\sum_{m=0}^{ML-1}\left\{S_{X^M Z_p^M}S^*_{X^M Z_r^M}\right\}\left(e^{2\pi i\frac{\phi-m}{ML}}\right) \\
&= f_s^2\sum_{m=0}^{ML-1}\left[\sum_{k\in\mathbb{Z}} S_{XZ_p}\left(f_s\left(\phi-m-kML\right)\right)\right. \qquad (4.22)\\
&\qquad \left.\times\sum_{l\in\mathbb{Z}} S^*_{XZ_r}\left(f_s\left(\phi-m-lML\right)\right)\right]. \qquad (4.23)
\end{aligned}$$

Since a version of the Gaussian stationary process $X(\cdot)$ with almost surely Riemann integrable paths exists [69], the expected MSE between $X(\cdot)$ and any typical reconstruction of it from $X^M[\cdot]$ (e.g. zero-order hold interpolation) converges to zero as $M\to\infty$. In addition, given an estimate $\widehat{X}^M[\cdot]$ of $X^M[\cdot]$ attaining the indirect DRF of $X^M[\cdot]$ given $\mathbf{Y}[\cdot]$, we can define $\widehat{X}(\cdot)$ to be an interpolated version of $X^M[\cdot]$ such that the expected MSE between $X(\cdot)$ and $\widehat{X}(\cdot)$ goes to zero. Therefore, the indirect DRF of $X^M[\cdot]$ given $\mathbf{Y}[\cdot]$ converges to the indirect DRF of $X(\cdot)$ given $\mathbf{Y}[\cdot]$.

We turn to evaluate (4.16) with (4.21) and (4.23). First note that we have $\sigma^2_{X^M}=\sigma^2_X$, as can easily be verified from properties of sampled signals. In addition, by a change of the integration variable from f to $\phi = f/f_s$, we can write (4.20):

$$D_{\mathsf{MB}(H_1,\ldots,H_L)}(f_s,R_\theta) = \sigma_X^2 - \sum_{l=1}^{L}\int_{-\frac{1}{2}}^{\frac{1}{2}}\left[\lambda_l\left(\bar{\mathbf{J}}\left(e^{2\pi i\phi}\right)\right)-\theta\right]^+ d\phi, \tag{4.24a}$$

$$R_\theta = \frac{f_s}{2}\sum_{l=1}^{L}\int_{-\frac{1}{2}}^{\frac{1}{2}}\log^+\left[\lambda_l\left(\bar{\mathbf{J}}\left(e^{2\pi i\phi}\right)\right)/\theta\right]d\phi \tag{4.24b}$$

where in $\phi \in \left(-\frac{1}{2}, \frac{1}{2}\right)$, the matrix $\bar{\mathbf{J}}\left(e^{2\pi i\phi}\right)$ is given by

$$\bar{\mathbf{J}}\left(e^{2\pi i\phi}\right) = \mathbf{S}_{\mathbf{Y}}^{-\frac{1}{2}*}\left(e^{2\pi i\phi}\right)\bar{\mathbf{K}}\left(e^{2\pi i\phi}\right)\mathbf{S}_{\mathbf{Y}}^{-\frac{1}{2}}\left(e^{2\pi i\phi}\right),$$

and $\bar{\mathbf{K}}\left(e^{2\pi i\phi}\right) = f_s^2\mathbf{K}(f_s\phi)$. It follows that in order to complete the proof, it is enough to show that the eigenvalues of $\mathbf{J}_M\left(e^{2\pi i\phi}\right)$, considered as $L_1\left(-\frac{1}{2}, \frac{1}{2}\right)$ functions in ϕ, converge to the eigenvalues of $\bar{\mathbf{J}}\left(e^{2\pi i\phi}\right)$. Since

$$\|\mathbf{S}_{\mathbf{Y}}{}^{-\frac{1}{2}*}\mathbf{K}_M\mathbf{S}_{\mathbf{Y}}{}^{-\frac{1}{2}} - \mathbf{S}_{\mathbf{Y}}{}^{-\frac{1}{2}*}\bar{\mathbf{K}}\mathbf{S}_{\mathbf{Y}}{}^{-\frac{1}{2}}\|_F \leq \|\mathbf{S}_{\mathbf{Y}}\|_F^{-1}\|f_s\mathbf{K}_M - \bar{\mathbf{K}}\|_F,$$

it is enough to prove convergence in $L_1\left(-\frac{1}{2}, \frac{1}{2}\right)$ for each entry, i.e. that

$$\lim_{M\to\infty}\int_{-\frac{1}{2}}^{\frac{1}{2}} \frac{\left|\left(\mathbf{K}_M\right)_{p,r}\left(e^{2\pi i\phi}\right) - \left(\bar{\mathbf{K}}\right)_{p,r}\left(e^{2\pi i\phi}\right)\right|}{\|\mathbf{S}_{\mathbf{Y}}\left(e^{2\pi i\phi}\right)\|_F} d\phi = 0 \tag{4.25}$$

for all $p, r = 1, \ldots, L$. We now use the following lemma:

Lemma 4.8 *Let $f_1(\varphi)$ and $f_2(\varphi)$ be two complex valued bounded functions such that $\int_{-\infty}^{\infty}|f_i(\varphi)|^2 d\varphi < \infty$, $i = 1, 2$. Then for any fixed $f_s > 0$,*

$$\int_{-\frac{1}{2}}^{\frac{1}{2}}\sum_{m=0}^{M-1}\sum_{k\in\mathbb{Z}} f_1(\phi + m + kM)\sum_{l\in\mathbb{Z}} f_2(\phi + m + lM)\, d\phi \tag{4.26}$$

converges to

$$\int_{-\frac{1}{2}}^{\frac{1}{2}}\sum_{n\in\mathbb{Z}} f_1(\phi - n) f_2(\phi - n)\, d\phi, \tag{4.27}$$

as M goes to infinity.

Proof [of Lemma 4.8] Equation (4.26) can be written as

$$\int_{-\frac{1}{2}}^{\frac{1}{2}}\sum_{m=0}^{M-1}\sum_{k\in\mathbb{Z}} f_1(\phi + m + kM) f_2(\phi + m + kM)\, d\phi \tag{4.28}$$

$$+\int_{-\frac{1}{2}}^{\frac{1}{2}}\sum_{m=0}^{M-1}\sum_{k\neq l} f_1(\phi + m + kM) f_2(\phi + m + lM)\, d\phi. \tag{4.29}$$

Since the term (4.28) is identical to (4.27), we must show that (4.29) vanishes as $M \to \infty$. Take M large enough such that

$$\int_{\mathbb{R}\setminus\left[-\frac{M+1}{2}, \frac{M+1}{2}\right]} |f_i(\phi)|^2 d\phi < \varepsilon^2, \quad i = 1, 2.$$

Without loss of generality assume that this M is even. By a change of variables (4.29) can be written as

$$\sum_{k\neq l}\int_{-\frac{M+1}{2}}^{\frac{M+1}{2}} f_1\left(\varphi+\frac{M}{2}+kM\right)f_2\left(\varphi+\frac{M}{2}+lM\right)d\varphi. \tag{4.30}$$

We split the indices in the last sum into three disjoint sets:

1. $\mathscr{I}=\{k,l\in\mathbb{Z}\setminus\{0,-1\},k\neq l\}$,

$$\begin{aligned}
&\left|\sum_{\mathscr{I}}\int_{-\frac{M+1}{2}}^{\frac{M+1}{2}} f_1\left(\varphi+\frac{M}{2}+kM\right)f_2\left(\varphi+\frac{M}{2}+lM\right)d\varphi\right|\\
&\overset{a}{\leq}\sum_{\mathscr{I}}\int_{-\frac{M+1}{2}}^{\frac{M+1}{2}}\left|f_1\left(\varphi+\frac{M}{2}+kM\right)\right|^2d\varphi+\sum_{\mathscr{I}}\int_{-\frac{M+1}{2}}^{\frac{M+1}{2}}\left|f_2\left(\varphi+\frac{M}{2}+lM\right)\right|^2d\varphi\\
&\leq\int_{\mathbb{R}\setminus[-\frac{M+1}{2},\frac{M+1}{2}]}|f_1(\varphi)|^2d\varphi+\int_{\mathbb{R}\setminus[-\frac{M+1}{2},\frac{M+1}{2}]}|f_2(\varphi)|^2d\varphi\leq 2\varepsilon^2,
\end{aligned}\tag{4.31}$$

where (a) is due to the triangle inequality and since for any two complex numbers a,b, $|ab|\leq\frac{|a|^2+|b|^2}{2}\leq|a|^2+|b|^2$.

2. $k=0,l=-1$,

$$\begin{aligned}
&\int_{-\frac{M+1}{2}}^{\frac{M+1}{2}} f_1\left(\varphi+\frac{M}{2}\right)f_2\left(\varphi-\frac{M}{2}\right)d\varphi\\
&=\int_{-\frac{M+1}{2}}^{0} f_1\left(\varphi+\frac{M}{2}\right)f_2\left(\varphi-\frac{M}{2}\right)d\varphi+\int_{0}^{\frac{M+1}{2}} f_1\left(\varphi+\frac{M}{2}\right)f_2\left(\varphi-\frac{M}{2}\right)d\varphi\\
&\overset{a}{\leq}\sqrt{\int_{-\frac{M+1}{2}}^{0} f_1^2\left(\varphi+\frac{M}{2}\right)d\varphi}\sqrt{\int_{-\frac{M+1}{2}}^{0} f_2^2\left(\varphi-\frac{M}{2}\right)d\varphi}\\
&\quad+\sqrt{\int_{0}^{\frac{M+1}{2}} f_1^2\left(\varphi+\frac{M}{2}\right)d\varphi}\sqrt{\int_{0}^{\frac{M+1}{2}} f_2^2\left(\varphi-\frac{M}{2}\right)d\varphi}\\
&\leq\sqrt{\int_{-\frac{M+1}{2}}^{0} f_1^2\left(\varphi+\frac{M}{2}\right)d\varphi}\sqrt{\int_{\mathbb{R}\setminus[-\frac{M+1}{2},\frac{M+1}{2}]} f_2^2(\phi)\,d\phi}\\
&\quad+\sqrt{\int_{\mathbb{R}\setminus[-\frac{M+1}{2},\frac{M+1}{2}]} f_1^2(\phi)\,d\phi}\sqrt{\int_{0}^{\frac{M+1}{2}} f_2^2\left(\varphi-\frac{M}{2}\right)d\varphi}
\end{aligned}\tag{4.32}$$

$$\leq\varepsilon\|f_1\|_2+\varepsilon\|f_2\|_2, \tag{4.33}$$

where (a) follows from the Cauchy-Schwartz inequality.

3. $k=-1, l=0$, using the same arguments as in the previous case,

$$\int_{-\frac{M+1}{2}}^{\frac{M+1}{2}} f_1\left(\varphi+\frac{M}{2}\right) f_2\left(\varphi-\frac{M}{2}\right) d\varphi \leq \varepsilon\left(\|f_1\|_2+\|f_2\|_2\right). \tag{4.34}$$

From (4.31), (4.33) and (4.34), the sum (4.30) can be bounded by

$$2\varepsilon\left(\|f_1\|_2+\|f_2\|_2\right)+2\varepsilon^2,$$

which can be made as close to zero as required. □

In order to complete the proof of Theorem 4.7, we apply Lemma 4.8 to (4.25) with

$$f_1(\phi)=\frac{S_{XZ_p}(f_s\phi)}{\sqrt{\|\mathbf{S_Y}(e^{2\pi i\phi})\|_F}},$$

$$f_2(\phi)=\frac{S^*_{XZ_r}(f_s\phi)}{\sqrt{\|\mathbf{S_Y}(e^{2\pi i\phi})\|_F}}.$$

and note that

$$\left(\bar{\mathbf{K}}\right)_{p,r}\left(e^{2\pi i\phi}\right)=f_s^2\left(\mathbf{K}\right)_{p,r}(f_s\phi)=f_s^2\sum_{k\in\mathbb{Z}}\left\{S_X^2H_iH_j^*\right\}(f_s(\phi-k)).$$

It follows from Lemma 4.8 that (4.25) converges to (4.20). □

4.3.3 Optimal Filter-Bank

A similar analysis to that we used for SB sampling will show that for a fixed R, the distortion is a non-increasing function of the eigenvalues of $\widetilde{\mathbf{S}}_{X|Y}(f)$. This implies that the optimal pre-sampling filters $H_1^\star,\ldots,H_L^\star$ that minimize the function $D_{\mathsf{MB}(H_1,\ldots,H_L)}(f_s,R)$, for a given R and f_s, are the exact same filters that minimize the MMSE of $X(\cdot)$ given $\mathbf{Y}[\cdot]$ in Theorem 2.7. Therefore, the following result applies:

Theorem 4.9 *The optimal pre-sampling filters $H_1^\star,\ldots,H_L^\star$ that minimize $D_{\mathsf{MB}(H_1,\ldots,H_L)}(f_s,R)$ are given by*

$$H_l^\star(f)=\begin{cases}1 & f\in F_l^\star,\\ 0 & f\notin F_l^\star,\end{cases}\quad l=1,\ldots,L, \tag{4.35}$$

where $F_l^\star,\ldots,F_L^\star$ satisfy conditions (i) and (ii) in Theorem 2.7. Moreover, with $H_1^\star,\ldots,H_L^\star$ the indirect DRF *of $X(\cdot)$ given $\mathbf{Y}[\cdot]$ is given by*

$$D^\star_{\mathsf{MB}(L)}(f_s,R_\theta) \triangleq D_{H_1^\star,\dots,H_L^\star}(f_s,R_\theta) = \mathsf{mmse}^\star_{MB}(f_s) + \sum_{l=1}^{L}\int_{F_l^\star} \min\left\{\frac{S_X^2(f)}{S_{X+\eta}(f)},\theta\right\} df,$$
$$= \sigma_X^2 - \sum_{p=1}^{L}\int_{F_l^\star}\left[\frac{S_X^2(f)}{S_{X+\eta}(f)}-\theta\right]^+ df. \tag{4.36a}$$

$$R_\theta = \frac{1}{2}\sum_{l=1}^{L}\int_{F_l^\star}\log^+\left[\frac{S_X^2(f)}{S_{X+\eta}(f)}/\theta\right]df \tag{4.36b}$$

Proof The filters $H_1^\star,\dots,H_L^\star$ given by Theorem 2.7 maximize the eigenvalues of the matrix $\widetilde{\mathbf{S}}_{X|\mathbf{Y}}(f)$ of (2.23) for every $f\in(-f_s/2,f_s/2)$. The same filters also minimize $D_{\mathsf{MB}(H_1,\dots,H_L)}(f_s,R)$, since it is monotone non-increasing in these eigenvalues. Finally, for this choice of filter, (4.20) reduces to (4.36). □

It follows from Corollary 4.4 that in the parametric reverse water-filling representation of the form (4.36), an increment in

$$\int_{\bigcup_{l=1}^L \mathscr{F}_l^\star}\frac{S_X^2(f)}{S_{X+\eta}(f)}df$$

reduces the overall distortion regardless of the other terms. However, since $\mu\left(\bigcup_{i=1}^L F_l^\star\right)\le f_s$, we conclude that the lower bound of (4.15) still holds under MB sampling:

Corollary 4.10 *For any $H_1,\dots,H_L$,*

$$\underline{D}(f_s,R)\le D^\star_{\mathsf{MB}(L)}(f_s,R)\le D_{\mathsf{MB}(H_1,\dots,H_L)}(f_s,R),$$

where $\underline{D}(f_s,R)$ is defined in (4.15).

4.3.4 Asymptotically many sampling branches

We now consider the behavior of $D^\star_{\mathsf{MB}(L)}(f_s,R)$ as the number of sampling branches L increases, which is the counterpart of Section 2.5.3 for the indirect DRF of $X(\cdot)$ given the output $\mathbf{Y}[\cdot]$ of a MB uniform sampler.

In Theorem 2.8 we have seen that as L increases, the pre-sampling filters $H_1,\dots,H_L$ can be chosen to approximate the set $F^\star$. As a result, we obtain the following theorem:

Theorem 4.11 *For any $f_s>0$ and $\varepsilon>0$, there exists $L\in\mathbb{N}$ and a set of LTI filters $H_1,\dots,H_L$ such that*

$$D_{\mathsf{MB}(H_1,\dots,H_L)}(f_s,R)-\varepsilon<\underline{D}(f_s,R). \tag{4.37a}$$

Proof In Theorem 2.8, we found that for f_s and $\varepsilon > 0$ there exists L large enough and a set of pre-sampling filters $H_1,\ldots,H_L$, each with an aliasing-free support $F_1,\ldots,F_L$, respectively, such that

$$\sum_{l=1}^{L}\int_{F_l}\frac{S_X^2(f)}{S_X(f)+S_\eta(f)}df+\delta > \int_{F^\star}\frac{S_X^2(f)}{S_{X+\eta}(f)}df.$$

Now, use Proposition 3.4 with $A = F^\star \cup F_1 \cup \ldots \cup F_L$,

$$f(x) = \mathbf{1}_{F^\star}\left(\frac{S_X^2(x)}{S_{X+\eta}(x)}\right),$$

and

$$g(x) = \mathbf{1}_{\cup_{l=1}^{L}F_l}\left(\frac{S_X^2(x)}{S_{X+\eta}(x)}+\delta\right).$$

Note that $\underline{D}(f_s,R)$ is a water-filling expression of the form (4.15) over $f(x)$ and A. Denote by D_δ the function defined by a water-filling expression over $g(x)$. Since $g(x) \geq f(x)$, it follows from Proposition 3.4 that

$$D_\delta \leq \underline{D}(f_s,R).$$

Since D_δ is continuous in δ and since $\lim_{\delta\to 0} D_\delta = D_{\mathsf{MB}(H_1,\ldots,H_L)}(f_s,R)$, for $\varepsilon > 0$ there exists L and δ such that $D_{\mathsf{MB}(H_1,\ldots,H_L)}(f_s,R)+\varepsilon > D_\delta \geq \underline{D}(f_s,R)$. □

4.3.5 Discussion

An immediate conclusion from Theorem 4.11 is that $D^\star_{\mathsf{MB}(L)}(f_s,R)$ converges to $\underline{D}(f_s,R)$ as the number of sampling branches L goes to infinity, as illustrated in Figure 4.8. In other words, with enough sampling branches and a careful selection of the pre-sampling filters according to (4.9), MB uniform sampling attains the distortion lower bound (4.15). It follows from its definition that the function $\underline{D}(f_s,R)$ is monotone in f_s. This last fact is in contrast to $D_{\mathsf{MB}(H_1,\ldots,H_L)}(f_s,R)$ and $D^\star_{\mathsf{MB}(L)}(f_s,R)$, which, as Figure 4.8 shows, are not guaranteed to be monotone in f_s. Figure 4.8 also suggests that multi-branch sampling can significantly reduce distortion for a given sampling frequency f_s and source coding rate R over single-branch sampling. Figure 4.8 also raises the possibility of reducing the sampling frequency without significantly affecting performance, as the function $\underline{D}^\star_{\mathsf{MB}(L)}(f_s,R)$ for $L > 1$ approximately achieves the asymptotic value of $D^\star_{\mathsf{MB}(L)}(f_{\mathsf{Nyq}},R)$ at $f_s \approx f_{\mathsf{Nyq}}/3$. In fact, as will be discussed in the next chapter, the minimal distortion subject to a bitrate constraint, described by the standard DRF of (3.3), is attained by sampling below the Nyquist rate of the source.

In the next section we show that the lower bound $\underline{D}(f_s,R)$ of (4.15) holds under any bounded linear sampler, and hence describes the fundamental distortion limit in ADX.

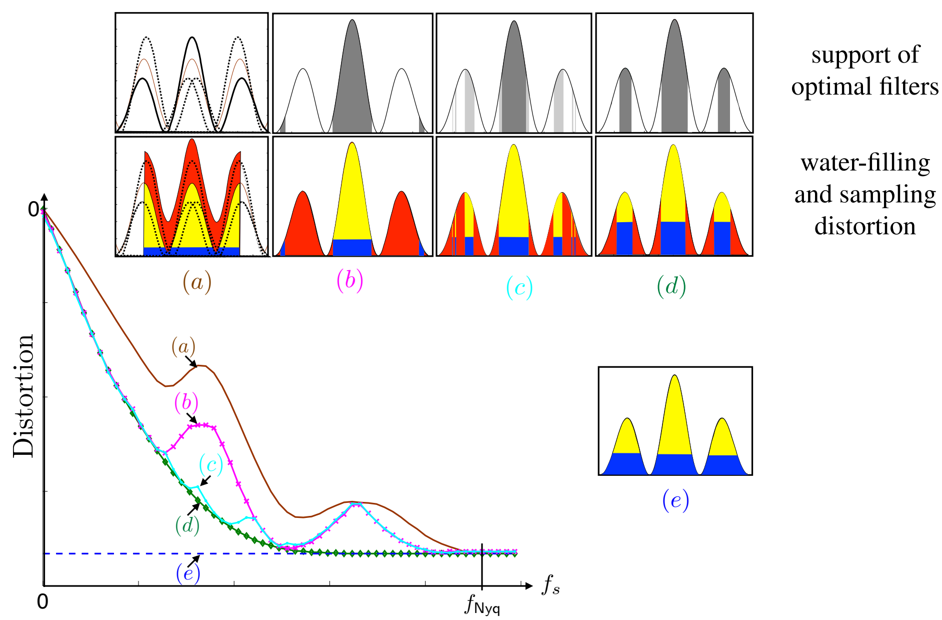

Figure 4.8: Minimal distortion versus the sampling rate f_s and a fixed value of R. (a) SB uniform sampler and all-pass pre-sampling filter (same situation as in Figure 4.3). (b) SB uniform and an optimal pre-sampling filter. The passband of the pre-sampling filter is given by shaded area over the PSD. (c) Two sampling branches with optimal filter-bank. Support of the pre-sampling filters corresponds to the two shaded areas over the PSD. (d) Five sampling branches achieves $D(f_s,R)$. (e) The standard DRF of (3.3). The measure of each of the shaded areas in (b)-(d) equals f_s.

4.4 General Bounded Linear Sampling

So far we have shown that the function $\underline{D}(f_s, R)$ of (4.15) provides an achievable lower bound for the distortion in ADX under MB sampling. In this section we extend this claim and show that the distortion in ADX under any bounded linear sampler is not smaller than $\underline{D}(f_s, R)$, provided f_s is replaced by the lower Beurling density of the sampling set.

Theorem 4.12 *Let $S(H,\Lambda)$ be a bounded linear sampler. Then the minimal distortion in the ADX setting with a Gaussian stationary input $X(\cdot)$ and a Gaussian stationary noise $\eta(\cdot)$ satisfies*

$$D_{S(H,\Lambda)}(R) \geq \underline{D}\left(d^-(\Lambda), R\right),$$

where $d^-(\Lambda)$ is the lower Beurling density of Λ.

Proof We first prove Theorem 4.12 under the following simplifying assumptions:

(i) The sampling set Λ is periodic, meaning there exists T_0 such that $\Lambda = \Lambda + T_0$.

(ii) Denote $g(t,\tau) = K_H(t, t-\tau)$. Then the kernel g is T_0-periodic in its first argument, namely, $g(t,\tau) = g(t+T_0,\tau)$ for any t and τ.

For Λ T_0-periodic, the number of points in Λ in any interval of length T_0 is denoted by P. It is easy to see that the Beurling density of Λ exists and equals $d(\Lambda) = P/T_0$. Denote by $t_0, \ldots, t_{P-1}$ the P members of Λ inside the interval $[0, T_0)$, and continue to enumerate the members of Λ in the positive direction in a similar manner. Similarly, enumerate the elements of Λ in the negative direction starting from t_{-1}, t_{-2}. By the periodicity of Λ, $t_{p+Pk} = t_p + T_0 k$ for all $p = 0, \ldots, P-1$ and $k \in \mathbb{Z}$. For $n = p + kP$ and $T > 0$ we have,

$$\begin{aligned} Y[n] &= \int_{-\infty}^{\infty} K_H(t_{p+Pk}, \tau)\,(X(\tau)+\eta(\tau))\,d\tau = \int_{-\infty}^{\infty} g(t_p + T_0 k, \tau - t_p + T_0 k)\,(X(\tau)+\eta(\tau))\,d\tau \\ &= \int_{-\infty}^{\infty} g(t_p, t_p + T_0 k - \tau)\,(X(\tau)+\eta(\tau))\,d\tau. \end{aligned}$$

We now down-sample the discrete-time index set by a factor P and replace the vector $Y[\cdot]$ by a vector valued process that contains the P indices $kP, kP+1, \ldots, kP+P-1$, namely

$$\mathbf{Y}[k] = \left(Y[Pk], Y[Pk+1], \ldots, Y[Pk+P-1]\right)).$$

For $p = 0, \ldots, P-1$ denote

$$g_p(\tau) \triangleq g(t_p, \tau - t_p), \quad t \in \mathbb{R}.$$

It follows that for any $p = 0,\ldots,P-1$ and $k \in \mathbb{Z}$ we have

$$Y_p[k] = Y[Pk+p] = \int_{-\infty}^{\infty} g(t_p, t_p + T_0 k - \tau)\,(X(\tau)+\eta(\tau))\,d\tau = \int_{-\infty}^{\infty} g_p(T_0 k - \tau)\,(X(\tau)+\eta(\tau))\,d\tau. \quad (4.38)$$

Since $h_p(\tau)$ defines an LTI system, it follows that sampling with the set Λ and the pre-processing system H is equivalent to P uniform sampling branches each of sampling frequency $1/T$ and a pre-sampling filter

$$H_p(f) \triangleq e^{2\pi i f t_p} G_p(f),$$

where $G_p(f)$ is the Fourier transform of $g_p(t)$ with respect to t, for $p = 0,\ldots,P-1$. Denote by $D_{\Lambda,g}(R)$ the indirect DRF of $X(\cdot)$ given the process $\mathbf{Y}_{\Lambda,g}[\cdot]$. If follows from Corollary 4.10 that

$$D_{S(H,\Lambda)}(R) \geq \underline{D}(T_0^{-1}, R) = \underline{D}(d(\Lambda), R).$$

We now remove assumptions (i) and (ii). For $T > 0$, define $\Lambda_T \triangleq [-T/2, T/2] \cap \Lambda$. Let $\delta > 0$ and let T_0 be such that for all $T > T_0(\delta)$ there exists $u_T \in \mathbb{R}$ such that

$$\frac{|[u_T, T+u_T] \cap \Lambda|}{T} - \delta < d^-(\Lambda).$$

Let $\varepsilon > 0$ and let $T = T(\varepsilon) > T_0(\delta)$ be such that

$$D_{S(H,\Lambda)}(R) + \varepsilon \geq D_T = \frac{1}{T}\int_{-T/2}^{T/2} \mathbb{E}\left(X(t) - \mathbb{E}\left[X(t)|f(Y_T)\right]\right)^2 dt, \quad (4.39)$$

where Y_T is the finite vector of samples obtained using $\Lambda_T = \Lambda \cap [-T/2, T/2]$ and f is an encoder of Y_T at rate R. Consider the periodic extension $\tilde{\Lambda} \triangleq \Lambda_T + T\mathbb{Z}$ of λ_T. Note that

$$d\left(\tilde{\Lambda}_T\right) = \frac{|[u_T, T+u_T] \cap \Lambda|}{T} < d^-(\Lambda) + \delta. \quad (4.40)$$

We also extend $g(t, t-\tau) = K_H(t,\tau)$ periodically as $\tilde{g}^T(t,\tau) \triangleq g([t],\tau)$, where here and henceforth $[t]$ denotes t *modulo the grid* $u_T + T\mathbb{Z}$ (i.e. $t = [t] + nT + u_T$ where $n \in \mathbb{Z}$ and $0 \leq [t] < T$). Recall from the first part of the proof that periodic sampling is equivalent to MB uniform sampling; and hence the asymptotic law of the samples and the estimator of $X(\cdot)$ from these samples exists. Denote by $Y_{\tilde{\Lambda},\tilde{g}}[\cdot]$ the asymptotic law of the process obtained by sampling with the periodic set $\tilde{\Lambda}$ and the periodic pre-processing system $\tilde{g}(t,\tau)$. For $t_n \in \tilde{\Lambda}$, we have

$$\begin{aligned} Y_{\tilde{\Lambda},\tilde{g}}[n] &= \int_{-\infty}^{\infty} \tilde{g}(t_n, t_n - \tau)\,(X(\tau)+\eta(\tau))\,d\tau \\ &= \begin{cases} \int_{-\infty}^{\infty} g(t_n, t_n - \tau)\,(X(\tau)+\eta(\tau))\,d\tau, & t_n \in \Lambda_T, \\ \int_{-\infty}^{\infty} g([t_n], [t_n] + nT + u_T - \tau)\,(X(\tau)+\eta(\tau))\,d\tau, & t_n \notin \Lambda_T. \end{cases} \end{aligned} \quad (4.41)$$

It follows from (4.41) that $Y_{\Lambda_T,g} \subset Y_{\tilde{\Lambda},\tilde{g}}$, and this implies

$$\frac{1}{T}\int_{-T/2}^{T/2} \mathbb{E}\left(X(t) - \mathbb{E}\left[X(t)|f(Y_T)\right]\right)^2 dt \geq D_{S(\tilde{g},\tilde{\Lambda})}(R). \tag{4.42}$$

for any encoder f of rate R. Note that the process $Y_{\tilde{\Lambda},\tilde{g}}[\cdot]$ is the result of sampling $X(\cdot)$ using the periodic pre-processing $\tilde{g}(t,\tau)$ system and the periodic sampling set $\tilde{\Lambda}$, both of which have period T. By the first part of the proof we have

$$D_{S(\tilde{g},\tilde{\Lambda})}(R) \geq \underline{D}\left(d(\tilde{\Lambda}),R\right). \tag{4.43}$$

From (4.39), (4.42) and (4.43) it follows that for any $\delta > 0$, $\varepsilon > 0$ and $T > T_0(\delta)$ that satisfy (4.39), we have

$$\begin{aligned} D_{S(H,\Lambda)}(R) + \varepsilon &\geq \frac{1}{T}\int_{-T/2}^{T/2} \mathbb{E}\left(X(t) - \mathbb{E}\left[X(t)|f(Y_T)\right]\right)^2 dt \geq D_{S(\tilde{g},\tilde{\Lambda})}(R) \\ &\geq \underline{D}\left(d(\tilde{\Lambda}),R\right) \geq \underline{D}\left(d^-(\Lambda)+\delta,R\right), \end{aligned} \tag{4.44}$$

where the last inequality in (4.44) is due to the monotonicity of $\underline{D}(f_s,R)$ in f_s, as follows from its definition in (4.15). It also follows from this definition that $\underline{D}(f_s,R)$ is continuous in f_s. Since ε and δ can be taken to be arbitrarily small, we have shown that $D_{S(H,\Lambda)}(R) \geq \underline{D}(d^-(\Lambda),R)$. □

It is important to note that Theorem 4.12 holds even if a characterization of $D_{S(H,\Lambda)}(f_s,R)$ in terms of an optimization over joint probability distributions subject to a mutual information rate constraint is unknown. Therefore, the function $\underline{D}(f_s,R)$ provides a fundamental lower bound for the distortion in ADX.

4.5 Chapter Summary

In this chapter we considered the optimal tradeoff among distortion, bitrate and sampling rate in the ADX setting under a bounded linear sampling. We showed that the distortion-rate function describing this tradeoff is given in general as the sum of two terms: (1) the minimal MSE in recovering the source signal from the output of the sampler, and (2) a water-filling expression over the part of the spectrum that is recovered by the estimator that estimates the source signal from this output. In the special case of time-invariant sampling, the minimal distortion in the ADX setting can be derived in a closed form in terms of the sampling rate f_s, the bitrate R, the pre-sampling filters $H_1,\ldots,H_L$, and the PSDs of the source and the noise. Increasing the number of sampling branches with an optimization over their passband has led to a lower bound $\underline{D}(f_s,R)$ that, for a fixed R and f_s, only depends on the PSDs of the noise and the source. Moreover, we showed that $\underline{D}(f_s,R)$ provides a lower bound on the distortion in ADX under any bounded linear sampler – not only uniform time-invariant samplers.

The procedure for attaining the distortion level $\underline{D}(f_s,R)$ follows from combining the optimal coding scheme deduced from Theorem 3.1 with the optimization over the sampling technique that minimizes the MMSE as explored in Section 2.5. This procedure is summarized in the following steps:

(i) Given the bitrate constraint $R+\varepsilon$ and sampling rate f_s, use a multi-branch uniform sampler with a sufficient number of sampling branches L such that the effective passband of all branches is the set $F^\star$. i.e., a set of Lebesgue measure f_s with maximal SNR.

(ii) Estimate the signal $X(\cdot)$ under a MSE criterion. As the time interval goes to infinity, this estimate is given by the process $\widetilde{X}(\cdot)$ defined in (2.20).

(iii) Encode the estimate from (ii) in an optimal manner using $T(R+\varepsilon)$ bits as in standard source coding [6]. Namely, generate $T(R+\varepsilon)$ i.i.d. reconstruction waveforms for $\widehat{X}_T(\cdot)$ from the distribution of the stationary process with PSD given by the preserved part of the spectrum in Figure 4.7. Use minimum distance encoding with respect to these reconstruction waveforms.

In the next chapter we explore properties of the fundamental distortion limit in ADX derived in this chapter and how these properties impact the joint optimization of sampling and lossy compression.

Chapter 5

Joint Optimization of Sampling and Lossy Compression

The minimal distortion in ADX describes the fundamental performance limit of any system involving noise, sampling, and quantization or limited bitrate operation. While the separate effect of each of these information inhibiting phenomena is well understood, the ADX setting allows us to examine their combined effect. In this chapter we explore the joint optimization of the sampling rate and the lossy compression scheme in the ADX setting. We focus in particular on two interesting phenomena that arise from this optimization.

The first phenomenon concerns the minimal sampling rate required in order to attain the minimal distortion in ADX. When the bitrate constraint R is relaxed and no noise is presence, as we saw in Chapter 2, this sampling rate is the spectral occupancy of the signal, i.e. its Landau rate. In Section 5.1 below, we explain that under a finite bitrate constraint R on the output of the encoder, the minimal distortion, described by the DRF of the source signal, is attained by sampling at a rate smaller than the Landau rate. That is, for every bitrate R, there exists a new critical sampling rate above which the minimal distortion in ADX is attained. This new critical sampling rate converges to the Landau rate as R goes to infinity. In fact, even when the bandwidth of the signal is unbounded (as in the case of a Markov process), for any finite bitrate the aforementioned critical sampling rate is still finite. Namely, while it is impossible to sample a non-bandlimited signal without additional distortion due to sampling, it is still possible to sample the signal and attain the minimal distortion resulting from encoding it.

This last observation leads to the second theoretic phenomenon considered in this chapter: the asymptotic number of bits per sample required in order to describe a stationary Gaussian process with vanishing distortion. For band-limited signals, any increase in the bitrate improves accuracy, while the distortion due to sampling vanished for sampling beyond the Landau or Nyquist rates. Therefore, the number of bits per sample may be unbounded in order to achieve vanishing distortion. In Section 5.2, we show that for some non-bandlimited signals, the asymptotic number of bits per sample required to attain vanishing distortion

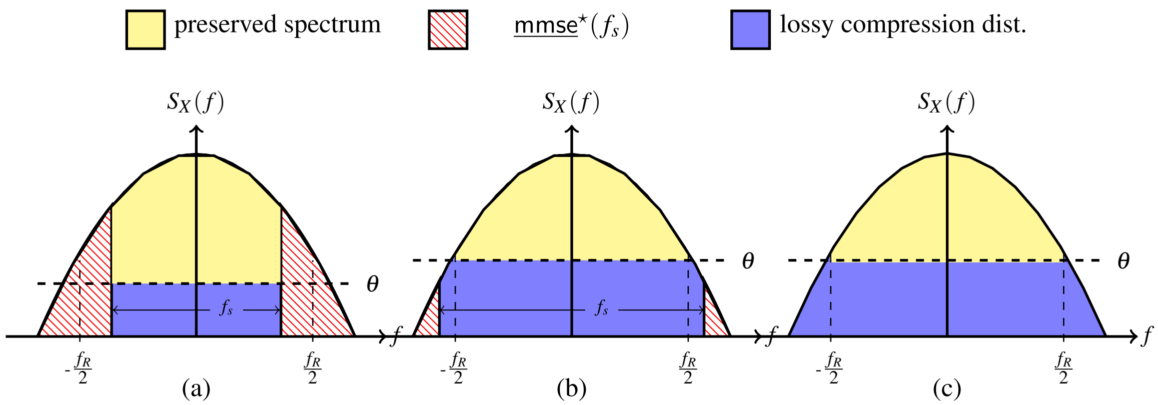

Figure 5.1: Water-filling interpretation for Theorem 5.1 with a unimodal PSD and zero noise. The distortion in each case is the sum of the MMSE and the lossy compression distortion. All figures correspond to the same bitrate R and different sampling rates: (a) $f_s < f_R$, (b) $f_s \geq f_R$ and (c) $f_s > f_{\mathsf{Nyq}}$. The DRF of $X(\cdot)$ is attained for all f_s greater than $f_R < f_{\mathsf{Nyq}}$

converges to a finite value. This observation suggests that in encoding these signals, increasing bit-resolution of the quantizer is futile unless accompanied by a proportional increase in sampling rate.

5.1 Optimal Sampling Rate subject to a Bitrate Constraint

Consider the expression $D_{\mathsf{SB}(H)}(f_s,R)$ of (4.6) for the unimodal PSD shown in Figure 5.1 with zero noise, where the pre-sampling filter H is an ideal LPF with cutoff frequency $f_s/2$. This LPF functions as an anti-aliasing filter, and therefore the part of $D_{\mathsf{SB}(H)}(f_s,R)$ associated with the sampling distortion is only due to those energy bands blocked by the filter. As a result, for this PSD, $D_{\mathsf{SB}(H)}(f_s,R) = \underline{D}(f_s,R)$ and it can be described by the sum of the MMSE and the lossy compression parts in Figure 5.1(a). Figure 5.1(b) describes the function $D_{\mathsf{SB}(H)}(f_s,R)$ under the same bitrate R and a higher sampling rate, while the cutoff frequency of the low-pass filter is adjusted to this higher sampling rate. As can be seen from the figure, at this higher sampling rate $D_{\mathsf{SB}(H)}(f_s,R)$ equals the DRF of $X(\cdot)$ in Figure 5.1(c), although this sampling rate is still below the Nyquist rate of $X(\cdot)$. In fact, it follows from Figure 5.1 that the DRF of $X(\cdot)$ is attained at some critical sampling rate f_R that equals the Nyquist rate of the preserved part in the Pinsker water-filling expression (3.3). The existence of this critical sampling rate can also be seen in Figure 5.2, which illustrates $D_{\mathsf{SB}(H)}(f_s,R)$ as a function of f_s with $H(f)$ a low-pass filter.

The phenomenon described above can be generalized to any Gaussian stationary source and noise in the ADX setting, per the following theorem:

Theorem 5.1 *Let $X(\cdot)$ be a Gaussian stationary process with PSD $S_X(f)$. For each point $(R,D) \in [0,\infty) \times$*

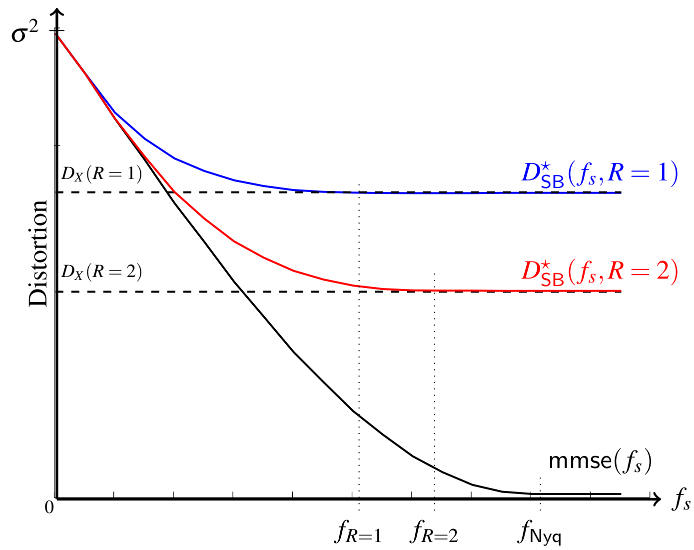

Figure 5.2: The function $D^\star_{\mathsf{SB}}(f_s,R)$ for the PSD of Figure 5.1 with H a LPF of cutoff frequency $f_s/2$ and zero noise, for two values of the bitrate R. Also shown is the DRF of $X(\cdot)$ at these values. Note that these DRFs are attained at the sub-Nyquist sampling rates marked by f_R.

$(0,\sigma_X^2)$ *on the graph of the function* $D_{X|X+\eta}(R)$ *of* (3.10) *associated with a water-level* θ*, let*

$$f_R = \mu\left(\{f : S_X(f) \leq \theta\}\right).$$

Then for all $f_s \geq f_R$*,*

$$D_{X|X+\eta}(R) = \underline{D}(f_s,R),$$

where $\underline{D}(f_s,R)$ *is the fundamental distortion limit in ADX defined in* (4.15).

We emphasize that the critical frequency f_R depends only on the PSDs $S_X(f)$ and $S_\eta(f)$ and on the operating point on the indirect distortion-rate curve. This operating point can be parametrized by either D, R, or the water-level θ using (3.10).

Proof [of Theorem 5.1] Let (R,D) be a point on the graph of $D_{X|X+\eta}(R)$. Let $S_{X|Y}(f) = S_X^2(f)/S_{X+\eta}(f)$ be the PSD of the estimator $\mathbb{E}[X(\cdot)|X(\cdot)+\eta(\cdot)]$. For θ such that

$$R = \frac{1}{2}\int_{-\infty}^{\infty} \log^+\left[S_{X|X+\eta}(f)/\theta\right] df,$$

denote $F_\theta \triangleq \left\{ f \in \mathbb{R} \,:\, S_{X|X+\eta}(f) > \theta \right\}$, so that $f_R = \mu(F_\theta)$,

$$R = \frac{1}{2}\int_{F_\theta} \log \frac{S_{X|X+\eta}(f)}{\theta} df,$$

and

$$D = \sigma_X^2 - \int_{F_\theta} \left(S_{X|X+\eta}(f) - \theta\right) df.$$

Let $F^\star \subset \mathbb{R}$ be such that

$$\underline{D}(f_R,R) = \sigma_X^2 - \int_{F^\star} \left[S_{X|X+\eta}(f)df - \theta\right]^+,$$

and

$$R = \frac{1}{2}\int_{F^\star} \log^+[S_{X|X+\eta}(f)/\theta] df.$$

From the definition of $\underline{D}(f_s,R)$ in (4.15) it follows that

$$\int_{F^\star} S_{X|X+\eta}(f)df \geq \int_{F_\theta} S_{X|X+\eta}(f)df.$$

Hence, Proposition 3.4 implies $\underline{D}(f_R,R) \leq D_{X|X+\eta}(R)$. The reverse inequality follows from Proposition (4.1)-(iii). □

From the definition of $\underline{D}(f_s,R)$ in (4.15), it follows that $\underline{D}(\theta) = \underline{\mathsf{mmse}}(f_s) + f_s\theta$ when $f_s < f_R$. This fact implies that for $D > \underline{\mathsf{mmse}}(f_s)$, the inverse of $\underline{D}(f_s,R)$ with respect to R can be written as

$$\underline{R}(f_s,D) = \frac{1}{2}\int_{-\frac{f_s}{2}}^{\frac{f_s}{2}} \log^+\left(\frac{f_s S_X(f)}{D - \underline{\mathsf{mmse}}(f_s)}\right) df. \tag{5.1}$$

Together with Theorem 5.1, (5.1) implies the following representation for R as a function of the distortion D, which we state as a theorem:

Theorem 5.2 (rate-distortion lower bound) *The bitrate required to encode the samples of a Gaussian stationary process $X(\cdot)$ from the output of a bounded linear sampler of rate f_s and recover $X(\cdot)$ with MSE at most D, is lower bounded by*

$$\underline{R}(f_s,D) = \begin{cases} \frac{1}{2}\int_{-\frac{f_s}{2}}^{\frac{f_s}{2}} \log^+\left(\frac{f_s S_X(f)}{D - \underline{\mathsf{mmse}}(f_s)}\right) df, & f_s < f_R, \\ R_X(D), & f_s \geq f_R, \end{cases} \tag{5.2}$$

for $D > \underline{\mathsf{mmse}}(f_s)$, where $R_{X|X+\eta}(D)$ is the indirect rate-distortion function of $X(\cdot)$ (defined by the inverse of the indirect distortion-rate function of (3.10)).

5.1.1 Discussion

From the definition of the function $\underline{D}(f_s,R)$ in (4.15), it immediately follows that $\underline{D}(f_s,R) = D_{X|X+\eta}(R)$ for all sampling rates above the Landau rate f_{Lan} of $X(\cdot)$. Theorem 5.1 implies that this equality is extended to sampling rates above f_R, which is never larger than f_{Lan}. In fact, the inequality $f_R \leq f_{\mathsf{Lan}}$ is strict whenever $S_{X|X+\eta}(f)$ is not uniformly distributed over $\mathrm{supp}(S_X)$, as in the PSDs $S_\Lambda(f)$, $S_\Omega(f)$, and $S_\omega(f)$ in Figures 5.3, but not $S_\Pi(f)$ there. For such a non-uniformly distributed PSD, there exists a region of values of R for which $f_R < f_{\mathsf{Lan}}$. As shown in Figure 5.1 for the case $\eta(\cdot) \equiv 0$, within this region the MMSE as a result of reduced-rate sampling can be exchanged with lossy compression distortion such that the overall distortion is unaffected.

As R goes to infinity, $\underline{D}(f_s,R)$ converges to $\underline{\mathsf{mmse}}(f_s)$, the water-level θ goes to zero, the set F_θ coincides with the support of $S_X(f)$ and f_R converges to f_{Lan}. Theorem 5.1 then implies that $\underline{\mathsf{mmse}}(f_s) = 0$ for all $f_s \geq f_{\mathsf{Lan}}$, which agrees with Landau's characterization of the condition for perfect recovery of signals in the Paley-Wiener space under nonuniform samples discussed in Section 2.4.3.

As a summary of Theorems 3.1, 4.7 ,4.11, and 5.1, it follows that the procedure for attaining the DRF of $X(\cdot)$ in the ADX setting is as follows:

(i) Fix a desired output bitrate R and a sampling rate $f_s \geq f_R$.

(ii) For $\varepsilon > 0$, choose L large enough that $D^\star_{\mathsf{MB}(L)}(f_s,R) - \underline{D}(f_R,R) < \varepsilon/3$. Use a MB sampler with L sampling branches and pre-sampling filters $H_1,\ldots,H_L$ that attain $D^\star_{\mathsf{MB}(L)}(f_s,R)$.

(iii) Choose a time horizon T large enough such that encoding the MMSE process $\widetilde{X}_T(\cdot)$ of (2.7) using TR bits leads to a distortion not exceeding $D_{\widetilde{X}}(R) + \varepsilon/3$.

Theorem 5.1 implies that the average distortion in reconstructing $X(\cdot)$ under the scheme above does not exceed $D_{X|X+\eta}(R) + \varepsilon$ (or $D_X(R) + \varepsilon$ when η is zero).

An intriguing way to explain the critical sampling rate subject to a bitrate constraint arising from Theorem 5.1, follows by considering the number of degrees of freedom in the representation of the analog signal pre- and post- sampling and lossy compression. Specifically, for stationary Gaussian signals with zero sampling noise, the degrees of freedom in the signal representation are those spectral bands in which the PSD is non-zero. When the signal energy is not uniformly distributed over these bands (unlike the example of the PSD (5.5)), the optimal lossy compression scheme described by Pinsker's expression (3.3) calls for discarding those bands with the lowest energy, i.e., the parts of the signal with the lowest uncertainty. The pre-sampling operation "aligns" the signal prior to sampling and therefore allows for sampling at a rate adjusted to the overall degrees of freedom in the lossy compressed signal representation.

The degree to which the new critical rate f_R is smaller than the Nyquist rate depends on the energy distribution of $X(\cdot)$ along its spectral support. The more uniform it is, the fewer degrees of freedom are reduced

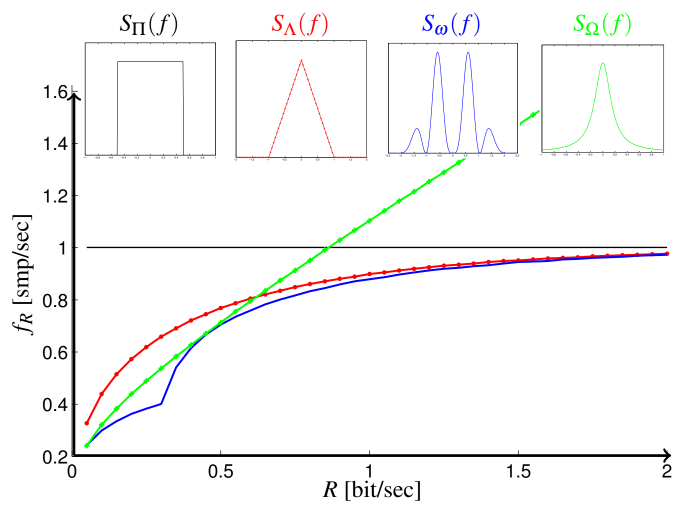

Figure 5.3: Reduction in the critical sampling rate f_R due to bitrate R constraint. For the bandlimited PSDs $S_\Pi(f)$, $S_\Lambda(f)$, and $S_\omega(f)$, the critical sampling rate is always below the Nyquist rate. The critical sampling rate is finite for any R even for the non-bandlimited PSD $S_\Omega(f)$.

due to optimal lossy compression, and therefore the critical sampling rate f_R is closer to the Nyquist rate. As R increases, fewer spectral bands are discarded due to lossy compression, and f_R approaches the Nyquist rate, which, from Corollary 2.6, is the critical sampling rate that allows zero MSE. Figure 5.3 illustrates the dependency of f_R on R for various PSD functions and $S_\eta(f) = 0$. Note that whenever the energy distribution is not uniform and the signal is bandlimited, the critical rate f_R converges to the spectral occupancy as R goes to infinity and converges to zero as R goes to zero.

5.1.2 Examples

In the following examples, the exact dependency of f_R on R and D is found for the various PSDs appearing in Figure 5.3. In all cases it is assumed that the sampling noise $\eta(\cdot)$ equals zero.

Example 4 (triangle PSD) *Consider a Gaussian stationary source with PSD*

$$S_\Lambda(f) = \frac{\sigma_X^2}{f_B}\left[1 - |f/f_B|\right]^+. \tag{5.3}$$

Let $(R,D) \in [0,\infty) \times [0,1]$ *be a point on the distortion-rate curve of* $X(\cdot)$. *It can be shown that* $F_\theta = f_B\left[-1 + f_B\theta, 1 - f_B\theta\right]$ *and* $f_R = 2f_B(1 - f_B\theta/\sigma_X^2)$ *for* $0 < \theta \le \sigma_X^2/f_B$. *The exact relation between* R *and*

f_R is given by

$$R = \frac{1}{2}\int_{-\frac{f_R}{2}}^{\frac{f_R}{2}} \log\left(\frac{1-|f/f_B|}{1-\frac{f_R}{2f_B}}\right) df$$
$$= f_B \log \frac{1}{1-\frac{f_R}{2\ln 2}} - \frac{f_R}{2\ln 2}, \quad 0 \le f_R \le 2f_B. \tag{5.4}$$

Expressing f_R as a function of D leads to $f_R = 2f_B\sqrt{1-D/\sigma_X^2}$.

Example 5 (rectangular PSD) *In the case where the PSD of $X(\cdot)$ is of the form*

$$S_\Pi(f) = \frac{\sigma_X^2}{2f_B}\begin{cases}1, & |f| \le f_B,\\ 0, & |f| > f_B,\end{cases} \tag{5.5}$$

we have that $F_\theta = [-f_B, f_B]$ for all f_s. This implies that $f_R = 2f_B$. Therefore, in this example, $\underline{D}(f_s,R) = D_X(R) = \sigma_X^2 2^{-R/f_B}$ only for f_s larger than the Nyquist rate. The expression for $\underline{D}(f_s,R)$ for $f_s < f_{\mathsf{Nyq}} = 2f_B$ is given by (4.10).

As shown in Figure 5.3, PSD $S_\Omega(f)$ is non-unimodal PSD for which $\underline{D}(f_s,R)$ is attained by MB sampling with a large enough number of sampling branches. The PSD $S_\Omega(f)$ corresponds to the Gauss-Markov process whose bandwidth is infinite. The relation between R and f_R for this process illustrates another interesting phenomenon that is explored in more detail in the next section.

5.2 Sampling Infinite Bandwidth Signals

A common practice in signal processing is to restrict attention to signals that are bandlimited since these can be perfectly represented by their discrete-time samples. Nevertheless, there are many important cases where this assumption does not hold, including Markov processes, autoregressive processes, the Wiener process (or other semi-martingales), and all processes that are used in prediction and filtering theory [48]. In this section we argue that an important contribution of the ADX theory is in describing the optimal tradeoff among distortion, sampling rate, and bitrate, even if the source signal is not non-bandlimited. This tradeoff is best explained by an example.

5.2.1 Example: sampling a Gauss-Markov process

Let $X_\Omega(t)$ be a Gaussian stationary process with PSD

$$S_\Omega(f) = \frac{1/f_0}{(\pi f/f_0)^2+1}, \quad f_0 > 0. \tag{5.6}$$

The signal $X_\Omega(t)$ is also a Markov process, and it is in fact the unique Gaussian stationary process that is also Markovian (a.k.a *Ornstein-Uhlenbeck* process). The relation between R and f_R for $X_\Omega(\cdot)$ can be found in a similar way to Examples 4 and 5 and is given by

$$R = \frac{1}{\ln 2}\left(f_R - f_0 \frac{\arctan\left(\pi f_R / f_0\right)}{\pi/2} \right). \tag{5.7}$$

The relation between f_R to R defined by (5.7) is illustrated in Figure 5.3. The interesting phenomenon illustrated by this figure is that although the Nyquist rate of $X_\Omega(\cdot)$ is infinite, for any finite R there exists a critical sampling frequency f_R, satisfying (5.7), such that the DRF of $X_\Omega(t)$ can be attained by sampling at or above f_R. Namely, when the non-bandlimited signal $X_\Omega(\cdot)$ is considered under a bitrate constraint R, there exists a sampling scheme at rate f_s such that the overall distortion in the system equals the minimal distortion subject only to the bitrate constraint.

Since the distortion in ADX is bounded from below by both the DRF of $X_\Omega(\cdot)$ and the MMSE (as follows from Proposition 4.1), both f_s and R must be taken to infinity in order to describe a realization of $X_\Omega(\cdot)$ with vanishing distortion. In this asymptotic regime, it follows from (5.7) that the relation between the critical sampling rate f_R and R is $R = f_s/\ln 2 + o(1/f_R)$. Therefore, for R sufficiently large, the digital representation of the samples of $X_\Omega(\cdot)$ must allocate $R/f_s = 1/\ln 2 \approx 1.45$ bits per sample in order to attain the DRF of $X_\Omega(\cdot)$. If the number of bits per sample goes below this value, then, as R and the sampling rate increase, the distortion is dominated by the lossy compression distortion as there are not enough bits to represent the information acquired by the sampler. Consequently, as illustrated in Figure 5.4, the ratio between $\underline{D}(f_s,R)$ and $D_X(R)$ goes to one in this case. If the number of bits per sample is above $1/\ln 2$, then the distortion is dominated by the sampling distortion, and the ratio between $\underline{D}(f_s,R)$ and $D_X(R)$ converges to a constant bigger than one as illustrated in Figure 5.4.

5.2.2 Classification of infinite bandwidth signals

The asymptotic value of bits per sample R/f_R distinguishes between two classes of signals:

- (C1) signals for which $R/f_R \to \infty$ as $R \to \infty$
- (C2) signals for which R/f_R is bounded as $R \to \infty$

The class (C1) includes all bandlimited signals, since for these signals f_R is bounded by f_{Lan}. In addition, (C1) also includes non-bandlimited signals whose PSD vanishes quickly. For example consider the PSD $S_\lambda(f) = e^{-0.5|f|}$. For this PSD, the relation between f_R and R is given by

$$R = \frac{1}{2}\left(1 - e^{-f_R/2}\right) + \frac{f_R^2}{4}, \tag{5.8}$$

which implies that R/f_R is unbounded as R goes to infinity.

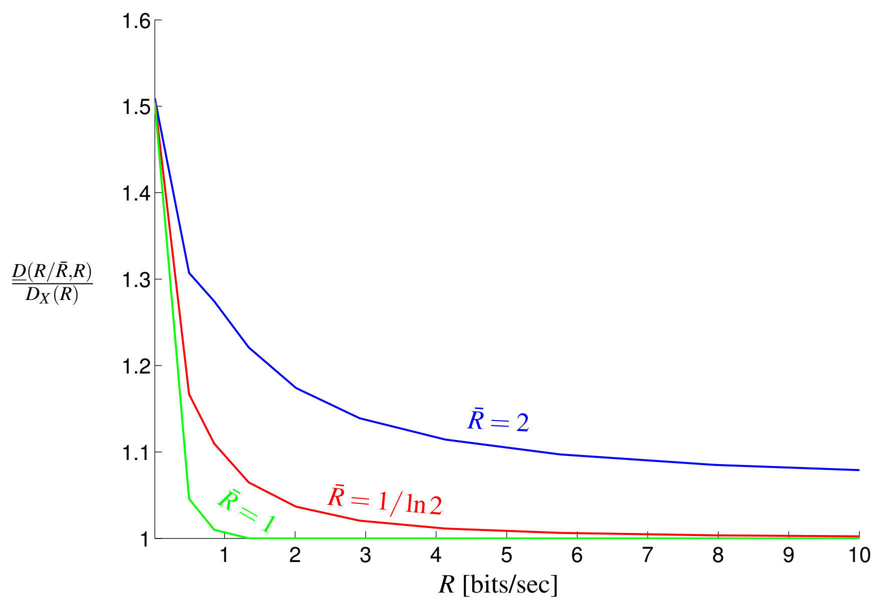

Figure 5.4: The ratio between the minimal distortion in ADX and the DRF for the Gauss-Markov process $X_\Omega(\cdot)$ with a constant number of bits per samples R/f_s. The ADX distortion vanishes at the same rate as the DRF whenever the number of bits per sample $\bar{R} = R/f_s$ is smaller than $1/\ln 2$.

For signals in (C1), there exists a critical sampling rate above which the DRF is attained for any R. Therefore, in order to encode a signal in (C1) with vanishing distortion it is enough to use a finite rate sampler, but the quantizer's bit-resolution must approaches infinity. In other words, the main challenge in encoding these signals with vanishing distortion is the design of quantizers with increasing bit-resolution.

The class (C2) contains only non-bandlimited signals such as the Gauss-Markov process $X_\Omega(\cdot)$. For signals in this class, the ratio $\underline{D}(f_s,R)/D_X(R)$ goes either to one or to a constant bigger than one, depending whether the number of bits per sample R/f_s exceeds the asymptotic ratio R/f_R or not, respectively. Hence, this asymptotic ratio is the minimal number of bits one is required to provide per sample to ensure that the distortion in ADX is not dominated by the lossy compression error. The main challenge in encoding sampling in this class is the requirement to constantly increase the sampling rate side by side with the bitrate in order to improve accuracy of sampled representations. We emphasize that the ADX distortion $\underline{D}(f_s,R)$ always goes to zero whenever both R and f_s goes to infinity.

5.3 Chapter Summary

In this chapter we considered the joint effect of a sampling constraint and a bitrate constraint on the ability to represent an analog signal under the ADX setting. We first showed that for a fixed bitrate R, the minimal distortion in ADX gives rise to a critical sampling rate f_R such that sampling at or above f_R attains the minimal distortion in encoding an analog signal at bitrate R. This rate extends the Landau and Nyquist rates that provide the necessary rates for sampling with distortion-free recovery in classical sampling theory under an infinite precision representation of the samples. When these samples are encoded or compressed in a lossy manner, the minimal distortion is described by the DRF or, if system noise is present, by the indirect DRF of the signal given its noisy version. These functions, however, can be attained by sampling at the rate f_R that is smaller than the Landau rate for signals with non-uniform spectral distribution.

In addition, we considered the asymptotic ratio between bitrate and sampling rate required to attain vanishing distortion under the ADX setting. For signals for which this ratio is finite, this ratio represents the number of bits per sample required to make the MMSE distortion and the lossy compression distortion vanish at the same rate.

Chapter 6

Conclusions and Future Directions

Processing, digital communication, and digital storage of an analog signal is achieved by first representing it as a bit sequence. The restriction on the bitrate of this sequence is the result of restrictions on power, memory, communication, and computation. In addition to bitrate restrictions, hardware and modeling constraints in processing analog information imply that the digital representation is obtained by first sampling the analog signal and then quantizing its samples in a lossy manner. That is, the transformation of an analog signal to bits involves the combination of sampling and lossy compression operations.

In order to account for the joint effect of sampling and bitrate constraints, we presented the analog-to-digital compression (ADX) setting. The ADX setting considers any encoding scheme of the samples at rate R bits per unit time in which these samples are obtained by a bounded linear sampler that produces no more than f_s samples per unit time on average. The class of bounded linear samplers covers most of the sampling systems encountered in practice, including linear time-invariant sampling and nonuniform sampling. Therefore, the minimal distortion in the ADX provides the fundamental performance limit of most of the real world systems that convert analog information to a bit sequence.

This thesis provided a full characterization of the minimal ADX distortion for the important case of a Gaussian stationary input signal corrupted by Gaussian stationary noise. Specifically, for a given bitrate constraint R and a sampling rate constraint f_s, we provided an achievable lower bound on the distortion under any bounded linear sampling system of sampling rate f_s and encoding scheme of bitrate R. This lower bound is given only in terms of the spectral properties of the source signal and the noise. By relaxing the bitrate constraint or the sampling rate constraint, this lower bound recovers the distortion under optimal estimation through linear filtering or optimal lossy compression, respectively. Hence, our characterization of the minimal distortion in ADX unified and generalized classical results in sampling theory and lossy source coding theory in various aspects.

The characterization of the minimal ADX was achieved by first considering the MMSE subject only to the sampling constraint and the noise. Next, we considered the indirect source coding problem of estimating an analog signal from an encoded version of its observations, and ignored the sampling constraint. The MMSE estimation problem and the indirect source coding problem were then combined to a single setting providing the minimal MSE distortion due to the joint effect of sampling and lossy compression. Finally, an optimization over the sampling structure to minimize the distortion was performed, leading to the minimal MSE distortion in ADX.

Our characterization of the minimal distortion in ADX has important theoretical and practical implications in signal processing and information theory. Specifically, the minimal ADX distortion provides an extension of the classical sampling theory of Whittaker-Kotelnikov-Shannon-Landau by describing the minimal sampling rate required for minimizing the distortion in sampling an analog signal. The ADX distortion also leads to a theory of representing signals of infinite bandwidth. Namely, it provides the exact number of bits required to represent each sample of the analog waveform as the number of bits and samples per unit time goes to infinity in order to decrease the distortion to zero. In addition, we concluded that sampling at the Nyquist rate is not necessary for minimizing distortion when working under a bitrate constraint. Arguably, almost any system that processes information is subject to such a constraint due to power, cost or memory limitations of hardware. Moreover, sampling at the critical sampling rate results in the most compact digital representation of the analog signal and provides a mechanism to remove redundancy at the sensing stage.

We conclude this thesis by discussing various extensions of its main setting and results. First in Section 6.1, we consider two specific research directions that follow from the ADX setting. More general extensions are discussed in Section 6.2.

6.1 Extensions

We now present two interesting research directions related to the ADX setting that provide specific examples for the way the results of this thesis can be extended.

6.1.1 Sampling sparse signals subject to a bitrate constraint

The pioneering work of [36] and [35] initiated much work in compressed sensing (CS), in which sparse vector is recovered from its noisy random linear projections. The main principle in CS is that a relatively small number of random linear projections is enough to represent the source, provided it has only few non-zero entries in some basis. The fact that sparse sources possess such a low-dimensional representation justifies the "compressed" part of CS. Nevertheless, reducing dimensions does not yet provide compression in the information theoretic sense, since it is still required to quantize this low-dimensional representation, i.e., to map it to a finite alphabet set. Therefore, the compressibility of a sparse signal can be studied using the ADX framework adapted to a CS setting, as was recently considered in [37].

An interesting open question that arises in the ADX setting under a sparse source signal assumption concerns the effect of the bit restriction on the sampling ratio required for optimal reconstruction. Specifically, let X^n be an n-dimensional i.i.d. vector in which each entry is standard normal with probability p, or 0 with probability $1-p$. It is well-known that with m random linear samples of infinite bit precision, perfect recovery of X^n in the noiseless case, or bounded noise sensitivity in the noisy case [70], is possible only if $p > m/n$ [71]. Therefore, p represents the number of degrees of freedom in representing a sparse signal from its random linear projections, and hence can be seen as the compressed sensing counterpart of the Landau rate. However, under an overall bit budget of nR bits to represent the samples, reconstruction is only possible to a distortion level equal to the DRF of X^n, obtained by an encoder that observes it directly.

An interesting open question is whether, subject to a bit per source symbol constraint, the DRF can be attained using fewer then np samples. If the optimal direct encoding scheme of X^n that attains its DRF reduces the degrees of freedom in the signal representation (in analogy to Pinsker's formula (3.3)), then its optimal lossy compressed version may have $n\widehat{p} < np$ non-zeros. As a result, its compressed version can be perfectly represented using a bit more than $n\widehat{p}$ measurements, and no additional distortion is introduced due to sampling at $\widehat{p} < p$. If, on the other hand, the number of degrees of freedom after optimal lossy compression of X^n is equivalent to np, then no reduction in sampling rate of the aforementioned kind is possible.

As a result, its lossy compressed version may have fewer than np non-zeros and, thus, can be perfectly represented using fewer than np measurements. On the other hand, a negative answer would imply that all degrees of freedom are used under direct encoding

6.1.2 Compress-and-Estimate – source coding with oblivious encoding

The characterization of the minimal ADX distortion implies that, in order to attain it, the encoder first estimates the analog signal from its samples and then encodes this estimate. This estimation prior to encoding, however, seems infeasible in many important scenarios. First, this estimate depends on the source's PSD, the sampling rate, and the noise intensity. In the absence of information about any of these factors, an optimal source estimate cannot be obtained. For example, this scenario occurs in statistical analysis and machine learning, since the underlying signal characterization is unknown at the time the data is collected. Moreover, the estimation of the source may require expensive resources that are unavailable to the encoder, due to power constraints for example.

In lieu of estimating the signal prior to encoding, the encoder may simply use an optimal source code adapted to its raw samples. The decoder, aware of the full source's statistics and not limited by resources, would try to estimate the source from this sub-optimal encoding. That is, as opposed to the optimal "estimate-and-compress" strategy in indirect source coding, the encoder and the estimator apply a "compress-and-estimate" (CE) strategy.

As shown in [72] for simple indirect settings, source coding under a CE strategy can lead to highly sub-optimal performance compared to the indirect distortion-rate function. Hence, it is suggested that the tradeoff among distortion, bitrate, and sampling rate in the ADX setting under a CE strategy would yield different conclusions than the ones reported here for the optimal strategy. As an example, consider the case in which a noisy source is sampled above its Nyquist rate. As shown in this thesis, the ADX distortion is fixed for any super-Nyquist sampling rate since the decoder estimates the original signal from the samples. On the other hand, under the CE scenario, the source encoder cannot separate the original analog signal from noise prior to encoding. Therefore, a higher sampling rate would introduce more noise into the system and may increase the overall distortion.

6.2 General Extensions

In addition to the two specific research directions presented above, the results of this thesis can be extended by generalizing the ADX framework. Several such extensions and the motivation for them are briefly discussed below.

6.2.1 Statistical inference under sampling and bitrate constraints

The ADX setting focuses on the estimation of analog signals under an MSE criterion. A more general setting explores the performance of any inference problem or decision that is based on a lossy compressed version of the samples of an analog signal. As an example, we may consider the tradeoff between probability of error, sampling rate and distortion in hypothesis testing based on the lossy compressed samples. In another example, we assume that the distribution of the analog signal depends on some vector of parameters and consider the MSE as a function of the sampling rate and bitrate in estimating this vector.

6.2.2 Optimal tradeoff between bit precision and sampling rate under non-ideal encoding

The minimal distortion in the ADX setting is achieved by using an ideal encoder. That is, given the noisy samples of the analog signal, the encoder provides the optimal description of the original signal to the decoder subject only to the bitrate constraints. In practice, analog-to-digital implementations may impose other restrictions such as limited memory at the encoder and causality. In order to understand the performance limits of such systems, it would be beneficial to extend our model to incorporate such restrictions. In particular, it is interesting to understand which restrictions lead to a non-trivial tradeoff between the average number of bits per second used to represent the process and the sampling rate of the system.

6.2.3 Multiterminal sampling and lossy compression

Another natural extension of the ADX setting is obtained by assuming that noisy samples of the analog signal are observed and encoded in more than one location. This scenario arises, for example, in sensor networks, where different noisy versions of the same underlying analog signal are observed by multiple sensors, where each sensor transmits information to a central estimation unit independently of the other sensors. The goal is to design a distributed encoding and sampling strategy for the sensors in order to minimize the distortion in estimating the underlying analog signal at the central estimator. This setting can be seen as a combined problem of sampling and the indirect multiterminal source coding setting of [73].

Appendix A

The Distortion-Rate Function of Cyclostationary Processes

The proof of Theorem 4.3 is based on an expression for the DRF of a pulse-amplitude modulated (PAM) process given in terms of the PSD of the baseband signal and the Fourier transform of the modulating pulse. This expression is a special case of a more general result that characterizes the DRF of any Gaussian cyclostationary process (CSP) in term its spectral properties. This general result and its application for the DRF of a PAM processes are dervied in this appendix. The content of this appendix first appeared in [19].

Definitions and Notations

Cyclostationary Processes

Throughout this appendix, we consider zero mean Gaussian processes in both discrete and continuous time. We use round brackets to denote a continuous time index and square brackets for a discrete time index.

The statistics of a zero mean Gaussian process $X(\cdot)$ is specified in terms of its autocorrelation function

$$R_X(t,\tau) \triangleq \mathbb{E}\left[X(t+\tau)X(t)\right].$$

We note that in [52] and in other references, the symmetric auto-correlation function

$$\tilde{R}_X(t,\tau) \triangleq \mathbb{E}\left[X(t+\tau/2)X(t-\tau/2)\right] = R_X(t-\tau/2,\tau),$$

the corresponding CPSD $\hat{\tilde{S}}^n_X(f)$ and TPSD $\tilde{S}^{\,t}_X(f)$, are used. The conversion between $\hat{S}^n(f)$ and the symmetric CPSD is given by $\hat{\tilde{S}}^n_X(f) = \hat{S}^n_X(f-n/(2T_0))$. If in addition the autocorrelation function is periodic in t with a fundamental period T_0,

$$R_X(t+T_0,\tau) = R_X(t,\tau),$$

then we say that $X(\cdot)$ is a *CSP* [49, 52]. We also assume that $R_X(t,\tau)$ is bounded and Riemann integrable on $[0,T_0]\times\mathbb{R}$, and therefore

$$\sigma_X^2 = \lim_{T\to\infty}\frac{1}{2T}\int_{-T}^{T}\mathbb{E}X(t)^2dt = \frac{1}{T_0}\int_0^{T_0}R_X(t,0)dt$$

is finite.

Suppose that $R_X(t,\tau)$ has a convergent Fourier series representation in t for almost any $\tau\in\mathbb{R}$. Then the statistics of $X(\cdot)$ are uniquely determined by the *cyclic autocorrelation* (CA) function:

$$\hat{R}_X^n(\tau) \triangleq \frac{1}{T_0}\int_{-T_0/2}^{T_0/2}R_X(t,\tau)e^{-2\pi int/T_0}dt, \quad n\in\mathbb{Z}. \tag{A.1}$$

The Fourier transform of $\hat{R}_X^n(\tau)$ with respect to τ is denoted as the *cyclic power spectral density* (CPSD) function:

$$\hat{S}_X^n(f) = \int_{-\infty}^{\infty}\hat{R}_X^n(\tau)e^{-2\pi i\tau f}d\tau, \quad -\infty\leq f\leq\infty. \tag{A.2}$$

If $\hat{S}_X^n(f)$ is identically zero for all $n\neq 0$, then $R_X(t,\tau)=R_X(0,\tau)$ for all $0\leq t\leq T$ and the process $X(\cdot)$ is stationary. In such a case $S_X(f)\triangleq\hat{S}_X^0(f)$ is the *power spectral density* (PSD) function of $X(\cdot)$. The *time-varying power spectral density* (TPSD) function [52, Sec. 3.3] of $X(\cdot)$ is defined by the Fourier transform of $R_X(t,\tau)$ with respect to τ, i.e.

$$S_X^t(f) \triangleq \int_{-\infty}^{\infty}R_X(t,\tau)e^{-2\pi if\tau}d\tau. \tag{A.3}$$

The Fourier series representation implies that

$$S_X^t(f) = \sum_{n\in\mathbb{Z}}\hat{S}_X^n(f)e^{2\pi int/T_0}. \tag{A.4}$$

Associated with every CSP $X(\cdot)$ with period T_0 is a set of stationary discrete time processes $X^t[\cdot]$, $0\leq t\leq T_0$, defined by

$$X^t[n] = X\left(T_0n+t\right), \quad n\in\mathbb{Z}. \tag{A.5}$$

These processes are called the *polyphase components* (PC) of the CSP $X(\cdot)$. The cross-correlation function of $X^{t_1}[\cdot]$ and $X^{t_2}[\cdot]$ is given by

$$\begin{aligned}R_{X^{t_1}X^{t_2}}[n,k] &= \mathbb{E}\left[X[T_0(n+k)+t_1]X[T_0n+t_2]\right]\\ &= R_X\left(T_0n+t_2, T_0k+t_1-t_2\right)\\ &= R_X\left(t_2, T_0k+t_1-t_2\right).\end{aligned} \tag{A.6}$$

Since $R_{X^{t_1}X^{t_2}}[n,k]$ depends only on k, this implies that $X^{t_1}[\cdot]$ and $X^{t_2}[\cdot]$ are jointly stationary. The PSD of $X^t[\cdot]$ is given by

$$\begin{aligned} S_{X^t}\left(e^{2\pi i\phi}\right) &\triangleq \sum_{k\in\mathbb{Z}} R_{X^tX^t}[0,k]e^{-2\pi i\phi k} \\ &= \sum_{k\in\mathbb{Z}} R_X\left(t,T_0k\right)e^{-2\pi i\phi k}, \quad -\frac{1}{2}\le\phi\le\frac{1}{2}. \end{aligned} \tag{A.7}$$

Using the spectral properties of sampled processes, we can use (A.7) and (A.4) to connect the functions $S_{X^t}\left(e^{2\pi i\phi}\right)$ and the CPSD of $X(\cdot)$ as follows:

$$S_{X^t}\left(e^{2\pi i\phi}\right) = \frac{1}{T_0}\sum_{k\in\mathbb{Z}} S_X^t\left(\frac{\phi-k}{T_0}\right) = \frac{1}{T_0}\sum_{k\in\mathbb{Z}}\sum_{n\in\mathbb{Z}} \hat{S}_X^n\left(\frac{\phi-k}{T_0}\right)e^{2\pi int/T_0}.$$

More generally, for $t_1,t_2\in[0,T_0]$ we have

$$\begin{aligned} S_{X^{t_1}X^{t_2}}\left(e^{2\pi i\phi}\right) &= \sum_{k\in\mathbb{Z}} R_{X^{t_1}X^{t_2}}[0,k]e^{-2\pi ik\phi} \\ &= \frac{1}{T_0}\sum_{k\in\mathbb{Z}} S_X^{t_2}\left(\frac{\phi-k}{T_0}\right)e^{2\pi i(t_1-t_2)\frac{\phi-k}{T_0}} \\ &= \frac{1}{T_0}\sum_{k\in\mathbb{Z}}\sum_{m\in\mathbb{Z}} \hat{S}_X^m\left(\frac{\phi-k}{T_0}\right)e^{2\pi i\left(m\frac{t_2}{T_0}+\frac{t_1-t_2}{T_0}(\phi-k)\right)}. \end{aligned} \tag{A.8}$$

We now turn to briefly describe the discrete-time counterpart of the CA, CPSD, TPSD and the polyphase components defined in (A.1), (A.2), (A.3) and (A.5), respectively.

A discrete time zero mean Gaussian process $X[\cdot]$ is said to be a CSP with period $M\in\mathbb{N}$ if its covariance function

$$R_X[n,k] = \mathbb{E}\left[X[n+k]X[n]\right]$$

is periodic in k with period M. For $m=0,\ldots,M-1$, the m^{th} cyclic autocorrelation (CA) function of $X[\cdot]$ is defined as

$$\hat{R}_X^m[k] \triangleq \sum_{n=0}^{M-1} R_X[n,k]e^{-2\pi inm/M}.$$

The m^{th} CPSD function is then given by

$$\hat{S}_X^m\left(e^{2\pi i\phi}\right) \triangleq \sum_{k\in\mathbb{Z}} \hat{R}_X^m[k]e^{-2\pi i\phi k},$$

and the discrete TPSD function is

$$S_X^n\left(e^{2\pi i\phi}\right) \triangleq \sum_{k\in\mathbb{Z}} R_X[n,k]e^{-2\pi i\phi k}.$$

Finally, we have the discrete time Fourier transform relation

$$S_X^n\left(e^{2\pi i\phi}\right) = \frac{1}{M}\sum_{m=0}^{M-1}\hat{S}_X^m\left(e^{2\pi i\phi}\right)e^{2\pi i\phi nm/M}.$$

The m-th stationary component $\bar{X}^m[\cdot]$, $0 \le m \le M-1$ of $X[\cdot]$ is defined by

$$X^m[n] \triangleq X[Mn+m], \quad n \in \mathbb{Z}. \tag{A.9}$$

For $0 \le m,r,n \le M-1$ and $k \in \mathbb{Z}$ we have

$$\begin{aligned} R_{X^mX^r}[n,k] &= \mathbb{E}\left[X^m[n+k]X^r[n]\right] \\ &= \mathbb{E}\left[X[Mn+Mk+m]X[Mn+r]\right] \\ &= R_X[Mn+r, Mk+m-r] \\ &= R_X[r, Mk+m-r]. \end{aligned} \tag{A.10}$$

Using properties of multi-rate signal processing:

$$S_{X^mX^r}\left(e^{2\pi i\phi}\right) = \sum_{k\in\mathbb{Z}} R_X[r,Mk+m-r]e^{-2\pi ik\phi} = \frac{1}{M}\sum_{n=0}^{M-1} S_X^r\left(e^{2\pi i\frac{\phi-n}{M}}\right)e^{2\pi i(m-r)\frac{\phi-n}{M}}. \tag{A.11}$$

The discrete-time counterpart of (A.8) is then

$$S_{X^mX^r}\left(e^{2\pi i\phi}\right) = \frac{1}{M}\sum_{k=0}^{M-1}\sum_{n=0}^{M-1}\hat{S}_X^n\left(e^{2\pi i\frac{\phi-k}{M}}\right)e^{2\pi i\frac{nr+(m-r)(\phi-k)}{M}}. \tag{A.12}$$

The functions $S_{X^mX^r}\left(e^{2\pi i\phi}\right)$, $0 \le m,r \le M-1$ define an $M \times M$ matrix $\mathbf{S}_X\left(e^{2\pi i\phi}\right)$ with $(m+1,r+1)^{\text{th}}$ entry $S_{X^mX^r}\left(e^{2\pi i\phi}\right)$. This matrix completely determines the statistics of $X[\cdot]$, and can be seen as the PSD matrix associated with the stationary vector valued process $\mathbf{X}^M[n]$ defined by the stationary components of $X[\cdot]$:

$$\mathbf{X}^M[n] \triangleq \left(X^0[n],\dots,X^{M-1}[n]\right), \quad n \in \mathbb{Z}. \tag{A.13}$$

We denote the autocorrelation matrix of $\mathbf{X}^M[\cdot]$ as the PSD-PC matrix. Note that the $(r+1,m+1)^{\text{th}}$ entry of the PSD-PC matrix is given by (A.10).

Example 6 (pulse-amplitude modulation (PAM)) *Consider a Gaussian stationary process $U(\cdot)$ modulated by a deterministic signal $p(t)$ as follows:*

$$X_{\mathsf{PAM}}(t) = \sum_{n\in\mathbb{N}} U(nT_0)p(t-nT_0). \tag{A.14}$$

This process is cyclostationary with period T_0 and CPSD [74, Eq. 49]

$$\hat{S}^n_{\mathsf{PAM}}(f) = \frac{1}{T_0} P(f) P^* \left(f - \frac{n}{T_0} \right) S_U(f), \quad n \in \mathbb{Z}, \tag{A.15}$$

where $P(f)$ is the Fourier transform of $p(t)$ and $P^*(f)$ is its complex conjugate. If T_0 is small enough such that the support of $P(f)$ is contained within the interval $\left(-\frac{1}{2T_0}, \frac{1}{2T_0}\right)$, then $\hat{S}^n_{\mathsf{PAM}}(f) = 0$ for all $n \neq 0$, which implies that $X_{PAM(\cdot)}$ is stationary.

The distortion-rate function

For a fixed $T > 0$, let X_T be the reduction of $X(\cdot)$ to the interval $[-T,T]$. Define the distortion between two waveforms $x(\cdot)$ and $y(\cdot)$ over the interval $[-T,T]$ by

$$d_T(x(\cdot), y(\cdot)) \triangleq \frac{1}{2T} \int_{-T}^{T} (x(t) - y(t))^2 dt. \tag{A.16}$$

We expand X_T by a Karhunen-Loève (KL) expansion [75, Ch 9.7] as

$$X_T(t) = \sum_{k=1}^{\infty} X_k f_k(t), \quad -T \leq t \leq T, \tag{A.17}$$

where $\{f_k\}$ is a set of orthogonal functions over $[-T,T]$ satisfying the Fredholm integral equation

$$\lambda_k f_k(t) = \frac{1}{2T} \int_{-T}^{T} K_X(t,s) f_k(s) ds, \quad t \in [-T,T], \tag{A.18}$$

with corresponding eigenvalues $\{\lambda_k\}$, and where

$$K_X(t,s) \triangleq \mathbb{E}X(t)X(s) = R_X(s, t-s).$$

Assuming a similar expansion as (A.17) to an arbitrary random waveform Y_T, we have

$$\mathbb{E}d_T(X_T, Y_T) = \frac{1}{2T} \int_{-T}^{T} \mathbb{E}(X(t) - Y(t))^2 dt = \sum_{n \in \mathbb{Z}} \mathbb{E}(X_n - Y_n)^2.$$

The mutual information between $X(\cdot)$ and $Y(\cdot)$ on the interval $[-T,T]$ is defined by

$$I_T(X(\cdot), Y(\cdot)) \triangleq \frac{1}{2T} \lim_{N \to \infty} I\left(\mathbf{X}^N_{-N}; \mathbf{Y}^N_{-N}\right),$$

where $\mathbf{X}^N_{-N} = (X_{-N}, \ldots, X_N)$, $\mathbf{Y}^N_{-N} = (Y_{-N}, \ldots, Y_N)$ and the X_ns and Y_ns are the coefficients in the KL expansion of $X(\cdot)$ and $Y(\cdot)$, respectively.

Denote by $\mathscr{P}_T$ the set of joint probability distributions $P_{X,\hat{X}}$ over the waveforms $\left(X(\cdot), \hat{X}(\cdot)\right)$, such that the marginal of $X(\cdot)$ agrees with the original distribution, and the average distortion $\mathbb{E}d_T\left(X(\cdot), \hat{X}(\cdot)\right)$ does not exceed D. The rate-distortion function (RDF) of $X(\cdot)$ is defined by

$$R(D) = \lim_{T\to\infty} R_T(D),$$

where

$$R_T(D) = \inf I_T\left(X(\cdot);\hat{X}(\cdot)\right),$$

and the infimum is over the set $\mathscr{P}_T$. It is well known that $R(D)$ and $R_T(D)$ are non-increasing convex functions of D, and therefore continuous in D over any open interval [6]. We define their inverse function as the distortion-rate functions $D(R)$ and $D_T(R)$, respectively. We note that by its definition, $D(R)$ is bounded from above by the average power of $X(\cdot)$ over a single period:

$$\begin{aligned}\sigma_X^2 &\triangleq \lim_{T\to\infty}\frac{1}{2T}\int_{-T}^{T}\mathbb{E}X^2(t)dt = \lim_{T\to\infty}\frac{1}{2T}\int_{-T}^{T}R_X(t,0)dt \\ &= \frac{1}{T_0}\int_0^{T_0}R_X(t,0)dt = \hat{R}_X^0(0).\end{aligned}$$

For Gaussian processes, we have the following parametric representation for $R_T(D)$ or $D_T(R)$ [75, Eq. 9.7.41]

$$D_T(\theta) = \sum_{k=1}^{\infty}\min\{\theta,\lambda_k\} \tag{A.19a}$$

$$R_T(\theta) = \frac{1}{2}\sum_{k=1}^{\infty}\log^+(\lambda_k/\theta), \tag{A.19b}$$

where $\log^+ x \triangleq \max\{\log x, 0\}$.

In the discrete-time case the DRF is defined in a similar way as in the continuous-time setting described above by replacing the continuous-time index in (A.16), (A.17) and (A.18), and by changing integration to summation. Since the KL transform preserves norm and mutual information, this definition of the DRF in the discrete-time case is consistent with standard expressions for the DRF of a discrete-time source with memory as in [6, Ch. 4.5.2]. Note that with these definitions, the continuous-time distortion is measured in MSE per time unit while the discrete-time distortion is measured in MSE per source symbol. Similarly, in continuous-time, R represents bitrate, i.e., the number of bits per time unit. In the discrete-time setting we use the notation $\bar{R}$ to denote bits per source symbol.

Since the distribution of a zero-mean Gaussian CSP with period T_0 is determined by its second moment $R_X(t,\tau)$, we observe form [65, Exc. 6.3.1] these processes are *asymptotic mean stationary* as defined in [Ch. 1.7][66]. A source coding theorem for this class of processes in discrete-time can be found in [Ch. 9][66]. Its extension to continuous-time follows immediately as long as the *flow* defined by the process, i.e. the mapping from the time index set $[-T,T]$ to the probability space, is measurable. This last condition is implicit in our definition of a continuous-time stationary process in terms of its finite dimensional probability distributions

[69].

Evaluating the DRF of a CSP

In the special case in which $X(\cdot)$ is stationary, it is possible to obtain $D(R)$ without explicitly solving the Fredholm equation (A.18) or evaluating the KL eigenvalues: in this case, the density of these eigenvalues converges to the PSD $S_X(f)$ of $X(\cdot)$. This leads to the celebrated *reverse water-filling* expression for the DRF of a stationary Gaussian process, originally derived by Pinsker [62]:

$$R(\theta) = \frac{1}{2}\int_{-\infty}^{\infty} \log^{+}\left[S_X(f)/\theta\right] df. \tag{A.20a}$$

$$D(\theta) = \int_{-\infty}^{\infty} \min\left\{S_X(f), \theta\right\} d\phi. \tag{A.20b}$$

The discrete-time version of (A.20) is given by

$$\bar{R}(\theta) = \frac{1}{2}\int_{-\frac{1}{2}}^{\frac{1}{2}} \log^{+}\left[S_X\left(e^{2\pi i\phi}\right)/\theta\right] d\phi. \tag{A.21a}$$

$$D(\theta) = \int_{-\frac{1}{2}}^{\frac{1}{2}} \min\left\{S_X\left(e^{2\pi i\phi}\right), \theta\right\} d\phi. \tag{A.21b}$$

Equations (A.20) and (A.21) define the distortion as a function of the rate through a joint dependency on the water level parameter θ.

We note that stationarity is not a necessary condition for the existence of a density function for the eigenvalues in the KL expansion. For example, such a density function is known for the Wiener process [76] which is a non-stationary process.

When $X(\cdot)$ is a general Gaussian CSP, its DRF can be evaluated, in principle, by the following procedure: compute the KL eigenvalues in (A.18) for each T, use (A.19) to obtain $D_T(R)$, and finally take the limit as T goes to infinity. For general CSPs, however, an easy way to describe the density of the KL eigenvalues is in general unknown. As a result, the evaluation of the DRF directly by the KL eigenvalues usually does not lead to a closed-form solution. We therefore turn to an alternative representation for the function $D(R)$ which is based on an approximation of the kernel $K_X(t,s)$ used in (A.18).

A Spectral Representation of the DRF

Our first observation is that in the discrete-time case, the DRF of a Gaussian CSP can be obtained by an expression for the DRF of a vector Gaussian stationary source. This expression is an extension of (A.21), which was derived in [67, Eq. (20) and (21)] and is given as follows:

$$D_{\mathbf{X}}(\theta) = \frac{1}{M}\sum_{m=1}^{M}\int_{-\frac{1}{2}}^{\frac{1}{2}} \min\left\{\lambda_m\left(e^{2\pi i\phi}\right),\theta\right\} d\phi \tag{A.22a}$$

$$R(\theta) = \frac{1}{M}\sum_{m=1}^{M}\int_{-\frac{1}{2}}^{\frac{1}{2}} \frac{1}{2}\log^{+}\left[\lambda_m\left(e^{2\pi i\phi}\right)/\theta\right] d\phi, \tag{A.22b}$$

where $\lambda_1\left(e^{2\pi i\phi}\right),...,\lambda_M\left(e^{2\pi i\phi}\right)$ are the eigenvalues of the PSD matrix $\mathbf{S_X}\left(e^{2\pi i\phi}\right)$ at frequency ϕ. We have the following result:

Theorem A.1 *Let $X[\cdot]$ be a discrete-time Gaussian CSP with period $M\in\mathbb{N}$. The distortion rate function of $X[\cdot]$ is given by*

$$D(\theta) = \frac{1}{M}\sum_{m=1}^{M}\int_{-\frac{1}{2}}^{\frac{1}{2}} \min\left\{\lambda_m\left(e^{2\pi i\phi}\right),\theta\right\} d\phi \tag{A.23a}$$

$$\bar{R}(\theta) = \frac{1}{2M}\sum_{m=1}^{M}\int_{-\frac{1}{2}}^{\frac{1}{2}} \log^{+}\left[\lambda_m\left(e^{2\pi i\phi}\right)/\theta\right] d\phi, \tag{A.23b}$$

where $\lambda_1\left(e^{2\pi i\phi}\right)\leq\ldots\leq\lambda_M\left(e^{2\pi i\phi}\right)$ are the eigenvalues of the PSD-PC matrix with $(m+1,r+1)^{th}$ entry given by

$$\mathbf{S}_{X^mX^r}\left(e^{2\pi i\phi}\right) = \frac{1}{M}\sum_{n=0}^{M-1} S_X^r\left(e^{2\pi i\frac{\phi-n}{M}}\right)e^{2\pi i(m-r)\frac{\phi-n}{M}}. \tag{A.24}$$

Proof Consider the vector valued process $\mathbf{X}^M[\cdot]$ defined in (A.13). The rate-distortion function of $\mathbf{X}^M[\cdot]$ is given by (A.22):

$$D(\theta) = \frac{1}{M}\sum_{m=1}^{M}\int_{-\infty}^{\infty} \min\left\{\lambda_m\left(e^{2\pi i\phi}\right),\theta\right\} d\phi, \tag{A.25a}$$

$$R(\theta) = \frac{1}{2}\sum_{m=1}^{M}\int_{-\infty}^{\infty} \log^{+}\left[\lambda_m\left(e^{2\pi i\phi}\right)/\theta\right] d\phi, \tag{A.25b}$$

where $0\leq\lambda_1\left(e^{2\pi i\phi}\right)\leq\ldots\leq\lambda_M\left(e^{2\pi i\phi}\right)$ are the eigenvalues of the spectral density matrix $\mathbf{S}_{\mathbf{X}^M}\left(e^{2\pi i\phi}\right)$ obtained by taking the Fourier transform of covariance matrix $\mathbf{R_X}[k]=\mathbb{E}\left[X^M[n+k](X^M[n])^T\right]$ entry-wise. The $(m,r)^{\text{th}}$ entry of $\mathbf{S}_{\mathbf{X}^M}\left(e^{2\pi i\phi}\right)$ is given by (A.11):

$$\left(\mathbf{S}_{\mathbf{X}^M}\left(e^{2\pi i\phi}\right)\right)_{m,r} = S_X^{m,r}\left(e^{2\pi i\phi}\right) = \frac{1}{M}\sum_{k=0}^{M-1} S_X^r\left(e^{2\pi i\frac{\phi-k}{M}}\right)e^{2\pi i(m-r)\frac{\phi-k}{M}}. \tag{A.26}$$

It is left to show that the DRF of $\mathbf{X}^M[\cdot]$ coincides with the DRF of $X[\cdot]$. By the source coding theorem for AMS processes [66, Thm. 11.4.1] it is enough to show that the operational block coding distortion-rate function ([66, Ch. 11.2]) of both processes is identical. Indeed, any N block codebook for $\mathbf{X}^M[\cdot]$ is an MN

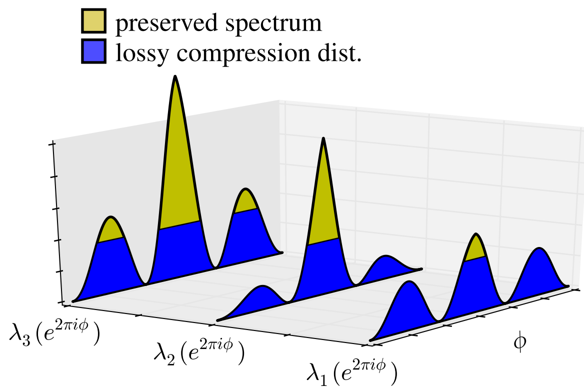

Figure A.1: water-filling interpretation of (A.23) for $M = 3$. The lossy compression distortion and the preserved spectrum are associated with equations (A.23a) and (A.23b), respectively.

block codebook for $X[\cdot]$ which achieves the same quadratic distortion averaged over the block. However, since $\mathbf{X}^M[\cdot]$ is stationary, by [66, Lemma. 11.2.3] we know that any distortion above the DRF of $\mathbf{X}^M[\cdot]$ is attained for large enough N. This implies that the same is true for $X[\cdot]$. □

Equation (A.23) has the water-filling interpretation illustrated in Figure A.1: the DRF is obtained by setting a single water-level over all eigenvalues of (A.24). These eigenvalues can be seen as the PSD of M independent processes obtained by the orthogonalization of the PC of $X[\cdot]$. The overall area below the water-level is the spectral density of the noise term in the test channel that attains the information expression of the DRF, while the area above this level is associated with the reconstructed signal in this channel [6]. Compared to the limit in the discrete-time version of the KL expansion, expression (A.23) exploits the cyclostationary structure of the process by using its spectral properties. These spectral properties capture information on the entire time-horizon and not only over a finite blocklength as in the KL expansion.

The following theorem explains how to extend the above evaluation to the continuous-time case.

Theorem A.2 *Let $X(\cdot)$ be a Gaussian CSP with period T_0 and correlation function $R_X(t,\tau)$ Lipschitz continuous in its second argument. For a given $M \in \mathbb{N}$, denote*

$$D_M(\theta_M) = \frac{1}{M}\sum_{m=1}^{M}\int_{-\frac{1}{2}}^{\frac{1}{2}} \min\left\{\lambda_m\left(e^{2\pi i\phi}\right), \theta_M\right\} d\phi \tag{A.27a}$$

$$R(\theta_M) = \frac{1}{2T_0}\sum_{m=1}^{M}\int_{-\frac{1}{2}}^{\frac{1}{2}} \log^+\left[\lambda_m\left(e^{2\pi i\phi}\right)/\theta_M\right] d\phi, \tag{A.27b}$$

where $\lambda_1\left(e^{2\pi i\phi}\right) \leq \ldots \leq \lambda_M\left(e^{2\pi i\phi}\right)$ are the eigenvalues of the matrix $\mathbf{S}_X\left(e^{2\pi i\phi}\right)$ with its $(m+1,r+1)^{th}$ entry *given by*

$$\begin{aligned} &\frac{1}{T_0}\sum_{k\in\mathbb{Z}} S_X^{rT_0/M}\left(\frac{\phi-k}{T_0}\right) e^{2\pi i(m-r)\frac{\phi-k}{M}} \\ &\quad = \frac{1}{T_0}\sum_{k\in\mathbb{Z}}\sum_{n\in\mathbb{Z}} \hat{S}_X^n\left(\frac{\phi-k}{T_0}\right) e^{2\pi i\frac{nr+(m-r)(\phi-k)}{M}}. \end{aligned} \tag{A.28}$$

Then the limit of D_M in M exists and the distortion-rate function of $X(\cdot)$ is given by

$$D(R) = \lim_{M\to\infty} D_M\left(\theta_M(R)\right). \tag{A.29}$$

Proof Given a Gaussian CSP $X(\cdot)$ with period $T_0 > 0$, we define the discrete-time process $\bar{X}[\cdot]$ obtained by uniformly sampling $X(\cdot)$ at intervals T_0/M, i.e.

$$\bar{X}[n] = X(nT_0/M), \quad n\in\mathbb{Z}. \tag{A.30}$$

The autocorrelation function of $\bar{X}[\cdot]$ satisfies

$$\begin{aligned} R_{\bar{X}}[n+M,k] &= \mathbb{E}\left[\bar{X}[n+M+k]\bar{X}[n+M]\right] \\ &= \mathbb{E}\left[X(nT_0/M+T_0+kT_0/M)X(nT_0/M+T_0)\right] \\ &= R_X(nT_0/M+T_0, kT_0/M+T_0) \\ &= R_X(nT_0/M, kT_0/M) \\ &= R_{\bar{X}}[n,k], \end{aligned}$$

which means that $\bar{X}[\cdot]$ is a discrete-time Gaussian CSP with period M. The TPSD of $\bar{X}[\cdot]$ is given by

$$S_{\bar{X}}^m(e^{2\pi i\phi}) = \frac{M}{T_0}\sum_{k\in\mathbb{Z}} S_X^{mT_0/M}\left(\frac{M}{T_0}(\phi-k)\right).$$

This means that the PSD of the m^{th} PC of $\bar{X}[\cdot]$ is

$$\begin{aligned} S_{\bar{X}}^m\left(e^{2\pi i\phi}\right) &= \frac{1}{M}\sum_{n=0}^{M-1} S_{\bar{X}}^m\left(e^{2\pi i\frac{\phi-n}{M}}\right) \\ &= \frac{1}{T_0}\sum_{n=0}^{M-1}\sum_{k\in\mathbb{Z}} S_X^{mT_0/M}\left(\frac{\phi-Mk-n}{T_0}\right) \\ &= \frac{1}{T_0}\sum_{l\in\mathbb{Z}} S_X^{mT_0/M}\left(\frac{\phi-l}{T_0}\right). \end{aligned}$$

By applying Theorem A.1 to $\bar{X}[\cdot]$, we obtain an expression for the DRF of $\bar{X}[\cdot]$ as a function of M:

$$D_M(\theta_M) = \frac{1}{M}\sum_{m=1}^{M}\int_{-\frac{1}{2}}^{\frac{1}{2}} \min\left\{\lambda_m\left(e^{2\pi i\phi}\right),\theta_M\right\} d\phi \tag{A.31a}$$

$$\bar{R}(\theta_M) = \frac{1}{2M}\sum_{m=1}^{M}\int_{-\frac{1}{2}}^{\frac{1}{2}} \log^+\left[\lambda_m\left(e^{2\pi i\phi}\right)/\theta_M\right] d\phi, \tag{A.31b}$$

where $\lambda_1\left(e^{2\pi i\phi}\right) \leq \ldots \leq \lambda_M\left(e^{2\pi i\phi}\right)$ are the eigenvalues of the matrix with $(m+1,r+1)^{\text{th}}$ entry

$$\begin{aligned} S_{\bar{X}^m\bar{X}^r}\left(e^{2\pi i\phi}\right) &= \frac{1}{M}\sum_{n=0}^{M-1} S_{\bar{X}}^r\left(e^{2\pi i\frac{\phi-n}{M}}\right)e^{2\pi i(m-r)\frac{\phi-n}{M}} \\ &= \frac{1}{T_0}\sum_{n=0}^{M-1}\sum_{k\in\mathbb{Z}} S_X^{rT_0/M}\left(\frac{\phi-n-kM}{T_0}\right)e^{2\pi i(m-r)\frac{\phi-n}{M}}, \\ &= \frac{1}{T_0}\sum_{l\in\mathbb{Z}} S_X^{rT_0/M}\left(\frac{\phi-l}{T_0}\right)e^{2\pi i(m-r)\frac{\phi-l}{M}}. \end{aligned} \tag{A.32}$$

In order to express the code-rate in bits per time unit, we multiply the number of bits per sample $\bar{R}$ by the sampling rate M/T_0. This shows that the DRF of $\bar{X}[\cdot]$, as measured in bits per unit time R, is given by (A.27).

In order to complete the proof we rely on the following lemma:

Lemma A.3 *Let $X(\cdot)$ be as in Theorem A.2 and let $\bar{X}[\cdot]$ be its uniformly sampled version at rate M/T_0 as in (A.30). Denote the DRF at rate R bits per time unit of the two processes by $D(R)$ and $\bar{D}(R)$, respectively. Then*

$$\lim_{M\to\infty}\bar{D}(R) = D(R).$$

The proof of Lemma A.3 is given as follows: let $K(t,s) = R_X(s,t-s)$ and $\bar{K}[n,k] = R_{\bar{X}}[n,k-n]$. For $M\in\mathbb{N}$, define

$$\widetilde{K}(t,s) = K\left(\lfloor tM/T_0\rfloor T_0/M, \lfloor sM/T_0\rfloor T_0/M\right).$$

For any fixed $T>0$, the kernel $\widetilde{K}(t,s)$ defines a Hermitian positive compact operator [77] on the space of square integrable functions over $[-T,T]$. The eigenvalues of this operator are given by the Fredholm integral equation

$$\tilde{\lambda}_l\tilde{f}_l(t) = \frac{1}{2T}\int_{-T}^{T}\widetilde{K}(t,s)\tilde{f}_l(s)ds, \quad -T\leq t\leq T, \tag{A.33}$$

where it can be shown that there are at most MT/T_0 non-zero eigenvalues $\{\tilde{\lambda}_l\}$ that satisfy (A.33). We define the function $\tilde{D}_T(R)$ by the following parametric expression:

$$\begin{aligned} \tilde{D}_T(\theta) &= \sum_{l=1}^{\infty} \min\left\{\tilde{\lambda}_l, \theta\right\} \\ R(\theta) &= \frac{1}{2}\sum_{l=1}^{\infty} \log^+\left(\frac{\tilde{\lambda}_l}{\theta}\right) \end{aligned} \tag{A.34}$$

(the eigenvalues in (A.34) are implicitly depend on T). Note that

$$\sum_{l=1}^{\infty} \tilde{\lambda}_l = \frac{1}{2T}\int_{-T}^{T} \widetilde{K}(t,t)dt = \frac{1}{2T}\sum_{n=-N}^{N} K(nT_0/M, nT_0/M), \tag{A.35}$$

where $N = MT/T_0$. Expression (A.35) converges to

$$\frac{1}{2T}\int_{-T}^{T} K(t,t)dt \leq \sigma_X^2$$

as M goes to infinity due to our assumption that $R(t,\tau)$ is Riemann integrable and therefore so is $K(t,s)$. Since we are interested in the asymptotic of large M, we can assume that (A.35) is bounded. This implies that $\tilde{D}_T(R)$ is bounded.

In order to claim that the eigenvalues $\{\tilde{\lambda}_l\}$ approximate the eigenvalues $\{\lambda_l\}$, we use the following lemma:

Lemma A.4 *Let $\{\lambda_l\}$ and $\{\tilde{\lambda}_l\}$ be the eigenvalues in the Fredholm integral equation of $K(t,s)$ and $\widetilde{K}(t,s)$, respectively. Assume that these eigenvalues are numbered in a descending order. Then*

$$\left|\lambda_l - \tilde{\lambda}_l\right| \leq 4CT_0/M, \quad l = 1,2,\ldots. \tag{A.36}$$

Proof of Lemma A.4

Approximations of the kind (A.36) can be obtained by Weyl's inequalities for singular values of operators defined by self-adjoint kernels [78]. In our case it suffices to use the following result [79, Cor. 1"]:

$$\left|\lambda_l - \tilde{\lambda}_l\right| \leq 2\sup_{t,s\in[-T,T]}\left|K(t,s) - \widetilde{K}(t,s)\right|, \quad l = 1,2,\ldots. \tag{A.37}$$

The assumption that $R_X(t,\tau)$ is Lipschitz continuous in τ implies that there exists a constant $C > 0$ such that for any $t_1, t_2, s \in \mathbb{R}$,

$$|K(t_1,s) - K(t_2,s)| = |R_X(s,t_1-s) - R_X(s,t_2-s)| \leq C|t_1 - t_2|.$$

We therefore conclude that $K_X(t,s)$ is Lipschitz continuous in both of its arguments from symmetry. Lipschitz continuity of $K(t,s)$ implies

$$\begin{aligned}
&|K(t_1,s_1)-K(t_2,s_2)| \\
&\leq |K(t_1,s_1)-K(t_1,s_2)| + \left|K(t_1,s_2)-\widetilde{K}(t_2,s_2)\right| \\
&\leq C|s_1-s_2| + C|t_1-t_2|.
\end{aligned}$$

As a result, (A.37) leads to

$$\begin{aligned}
\left|\lambda_l-\tilde{\lambda}_l\right| &\leq 2\sup_{t,s}\left|K(t,s)-\widetilde{K}(t,s)\right| \\
&= 2\sup_{t,s\in[-T,T]}\left|K(t,s)-K\left(\lfloor tM/T_0\rfloor T_0/M, \lfloor sM/T_0\rfloor T_0/M\right)\right| \\
&\leq 2C\left(|t-\lfloor tM/T_0\rfloor T_0/M| + |t-\lfloor sM/T_0\rfloor T_0/M|\right) \\
&\leq 4CT_0/M,
\end{aligned}$$

which proves Lemma A.4.

The significance of Lemma A.4 is in the fact that the eigenvalues of the kernel $K(t,s)$ used in the expression for the DRF of $X(\cdot)$ can be approximated by the eigenvalues of $\widetilde{K}(t,s)$, where the error in each of these approximations converge, uniformly in T, to zero as M increases. Since only a finite number of eigenvalues participate in (A.19) and since both $D_T(R)$ and $\tilde{D}_T(R)$ are bounded continuous functions of their eigenvalues, we conclude that $\tilde{D}_T(R)$ converges to $D_T(R)$ uniformly in T.
Now let $\varepsilon > 0$ and fix M_0 large enough such that for all $M > M_0$ and for all T,

$$\left|D_T(R)-\tilde{D}_T(R)\right| \leq \varepsilon. \tag{A.38}$$

Recall that in addition to (A.21), the DRF of $\bar{X}[\cdot]$, denoted here as $\bar{D}(\bar{R})$, can also be obtained as the limit in N of the expression

$$\begin{aligned}
\bar{D}_N(\theta) &= \sum_{l=1}^{\infty}\min\left\{\bar{\lambda}_l,\theta\right\} \\
\bar{R}(\theta) &= \frac{1}{2}\sum_{l=1}^{\infty}\log^+\left(\bar{\lambda}_l/\theta\right),
\end{aligned}$$

where $\bar{\lambda}_1,\bar{\lambda}_2,\dots$ are the eigenvalues in the KL expansion of $\bar{X}$ over $n=-N,\dots,N$:

$$\bar{\lambda}_l f_l[n] = \frac{1}{2N+1}\sum_{k=-N}^{N}K_{\bar{X}}[n,k]f_l[k], \quad l=1,\dots,N, \tag{A.39}$$

(there are actually at most $2N+1$ distinct non-zero eigenvalues that satisfies (A.39)). Letting $T_N = T_0 M/N$ and $\tilde{f}_l(t) = f_l\left(\lfloor t/T_0 \rfloor M\right)$ (A.39) can also be written as

$$\bar{\lambda}_l f_l[n] = \int_{-T_N}^{T_N} \widetilde{K}_X(nT_0/M, s) f_l\left[\lfloor s/T_0 \rfloor M\right] ds, \quad l = 1,2,\ldots,$$
$$\bar{\lambda}_l \tilde{f}_l(t) = \int_{-T_n}^{T_N} \widetilde{K}(t,s)\tilde{f}_l(s) ds, \quad -T_N < t < T_N.$$

From the uniqueness of the KL expansion, we obtain that for any N, the eigenvalues of $\widetilde{K}(t,s)$ over $T_N = T_0 M/N$ are given by the eigenvalues of $\bar{K}[n,k]$ over $-N,\ldots,N$. We conclude that

$$\bar{D}_N(\bar{R}) = \tilde{D}_{T_N}(R), \tag{A.40}$$

where $R = \bar{R}T_0/M$. Now take N large enough such that

$$|\bar{D}_N(R) - \bar{D}(R)| < \varepsilon,$$

and

$$|D_{T_N}(R) - D(R)| < \varepsilon.$$

For all $M \geq M_0$ we have

$$\begin{aligned} |D(R) - \bar{D}(R)| &= \big|D(R) - D_{T_N}(R) + D_{T_N}(R) + \tilde{D}_{T_N}(R) \\ &\quad - \tilde{D}_{T_N}(R) + \bar{D}_N(R) - \bar{D}_N(R) - \bar{D}(R)| \\ &\leq |D(R) - D_{T_N}(R)| \quad (A.41) \\ &\quad + \big|D_{T_N}(R) - \tilde{D}_{T_N}(R)\big| \quad (A.42) \\ &\quad + \big|\tilde{D}_{T_N}(R) - \bar{D}_N(R)\big| \quad (A.43) \\ &\quad + |\bar{D}_N(R) - \bar{D}(R)| \leq 3\varepsilon, \quad (A.44) \end{aligned}$$

where the last transition is because: (A.41) and (A.44) are smaller than ε by the choice of N, (A.42) is smaller than ε from (A.38). and (A.43) equals zero from (A.40). □

DRF of PAM-modulated signals

Proposition A.5 *Let $X_{\mathsf{PAM}}(\cdot)$ be defined by*

$$X_{\mathsf{PAM}}(t) = \sum_{n\in\mathbb{Z}} U(nT_0)p(t - nT_0), \quad t \in \mathbb{R}, \tag{A.45}$$

where $U(\cdot)$ is a Gaussian stationary process with (Although we only use the value of $U(t)$ at $t \in \mathbb{Z}T_0$, it is convenient to treat $U(\cdot)$ as continuous-time source so that the expressions emerging have only continuous-time spectrum) PSD $S_U(f)$ and $p(t)$ is an analog deterministic signal with $\int_{-\infty}^{\infty} |p(t)|^2 dt < \infty$ and Fourier transform $P(f)$. Assume moreover, that the covariance function $R_{X_{\mathsf{PAM}}}(t,\tau)$ of $X_{\mathsf{PAM}}(\cdot)$ is Lipschitz continuous in its second argument. The distortion-rate function of $X_{\mathsf{PAM}}(\cdot)$ is given by

$$D(\theta) = \frac{1}{T_0} \int_{-\frac{1}{2T_0}}^{\frac{1}{2T_0}} \min\left\{\widetilde{S}(f), \theta\right\} df \tag{A.46a}$$

$$R(\theta) = \frac{1}{2} \int_{-\frac{1}{2T_0}}^{\frac{1}{2T_0}} \log^+ \left[\widetilde{S}(f)/\theta\right] df, \tag{A.46b}$$

where

$$\widetilde{S}(f) \triangleq \sum_{k\in\mathbb{Z}} |P(f-k/T_0)|^2 S_U(f-k/T_0). \tag{A.47}$$

Proof The entries of the matrix $\mathbf{S}\left(e^{2\pi i\phi}\right)$ in Theorem A.2 are obtained by using the CPSD of the PAM process (A.15) in (A.28). For all $M \in \mathbb{N}$, this leads to

$$\mathbf{S}_{m+1,r+1}\left(e^{2\pi i\phi}\right) =$$

$$\frac{1}{T_0^2} \sum_{k\in\mathbb{Z}} \left[P\left(\frac{\phi-k}{T_0}\right) S_U\left(\frac{\phi-k}{T_0}\right) e^{2\pi i(\phi-k)\frac{m-r}{M}} \sum_{n\in\mathbb{Z}} P^*\left(\frac{\phi-n-k}{T_0}\right) e^{2\pi i\frac{nr}{M}} \right] \tag{A.48}$$

$$= \frac{1}{T_0^2} \sum_{k\in\mathbb{Z}} P\left(\frac{\phi-k}{T_0}\right) S_U\left(\frac{\phi-k}{T_0}\right) e^{2\pi i(\phi-k)\frac{m}{M}} \sum_{l\in\mathbb{Z}} P^*\left(\frac{\phi-l}{T_0}\right) e^{-2\pi i(\phi-l)\frac{r}{M}}. \tag{A.49}$$

The expression (A.49) consists of the product of a term depending only on m and a term depending only on r. We conclude that the matrix $\mathbf{S}\left(e^{2\pi i\phi}\right)$ can be written as the outer product of two M dimensional vectors, and thus it is of rank one. The single non-zero eigenvalue $\lambda_M\left(e^{2\pi i\phi}\right)$ of $\mathbf{S}\left(e^{2\pi i\phi}\right)$ is given by the trace of the matrix, which, by the orthogonality of the functions $e^{2\pi i\frac{nr}{M}}$ in (A.48), is evaluated as

$$\lambda_M\left(e^{2\pi i\phi}\right) = \frac{M}{T_0^2} \sum_{k\in\mathbb{Z}} \left| P\left(\frac{\phi-k}{T_0}\right)\right|^2 S_U\left(\frac{\phi-k}{T_0}\right). \tag{A.50}$$

We now use (A.50) in (A.27). In order to obtain (A.46), we change the integration variable from ϕ to $f = \phi/T_0$ and the water-level parameter θ to $T_0\theta/M$. Note that the final expression is independent of M, so the limit in (A.29) is already given by this expression. □

Bibliography

[1] E. T. Wittaker, "On the functions which are represented by the expansions of the interpolatory theory," in *Proc. Roy. Soc. Edinburgh*, vol. 35, 1915, pp. 181–194.

[2] C. E. Shannon, "Communication in the presence of noise," *IRE Transactions on Information Theory*, vol. 37, pp. 10–21, 1949.

[3] H. Landau, "Sampling, data transmission, and the Nyquist rate," *Proceedings of the IEEE*, vol. 55, no. 10, pp. 1701–1706, Oct 1967.

[4] C. E. Shannon, "A mathematical theory of communication," *Bell System Tech. J.*, vol. 27, pp. 379–423, 623–656, 1948.

[5] ——, "Coding theorems for a discrete source with a fidelity criterion," *IRE Nat. Conv. Rec*, vol. 4, no. 142-163, p. 1, 1959.

[6] T. Berger, *Rate-distortion theory: A mathematical basis for data compression*. Englewood Cliffs, NJ: Prentice-Hall, 1971.

[7] J. Yick, B. Mukherjee, and D. Ghosal, "Wireless sensor network survey," *Computer Networks*, vol. 52, no. 12, pp. 2292 – 2330, 2008.

[8] B. P. Ginsburg, "Energy-efficient analog-to-digital conversion for ultra-wideband radio," Ph.D. dissertation, Massachusetts Institute of Technology, 2007.

[9] R. Walden, "Analog-to-digital converter survey and analysis," *IEEE Journal on Selected Areas in Communications*, vol. 17, no. 4, pp. 539–550, Apr 1999.

[10] T. S. Rappaport, R. W. Heath Jr, R. C. Daniels, and J. N. Murdock, *Millimeter wave wireless communications*. Pearson Education, 2014.

[11] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature*, vol. 521, no. 7553, pp. 436–444, 2015.

[12] I. Karatzas and S. E. Shreve, *Methods of mathematical finance*. Springer Science & Business Media, 1998, vol. 39.

[13] M. W. Mahoney, "Randomized algorithms for matrices and data," *Foundations and Trends in Machine Learning*, vol. 3, no. 2, pp. 123–224, 2011.

[14] Y. C. Eldar, *Sampling Theory: Beyond Bandlimited Systems*. Cambridge University Press, 2015.

[15] R. Dobrushin and B. Tsybakov, "Information transmission with additional noise," *IRE Transactions on Information Theory*, vol. 8, no. 5, pp. 293–304, 1962.

[16] J. Wolf and J. Ziv, "Transmission of noisy information to a noisy receiver with minimum distortion," *IEEE Transactions on Information Theory*, vol. 16, no. 4, pp. 406–411, 1970.

[17] A. Kipnis, A. J. Goldsmith, Y. C. Eldar, and T. Weissman, "Distortion rate function of sub-Nyquist sampled Gaussian sources," *IEEE Transactions on Information Theory*, vol. 62, no. 1, pp. 401–429, Jan 2016.

[18] A. Kipnis, A. J. Goldsmith, T. Weissman, and Y. C. Eldar, "Distortion rate function of sub-nyquist sampled gaussian sources corrupted by noise," in *51st Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, Oct 2013, pp. 901–908.

[19] A. Kipnis, A. J. Goldsmith, and Y. C. Eldar, "Rate distortion function of cyclostationary Gaussian processes," *IEEE Transactions on Information Theory*, to appear.

[20] A. Kipnis and A. J. Goldsmith, "Distortion rate function of cyclostationary Gaussian processes," in *2014 IEEE International Symposium on Information Theory*, June 2014, pp. 2834–2838.

[21] A. Kipnis, Y. C. Eldar, and A. J. Goldsmith, "Fundamental distortion limits of analog-to-digital compression," 2016.

[22] A. Kipnis, A. J. Goldsmith, and Y. C. Eldar, "Gaussian distortion-rate function under sub-Nyquist nonuniform sampling," in *52nd Annual Allerton Conference on Communication, Control, and Computing (Allerton)*. IEEE, 2014, pp. 874–880.

[23] ——, "Sub-Nyquist sampling achieves optimal rate-distortion," in *Information Theory Workshop (ITW), 2015 IEEE*, April 2015, pp. 1–5.

[24] ——, "Optimal trade-off between sampling rate and quantization precision in sigma-delta a/d conversion," in *2015 International Conference on Sampling Theory and Applications (SampTA)*, May 2015, pp. 627–631.

[25] ——, "The distortion-rate function of sampled Wiener processes," 2016, to appear. [Online]. Available: http://arxiv.org/abs/1608.04679

[26] A. Kipnis, Y. C. Eldar, and A. J. Goldsmith, "Information rates of sampled Wiener processes," in *2016 IEEE International Symposium on Information Theory*, July 2016, pp. 740–744.

[27] E. T. Whittaker, "On the functions which are represented by the expansions of the interpolation-theory," *Proceedings of the Royal Society of Edinburgh*, vol. 35, pp. 181–194, 1915.

[28] V. A. Kotelnikov, "The carrying capacity of "ether" and wire in electrical communications," in *Materials for the 1st All-Union Congress on Problems of Technical Reconstruction of Communications and Development of the Low-Current Industry [in Russian]*, 1933.

[29] C. Shannon, "Communication in the presence of noise," *Proceedings of the IEEE*, vol. 86, no. 2, pp. 447–457, Feb 1998.

[30] H. Landau, "Necessary density conditions for sampling and interpolation of certain entire functions," *Acta Mathematica*, vol. 117, no. 1, pp. 37–52, 1967.

[31] R. G. Vaughan, N. L. Scott, and D. R. White, "The theory of bandpass sampling," *IEEE Transactions on Signal Processing*, vol. 39, no. 9, pp. 1973–1984, 1991.

[32] R. Venkataramani and Y. Bresler, "Optimal sub-nyquist nonuniform sampling and reconstruction for multiband signals," *IEEE Transactions on Signal Processing*, vol. 49, no. 10, pp. 2301–2313, Oct 2001.

[33] F. Marvasti, *Nonuniform sampling: theory and practice*. Springer Science & Business Media, 2012.

[34] Y. C. Eldar and T. Michaeli, "Beyond bandlimited sampling," *IEEE Signal Processing Magazine*, vol. 26, no. 3, pp. 48–68, 2009.

[35] D. L. Donoho, "Compressed sensing," *IEEE Transactions on Information Theory*, vol. 52, no. 4, pp. 1289–1306, 2006.

[36] E. Candes, J. Romberg, and T. Tao, "Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information," *IEEE Transactions on Information Theory*, vol. 52, no. 2, pp. 489–509, Feb 2006.

[37] A. Kipnis, G. Reeves, Y. C. Eldar, and A. J. Goldsmith, "Fundamental limits of compressed sensing under optimal quantization," in *2017 IEEE International Symposium on Information Theory*, June 2017.

[38] R. G. Gallager, *Stochastic processes: theory for applications*. Cambridge University Press, 2013.

[39] H. Dym and H. McKean, "Gaussian processes, function theory, and the inverse spectral problem," *Bull. Amer. Math. Soc. 84 (1978), 260-262*, pp. 0002–9904, 1978.

[40] D. Alpay and A. Kipnis, "Wiener chaos approach to optimal prediction," *Numerical Functional Analysis and Optimization*, vol. 36, no. 10, pp. 1286–1306, 2015. [Online]. Available: http://dx.doi.org/10.1080/01630563.2015.1065273

[41] I. Gelfand and G. Shilov, *Generalized functions. Volume 2*. Academic Press, 1968.

[42] R. Melrose, "Introduction to microlocal analysis," *Lecture notes from courses taught at MIT*, 2007.

[43] Y. Lyubarskii and K. Seip, "Weighted paley-wiener spaces," *Journal of the American Mathematical Society*, vol. 15, no. 4, pp. 979–1006, 2002.

[44] M. Matthews, "On the linear minimum-mean-squared-error estimation of an undersampled wide-sense stationary random process," *IEEE Transactions on Signal Processing*, vol. 48, no. 1, pp. 272–275, 2000.

[45] Y. Chen, Y. C. Eldar, and A. J. Goldsmith, "Shannon meets Nyquist: Capacity of sampled Gaussian channels," *IEEE Transactions on Information Theory*, vol. 59, no. 8, pp. 4889–4914, 2013.

[46] D. Chan and R. Donaldson, "Optimum pre-and postfiltering of sampled signals with application to pulse modulation and data compression systems," *IEEE Transactions on Communication Technology*, vol. 19, no. 2, pp. 141–157, April 1971.

[47] T. Michaeli and Y. C. Eldar, "High-rate interpolation of random signals from nonideal samples," *IEEE Transactions on Signal Processing*, vol. 57, no. 3, pp. 977–992, 2009.

[48] T. Kailath, "A view of three decades of linear filtering theory," *IEEE Transactions on Information Theory*, vol. 20, no. 2, pp. 146–181, 1974.

[49] W. A. Gardner, A. Napolitano, and L. Paura, "Cyclostationarity: Half a century of research," *Signal Process.*, vol. 86, no. 4, pp. 639–697, Apr. 2006.

[50] A. Peled and B. Liu, "Digital signal processing: theory, design, and implementation," *New York, John Wiley and Sons, Inc., 1976. 319 p.*, 1976.

[51] M. Bellanger, G. Bonnerot, and M. Coudreuse, "Digital filtering by polyphase network: Application to sample-rate alteration and filter banks," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. 24, no. 2, pp. 109–114, Apr 1976.

[52] W. Gardner, *Cyclostationarity in communications and signal processing*, ser. Electrical engineering, communications and signal processing. IEEE Press, 1994.

[53] T. Kailath, A. H. Sayed, and B. Hassibi, *Linear estimation*. Prentice Hall Upper Saddle River, NJ, 2000, vol. 1.

[54] S. P. Lloyd, "A sampling theorem for stationary (wide sense) stochastic processes," *Transactions of the American Mathematical Society*, vol. 92, no. 1, pp. pp. 1–12, 1959.

[55] A. Balakrishnan, "A note on the sampling principle for continuous signals," *IRE Transactions on Information Theory*, vol. 3, no. 2, pp. 143–146, June 1957.

[56] P. Vaidyanathan, "Multirate digital filters, filter banks, polyphase networks, and applications: a tutorial," *Proceedings of the IEEE*, vol. 78, no. 1, pp. 56–93, 1990.

[57] R. Young, *An Introduction to Nonharmonic Fourier Series*, ser. An Introduction to Nonharmonic Fourier Series. Academic Press, 2001.

[58] I. Gohberg, S. Goldberg, and M. Kaashoek, *Basic Classes of Linear Operators*. Birkhäuser Basel, 2003.

[59] H. L. Royden and P. Fitzpatrick, *Real analysis*. Macmillan New York, 1988, vol. 198, no. 8.

[60] A. Perez, “Extensions of shannon-mcmillan’s limit theorem to more general stochastic processes,” in *Trans. Third Prague Conference on Information Theory, Statistical Decision Functions and Random Processes*, 1964, pp. 545–574.

[61] T. Berger, “Rate distortion theory for sources with abstract alphabets and memory,” *Information and Control*, vol. 13, no. 3, pp. 254 – 273, 1968.

[62] A. Kolmogorov, “On the shannon theory of information transmission in the case of continuous signals,” *IRE Transactions on Information Theory*, vol. 2, no. 4, pp. 102–108, December 1956.

[63] T. Berger and J. Gibson, “Lossy source coding,” *IEEE Transactions on Information Theory*, vol. 44, no. 6, pp. 2693–2723, 1998.

[64] R. Gray, D. Neuhoff, and J. Omura, “Process definitions of distortion-rate functions and source coding theorems,” *IEEE Transactions on Information Theory*, vol. 21, no. 5, pp. 524–532, Sep 1975.

[65] R. Gray, *Probability, Random Processes, and Ergodic Properties*. Springer, 2009.

[66] ——, *Entropy and information theory*. Springer-Verlag, 1990.

[67] V. Kafedziski, “Rate distortion of stationary and nonstationary vector Gaussian sources,” in *2005 IEEE/SP 13th Workshop on Statistical Signal Processing*, 2005, pp. 1054–1059.

[68] Y.-P. Lin and P. Vaidyanathan, “Periodically nonuniform sampling of bandpass signals,” *IEEE Transactions on Circuits and Systems—Part II: Analog and Digital Signal Processing*, vol. 45, no. 3, pp. 340–351, Mar 1998.

[69] I. Karatzas and S. Shreve, *Brownian Motion and Stochastic Calculus*, ser. Graduate Texts in Mathematics. Springer New York, 1991.

[70] D. L. Donoho, A. Maleki, and A. Montanari, “The noise-sensitivity phase transition in compressed sensing,” *IEEE Transactions on Information Theory*, vol. 57, no. 10, pp. 6920–6941, 2011.

[71] Y. Wu and S. Verdu, “Optimal phase transitions in compressed sensing,” *IEEE Transactions on Information Theory*, vol. 58, no. 10, pp. 6241–6263, Oct 2012.

[72] A. Kipnis, S. Rini, and A. J. Goldsmith, “Compress-and-estimate in multiterminal source coding,” 2017, to appear. [Online]. Available: https://arxiv.org/abs/1602.02201

[73] T. Berger, Z. Zhang, and H. Viswanathan, "The CEO problem [multiterminal source coding]," *IEEE Transactions on Information Theory*, vol. 42, no. 3, pp. 887–902, 1996.

[74] W. Gardner, "Spectral correlation of modulated signals: Part I–analog modulation," *IEEE Transactions on Communications*, vol. 35, no. 6, pp. 584–594, 1987.

[75] R. Gallager, *Information theory and reliable communication*. Wiley, 1968.

[76] T. Berger, "Information rates of Wiener processes," *IEEE Transactions on Information Theory*, vol. 16, no. 2, pp. 134–139, 1970.

[77] H. König, *Eigenvalue Distribution of Compact Operators*, ser. Operator Theory: Advances and Applications. Birkhäuser Basel, 2013.

[78] H. Weyl, "Das asymptotische verteilungsgesetz der eigenwerte linearer partieller differentialgleichungen (mit einer anwendung auf die theorie der hohlraumstrahlung)," *Mathematische Annalen*, vol. 71, no. 4, pp. 441–479, 1912.

[79] H. Wielandt, "Error bounds for eigenvalues of symmetric integral equations," in *Proc. Sympos. Appl. Math*, vol. 6, 1956, pp. 261–282.