





Blind Modulo Analog-to-Digital Conversion

Amir Weiss , Member, IEEE, Everest Huang , Member, IEEE, Or Ordentlich , Member, IEEE, and Gregory W. Wornell , Fellow, IEEE

Abstract—In a growing number of applications, there is a need to digitize signals whose spectral characteristics are challenging for traditional analog-to-digital converters (ADCs). Examples, among others, include systems where the ADC must acquire at once a very wide but sparsely and dynamically occupied bandwidth supporting diverse services, as well as systems where the signal of interest is subject to strong narrowband co-channel interference. In such scenarios, the resolution requirements can be prohibitively high. As an alternative, the recently proposed *modulo-ADC* architecture can in principle require dramatically fewer bits in the conversion to obtain the target fidelity, but requires that information about the spectrum be known and explicitly taken into account by the analog and digital processing in the converter, which is frequently impractical. To address this limitation, we develop a *blind* version of the architecture that requires no such knowledge in the converter, without sacrificing performance. In particular, it features an automatic modulo-level adjustment and a fully adaptive modulo unwrapping mechanism, allowing it to asymptotically match the characteristics of the unknown input signal. In addition to detailed analysis, simulations demonstrate the attractive performance characteristics in representative settings.

Index Terms—Data conversion, automatic gain control, blind signal processing, adaptive filtering, least-mean-squares.

I. INTRODUCTION

THE available spectrum for a communication system is increasingly congested and varies widely by location and time. One strategy for operation in these dynamic conditions is to scan the spectrum to find unoccupied bandwidth within which to transmit. Historically, this has been a difficult task since the fraction of all the potentially usable bandwidth that can be scanned simultaneously is limited by the bandwidth of the receiver front-end. In traditional system architectures, a narrow analog filter matched to the desired communication band rejects out-of-band transmissions from overwhelming the analog-to-digital converter (ADC) prior to any digital processing. This results in either a fixed-frequency system, or an expensive




Fig. 1. A schematic block diagram illustration of the mod-ADC.

frequency-agile analog front-end for both the transmitter and receiver. For narrowband systems, the fraction of potentially available bandwidth that can be monitored at any instant can be small, which slows down the response to potentially rapidly changing channel conditions. The system is further complicated by the need to coordinate between communication nodes what frequencies are being used when.

The emergence of high speed ADCs with multiple GHz of bandwidth enables affordable systems to be built, that can simultaneously scan large regions of the spectrum for unutilized bandwidth to transmit in [1]. The congested nature of the spectrum, however, requires robust front-end processing to accommodate the large dynamic range required from multiple possibly strong interfering sources [2]. A wide or changing frequency allocation will by necessity allow these outside sources to be sampled as well. While digital processing can in principle remove the effect of the undesired signals, the ADC must still be able to faithfully sample the entire bandwidth, containing all signals, prior to any subsequent digital manipulation. Thus, despite the strong structure and high predictability of the sampled signal, a traditional ADC requires a high number of bits per-second in order to allow for high-quality reconstruction.

One possible approach for addressing the inefficiency described above, is using *modulo ADCs* [3] instead of a traditional ADC. A modulo ADC first folds each sample of the input process modulo Δ , where Δ is a design parameter, and only then quantizes the result using a traditional uniform quantizer. See Fig. 1 for a schematic description. The modulo operation limits the dynamic range of the signal to be quantized, which in turns results in a quantization error whose magnitude is proportional to Δ , rather than to the dynamic range of the original, unfolded signal. In [3] it is shown that the obtained signal can be reliably unfolded, provided that Δ is appropriately chosen proportionally to the standard deviation of the *prediction error* in predicting the (quantized) input from its past. Thus, when using a modulo ADC for digitizing a highly predictable process, one can attain high resolution using far fewer bits than for a white process. Simple recovery algorithms based on linear prediction are also given in [3].

A major caveat of the modulo ADC framework developed in [3] is that it assumes knowledge of the second-order statistics

Manuscript received 10 December 2021; revised 15 July 2022; accepted 18 July 2022. Date of publication 11 August 2022; date of current version 27 September 2022. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Hassan Mansour. This work was supported in part by ISF under Grant 1791/17 and in part by NSF under Grants CCF-1717610 and CCF-1816209. (Corresponding author: Amir Weiss.)

Amir Weiss and Gregory W. Wornell are with the Electrical Engineering and Computer Science Department, Massachusetts Institute of Technology, Cambridge 02139 USA (e-mail: amirweiss15@gmail.com; gww@mit.edu).

Everest Huang is with the Advanced Satcom Systems and Operations Group at MIT Lincoln Laboratory, Lexington, MA 02421 USA (e-mail: everest@ll.mit.edu).

Or Ordentlich is with the School of Computer Science and Engineering, Hebrew University of Jerusalem - Edmond J. Safra Campus, Jerusalem 91904, Israel (e-mail: or.ordentlich@mail.huji.ac.il).

Digital Object Identifier 10.1109/TSP.2022.3198184

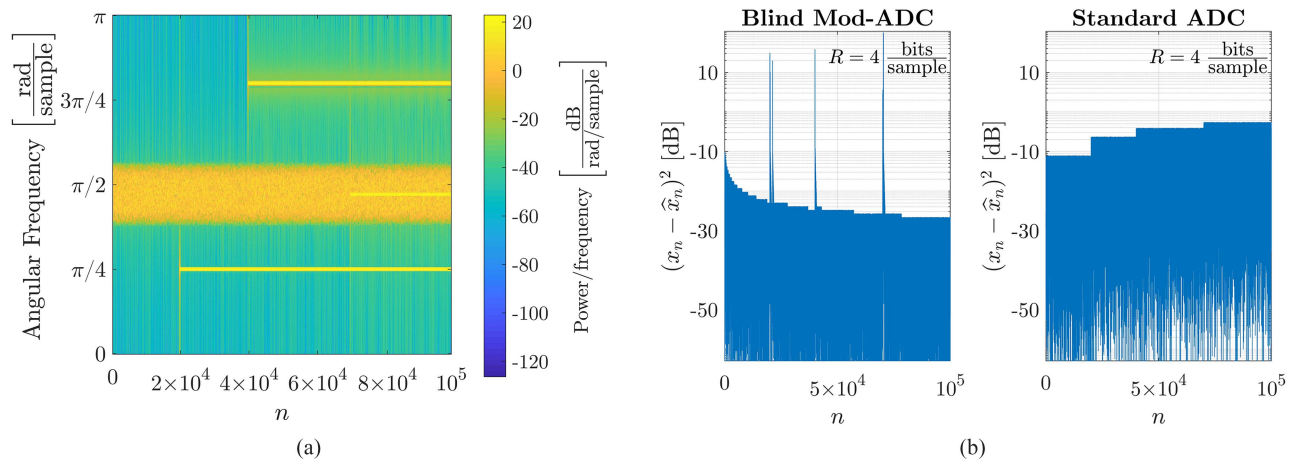


Fig. 2. An illustrative scenario of interest. (a) Spectrogram of the ADC input: a bandlimited signal of interest in the presence of narrowband interferences (b) The instantaneous squared errors of the outputs of a standard ADC and our proposed blind mod-ADC. Evidently, the blind mod-ADC achieves a significantly lower MSE relative to a standard uniform ADC with a perfect (“oracle”) AGC mechanism.

(SOSs) of the input process. Such knowledge is crucial for optimizing the modulo ADC parameter Δ , as well as for optimizing the coefficients of the prediction filter used in the unwrapping process. This is a significant impediment to the implementation of modulo ADCs in practice, as commercial ADCs must be robust to the characteristics of the process to be digitized. In particular, traditional ADCs employ an automatic gain control (AGC) mechanism [4], [5] for adapting the quantizers dynamic range to that of the input signal.

In this work we develop a *blind* mechanism for modulo ADCs, which adapts the effective modulo size (analogously to an AGC mechanism in a standard ADC, see Subsection I-C), as well as the coefficients of the unwrapping algorithm prediction filter, to the unknown statistics of the input signal, resulting in a robust ADC architecture. For a stationary input process, the performance of the developed architecture converges to that of the “informed” architecture in [3]. The developed architecture is also dynamic, and quickly adapts to changes in the characteristics of the input signal.

A. A Motivating Example

To illustrate the challenges in digitizing communication signals whose locations within the frequency band are unknown, we consider the signal whose spectrogram is depicted in Fig. 2(a). This spectrogram corresponds to a sampled binary phase-shift keying (BPSK) signal together with three narrowband interfering signals (specifically, pure tones), with each interferer initiated at a different time. If the carrier frequency of the BPSK signal were known in advance, one could first down-convert it and use an analog low-pass filter to cancel out all interference outside the frequency band it occupies, and only then sample at the corresponding Nyquist rate. The discrete-time signal resulting from this process is essentially “white” (i.e., temporally uncorrelated), at least before the interfering tone is initiated, and a standard uniform ADC would efficiently convert it to a sequence of bits.

Unfortunately, as described earlier, estimating the carrier frequency of the communication signal of interest is often a highly challenging task under the required latency constraints. Hence, the ADC must be applied to the sampled signal depicted in Fig. 2(a). A standard ADC is extremely inefficient for such a

signal, as it fails to exploit its sparsity in the frequency domain. The mean-squared error (MSE) attained by a standard ADC is governed by its dynamic range, which is determined via an AGC mechanism. In order to prevent overload errors that result in saturation of the ADC, the dynamic range is set proportionally to the signal’s average power. The signal’s power is unaffected by the fact that the sampling rate is significantly higher than the size of the essential support of the signal in the frequency domain. Thus, the number of bits per-second a standard ADC must output in order to reach some target MSE is significantly increased due to the uncertainty in the carrier frequency of the signal of interest.

To tackle this shortcoming of standard ADCs, this work develops a robust architecture, based on the emerging modulo ADC framework, which efficiently exploits the underlying structure of the input signal. Previous work [3] has shown that when the statistics of the signal to be acquired are known, or when its power spectral density (PSD) is at least confined to a particular frequency interval known in advance, modulo ADCs attain significant performance gains over standard uniform ADCs. Here, we develop an “AGC equivalent” mechanism for modulo ADCs, that obviates the need for prior knowledge of the signal’s statistics; See Fig. 3. This mechanism, together with suitably designed adaptive filtering, results in a *blind* modulo ADC, whose performance approaches that reported in [3] for stationary signals with known PSD.

Fig. 2(b) depicts the instantaneous squared error attained by the developed blind modulo ADC architecture for the signal from Fig. 2(a). For comparison, we also plot the squared error attained by a standard uniform ADC for the same signal. It is assumed that a perfect AGC is used for the standard ADC, such that its dynamic range is equal to $\kappa \sqrt{\text{Var}(x_n)}$, where $\{x_n\}$ is the input signal, and κ is a confidence parameter determining the overload probability. The same value of κ is used for both the standard and the modulo ADC systems, such that the overload probabilities for the two systems are similar.¹ Furthermore, both the standard and the blind modulo ADC systems use $R = 4$ bits per sample. It is evident that: (i) The developed blind modulo ADC architecture attains a significantly smaller MSE

¹The overload event for a modulo ADC is the event that the prediction error’s magnitude exceeds the dynamic range, as will be explained in detail, and explicitly defined in the sequel.

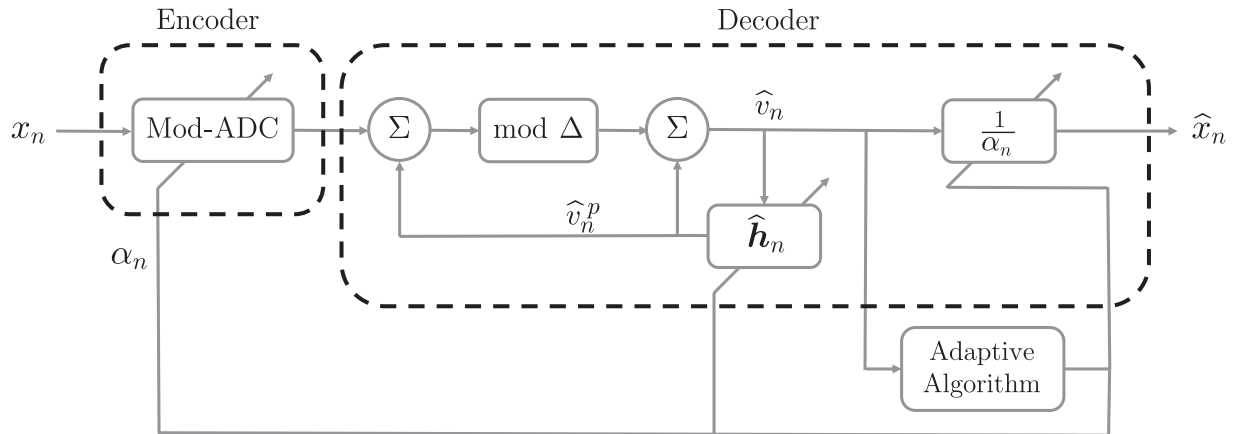


Fig. 3. A schematic block diagram of the blind mod-ADC encoder-decoder. A block with a diagonal arrow represents an adaptive operation (e.g., filtering).

than the one attained by a standard ADC; (ii) It quickly adapts to changes in the characteristics of the input signals, as reflected by Fig. 2(b); and (iii) While the addition of narrowband interferers strongly degrades the performance of the standard ADC, the MSE attained by the blind modulo ADC architecture is largely unaffected.

B. Related Work

The idea of using modulo ADCs/quantizers for exploiting temporal correlations within a stationary input process towards reducing the quantization rate R , dates back, at least, to [6], where a quantization scheme, called modulo-PCM, was introduced. Under the so called “high-resolution” assumption, which restricts the quantization’s error PSD to be much smaller than that of the signal *for all* frequencies [7], the analysis in [6] has shown that this scheme can attain distortion almost as small as the fundamental information theoretic lower bounds. Unfortunately, the “high-resolution” assumption breaks down completely for processes whose PSD function is not supported on the entire spectrum. Such processes include, for example, the process from Fig. 2(a), as well as any oversampled process. To that end, building on [7], a different modulo unwrapping algorithm was developed in [3], and the resulting modulo ADC system was shown to attain distortion close to the fundamental information theoretic lower bounds, even for processes for which the “high-resolution” assumption fails. Furthermore, relying on the unwrapping techniques developed in [8], a modulo ADC framework accompanied by an unwrapping algorithm was developed for vector processes, that are correlated in both space and time. Finally [3] also developed an architecture for a ring-oscillators-based circuit implementing a modulo ADC.

It should be noted that the results mentioned above rely on *complete* knowledge of the statistical law governing the inputs to the ADCs, with the exception of the result in [3], Section III, which is robust, but is of a minimax nature, in contrast to the pointwise optimality we seek here. For the case of temporally uncorrelated vector processes, Romanov et al. [9] developed a blind unwrapping algorithm which achieves performance close to that of an informed unwrapping algorithm, fully aware of the statistics. While a stationary process in time can be treated as a vector process in high dimensions, the scaling of the sample complexity of the algorithm from [9] renders it prohibitive for the blind modulo ADC problem of time processes under consideration in this work.

The line of research described above considers the improvement modulo ADCs offer over standard ADCs in terms of the trade-off between quantization rate and MSE distortion. The current paper continues this line of work. Another line of work which has received attention recently is that of the so-called “unlimited sampling”. Under the unlimited sampling framework, the quantization noise is usually not accounted for, and the focus is on characterizing the conditions which guarantee that a signal can be reconstructed from its folded version [10], [11], [12], [13]. In particular, it was shown that under mild conditions, a continuous time bandlimited signal can be recovered from its modulo reduced samples, provided that the sampling rate exceeds Nyquist’s rate [12], [14], [15], regardless of the modulo size. Some of the more recent work on unlimited sampling [16], [17] does take quantization noise into account, but adopts a worst-case model for the input signal (over a predefined class of signals), whereas here we model the input signal to the ADC as a stochastic process, and accordingly, analyze the statistical behavior of the MSE.

C. Contributions

In light of all the above, it is clear that a significant step towards realizing the modulo ADC technology is by developing the algorithmic framework, which will provide the essential robustness with respect to different types of signals and dynamic environments. Hence our motivation is developing an architecture with the appropriate algorithmic framework, which on one hand will be able to adapt quickly to changes reflected in the temporal structure of the input signal, and on the other hand will still provide reliable and stable high-resolution analog-to-digital conversion.

In order to appreciate our contributions, it is instructive to consider the trade-offs exhibited by several ADC architectures, as summarized in Table I. The table compares between a standard ADC, a standard ADC with AGC, an informed modulo ADC as described in [3], and the blind modulo ADC architecture we develop here. We compare the four solutions in terms of the statistical knowledge they require, and their performance guarantees. To simplify the exposition, suppose, for example, that the input to the ADC $\{x_n\}$ is a zero-mean stationary Gaussian process, with a (possibly unknown) PSD.

A standard (uniform) ADC has a fixed dynamic range. In order for overload events to be rare, such that the ADC is usually not saturated, the dynamic range must be greater than the standard

TABLE I
COMPARISON OF STANDARD AND MODULO ADCs WITH AND WITHOUT SIDE INFORMATION (SI) FOR STATIONARY GAUSSIAN PROCESSES INPUT SIGNALS

| | Standard ADC | Informed Mod-ADC |
|-------------|--|-------------------------|
| Encoder SI | Input signal variance, σ_x^2 | Innovation variance |
| Decoder SI | Input signal variance, σ_x^2 | PSD of the input signal |
| Performance | Near minimax optimal with variance constraint σ_x^2 | Near point-wise optimal |
| | Standard ADC with AGC | Blind Mod-ADC |
| Encoder SI | None | None |
| Decoder SI | None | None |
| Performance | Near minimax optimal for the (unknown) variance σ_x^2 | Near point-wise optimal |

deviation of $\{x_n\}$, denoted by σ_x , by some constant factor. Thus, in the design of the encoder and decoder, it is implicitly assumed that (an upper bound on) σ_x is known. The standard ADC cannot exploit any “memory” in the process $\{x_n\}$, but for an i.i.d. process it attains a rate-distortion trade-off which is quite close to the fundamental information theoretic limits, characterized by the rate-distortion function of the source [18]. Thus, it is near minimax optimal with respect to the class of all PSDs with variance σ_x^2 . A standard ADC with an AGC automatically adapts its dynamic range to σ_x , and does not require prior knowledge of it. Consequently, it attains near minimax optimality for the class of all PSDs with variance σ_x^2 , simultaneously for all values of σ_x^2 .

The informed modulo ADC from [3] requires the encoder to set the modulo size (or the signal scaling) appropriately, which requires knowledge of the variance of the innovation process (i.e., the error process due to optimal prediction). The decoder requires knowledge of the entire PSD in order to compute the coefficients of the optimal prediction filter it uses. It was shown [3] that for input processes of finite differential entropy rate, the rate-distortion trade-off this architecture attains is near optimal, as the quantization rate increases.

Clearly, a commercial ADC cannot be designed under the assumption that the innovation variance and the entire PSD of the input process is known in advance. In this paper, we close this gap and develop the blind modulo ADC architecture that makes no assumptions on the input process in the design of the encoder and the decoder, but nevertheless attains the same asymptotic performance as the modulo ADC architecture from [3]. In particular, our developed architecture asymptotically nearly attains the optimal rate-distortion trade-off *simultaneously* for all process with a finite differential entropy rate. The blind modulo ADC scheme we develop here adapts the modulo size / signal scaling at the encoder according to the associated innovation variance of the input process. Note that this task is considerably more challenging than that of an AGC in a standard ADC, since estimating the innovation variance is more involved than estimating the variance itself. Moreover, the decoder in a blind modulo ADC is implicitly estimating the necessary SOSs (for means of optimal prediction) beyond merely variance, i.e., cross correlations between past and present samples. Furthermore, the decoder blindly unwraps the quantized signal from the modulo measurements.

Our two main contributions in this work are the following:

- *Adaptive Algorithm for Blind Modulo Unwrapping:* We propose a feedback solution algorithm for a modulo ADC encoder-decoder, which blindly unwraps the modulo folding of the input signal. That is, our algorithm does not use prior knowledge on the temporal structure (i.e., the

autocorrelation function) of input signal. Nevertheless, using the least mean squares (LMS) algorithm [19], we are able to learn (only) the required SOSs, which allow us to exploit the unknown temporal structure, and gradually increase the resolution of the modulo ADC. Consequently, our developed blind modulo ADC architecture is more robust and practical than the one proposed in [3], which is designed based on such prior knowledge.

- *Asymptotic Performance Analysis of the Blind Modulo ADC Architecture:* We analyze the asymptotic performance of the developed algorithm in terms of the attainable resolution. We derive and present an insightful closed-form expression for the MSE distortion, which not only forecasts the best attainable performance under the specified conditions (dictated by the system parameters), but also intuitively explains the fundamental accuracy-stability trade-off inherent to the blind nature of the problem under consideration. Moreover, a steady state detector naturally stems from this analysis, allowing us to estimate the time at which the adaptive process can be (locally) paused. Consequently, the stability of the proposed method is increased, and as a (positive) byproduct, the overall computational load is reduced.

D. Paper Organization

The rest of the paper is organized as follows. The remainder of this section is devoted to a short outline of our notations. Section II is devoted to a brief review of the modulo ADC framework previously presented in [3], setting the premises for the current work. In Section III we formulate the problem of blind modulo ADC. Our proposed adaptive solution algorithm is presented in Section IV, where we derive the different algorithmic components, discuss key system parameters, trade-offs, and the asymptotic performance. Simulation results, corroborating our analytical derivation, are presented in Section V, and concluding remarks are given in Section VI.

E. Notations

We use x and \mathbf{x} for a scalar and a column vector, respectively. The superscript $(\cdot)^T$ denotes the transposition. We use $\mathbb{1}_{\mathcal{A}}$ to denote the indicator function of the event \mathcal{A} , namely $\mathbb{1}_{\mathcal{A}} = 1$ if \mathcal{A} is true, and $\mathbb{1}_{\mathcal{A}} = 0$ otherwise. $\mathbb{E}[\cdot]$ and $\text{Var}(\cdot)$ denote expectation and variance, respectively, $\text{Tr}(\cdot)$ denotes the trace operator, and $\|x_n\|_{\ell_N^\infty} \triangleq \sup_{n \in \{1, \dots, N\}} |x_n|$. We use $\hat{\cdot}$ to denote an estimator, e.g., \hat{x} is an estimator of x .

II. REVIEW ON A MODULO ADC

In this section, we briefly review the modulo ADC (encoding-decoding) algorithm previously proposed in [3] for scalar stationary processes. As our proposed blind method relies on some similar fundamental concepts, it is instructive to review the “informed” algorithm, which is described below.

For a positive number $\Delta \in \mathbb{R}^+$, we define

$$[x] \bmod \Delta \triangleq x - \Delta \cdot \left\lfloor \frac{x}{\Delta} \right\rfloor \in [0, \Delta), \quad \forall x \in \mathbb{R},$$

as the $[\cdot] \bmod \Delta$ operator, where $[x]$ is the floor operation, which returns the largest integer smaller than or equal to x . An R -bit modulo ADC with resolution parameter α , termed (R, α) mod-ADC, produces its output by first computing

$$[x]_{R, \alpha} \triangleq \llbracket \alpha x \rrbracket \bmod 2^R \in \{0, 1, \dots, 2^R - 1\}, \quad (1)$$

and then producing the binary representation of (1). A schematic illustration of the mod-ADC is given in Fig. 1.

Notice that when writing $[x]_{R,\alpha}$ as

$$[x]_{R,\alpha} = [\alpha x + \underbrace{([\alpha x] - \alpha x)}_{\triangleq \tilde{z}}] \bmod 2^R = [\alpha x + \tilde{z}] \bmod 2^R,$$

we identify $\tilde{z} \in (-1, 0]$ as the quantization error of a uniform scalar quantizer [20]. Although \tilde{z} is a deterministic function of x , this quantization error can be modeled quite accurately as additive *random* uniform noise. For details on the justification of this (standard) assumption by using subtractive dithers [21], Ch. 4], see [3], Section II]. We note, however, that in addition to the fact that the mild conditions underlying this modeling assumption are generally satisfied in our setting, the simulation results in Section V also demonstrate that it is a good approximation to the true nature of the quantization error. Under this assumption, an (R, α) mod-ADC is viewed as a stochastic channel, whose output y for an input x is given by

$$y = [\alpha x + z] \bmod 2^R,$$

where $z \sim \text{Unif}((-1, 0])$. Obviously, since the modulo operation is a form of *lossy* compression, it is generally impossible to recover the unfolded signal $\alpha x + z$ from its folded version $y = [\alpha x + z] \bmod 2^R$. Nevertheless, under relatively mild conditions, when the input signal is “temporally-predictable” to a sufficient degree, e.g., a correlated random process [3] or a deterministic bandlimited signal [10], [14], it is in fact possible to *perfectly* recover the unfolded signal² from its past samples and its current folded sample via causal processing.

More specifically, consider an (R, α) mod-ADC whose input signal x_n is a zero-mean stationary random process, with a known autocorrelation function $R_x[\ell] \triangleq \mathbb{E}[x_n x_{n-\ell}] \in \mathbb{R}$, whose one-sided support is assumed to be at least of (discrete) length $p \in \mathbb{N}^+$. The output of the mod-ADC is given by

$$y_n = [\alpha x_n + z_n] \bmod 2^R, \quad \forall n \in \mathbb{N}^+,$$

where $\{z_n \sim \text{Unif}((-1, 0])\}$, modeling the quantization noise, is an independent, identically distributed (i.i.d.) stochastic process. Further, define the unfolded quantized signal,

$$v_n \triangleq \alpha x_n + z_n, \quad \forall n \in \mathbb{N}^+,$$

and assume that the decoder has access to $\{v_{n-1}, \dots, v_{n-p}\}$, which is equivalent to assuming that the last p samples of y_n were correctly decoded. This can be achieved, for example, by proper initialization with a sufficiently small resolution parameter α , a notion that will also be used as part of our proposed method. For additional justifications of this assumption, see [3], Section II-A. Note that once v_n is recovered, x_n is readily estimated as $\hat{x}_n = \frac{v_n + \frac{1}{2}}{\alpha}$. Thus, the focus is on recovering v_n from y_n and $v_n \triangleq [v_{n-1} \cdots v_{n-p}]^T \in \mathbb{R}^{p \times 1}$.

The decoding algorithm proposed in [3] for recovering v_n w.h.p. when $R_x[\ell]$ is *known*, here referred to as *oracle* modulo unfolding, is given in Algorithm 1. The main idea behind the prescribed technical steps is the following. Every number $\gamma \in \mathbb{R}^+$ (similarly for $\gamma \in \mathbb{R}^-$) can be represented as

$$\gamma = \underbrace{(\Delta + \Delta + \dots + \Delta)}_{K_\gamma \in \mathbb{N}^+ \text{ times}} + \underbrace{[\gamma]}_{\triangleq \varepsilon_\gamma} \bmod \Delta = K_\gamma \Delta + \varepsilon_\gamma, \quad (2)$$

²With high probability (w.h.p.) for random signals, and to an arbitrary precision for deterministic bandlimited signals (“w.h.p.” in the sense that the probability of perfect recovery can be made arbitrarily large by increasing R).

Algorithm 1: Oracle Modulo Unfolding

$\hat{v}_{\text{oracle},n} = \text{ModUnfold}(y_n, v_n, R_x[\ell], \alpha, R)$.

Input: $y_n, v_n, R_x[\ell], \alpha, R$ **Output:** $\hat{v}_{\text{oracle},n}$

1 Compute the linear minimum MSE (LMMSE) predictor of v_n based on v_n

$$\hat{v}_{\text{LMMSE},n}^p = \mathbf{h}_{\text{opt}}^T (v_n + \frac{1}{2}\mathbf{1}) - \frac{1}{2}, \quad (3)$$

where $\mathbf{h}_{\text{opt}} \in \mathbb{R}^{p \times 1}$ is the length- p finite impulse response (FIR) filter yielding the LMMSE predictor, computed based on $R_x[\ell]$, and the shifts are to compensate for the non-zero mean $\mathbb{E}[z_n] = -\frac{1}{2}$;

2 Compute

$$w_{\text{LMMSE},n} = [y_n - \hat{v}_{\text{LMMSE},n}^p] \bmod 2^R,$$

$$\hat{e}_{\text{LMMSE},n}^p \triangleq ([w_{\text{LMMSE},n} + \frac{1}{2}2^R] \bmod 2^R) - \frac{1}{2}2^R;$$

3 Return $\hat{v}_{\text{oracle},n} = \hat{v}_{\text{LMMSE},n}^p + \hat{e}_{\text{LMMSE},n}^p$.

where, intuitively, K_γ and ε_γ correspond to coarse and fine information, respectively, in the “ Δ -representation” (2). The mod-ADC records only the fine information ε_γ in γ . Hence, for perfect reconstruction, only K_γ is required (assuming Δ is known). Conceptually, this means that as long as an estimator of v_n (possibly linear) has a minimal accuracy level, such that its residual estimation error lie in $[0, \Delta)$, K_γ can be recovered, which, in turn, means that v_n can be perfectly recovered.

An elaborate analysis of Algorithm 1 is provided in [3], wherein analytical performance guarantees are derived in the form of upper bounds on the probability of the *overflow* event, which inflicts $\hat{v}_n \neq v_n$, and is defined as

$$\mathcal{E}_{\text{OL},n}^* \triangleq \{ |e_{\text{LMMSE},n}^p| \geq \frac{1}{2}2^R \} = \{ \hat{e}_{\text{LMMSE},n}^p \neq e_{\text{LMMSE},n}^p \}, \quad (4)$$

where $e_{\text{LMMSE},n}^p \triangleq v_n - \hat{v}_{\text{LMMSE},n}^p$, and on the conditional MSE,

$$D \triangleq \mathbb{E} \left[(x_n - \hat{x}_n)^2 | \mathcal{E}_{\text{OL},n}^* \right]. \quad (5)$$

Specifically, it was shown that [3, Proposition 1],

$$\Pr(\mathcal{E}_{\text{OL},n}^*) \leq 2 \exp \left\{ -\frac{3}{2} 2^{2(R - \frac{1}{2} \log_2(12\sigma_{\text{LMMSE},p}^2))} \right\},$$

$$D \leq [12\alpha^2 (1 - \Pr(\mathcal{E}_{\text{OL},n}^*))]^{-1}, \quad (6)$$

where $\sigma_{\text{LMMSE},p}^2 \triangleq \mathbb{E}[(e_{\text{LMMSE},n}^p)^2]$ is the MSE of the LMMSE predictor based on the previous p samples, as in (3).

Algorithm 1, along with its information-theoretic analysis [3], provide strong evidence regarding the potential feasibility and merits of mod-ADCs, which are attractive for approaching the minimal number of raw output bits per sample, for a given sampling frequency f_s and a prespecified distortion level D .

Yet, devices such as ADCs usually operate under dynamic conditions, giving rise to a wide range of possible inputs with unknown characteristics, and must still maintain proper operation. Therefore, one significant step towards implementing mod-ADCs for real-life applications can be made by relaxing the (sometimes too restrictive) assumption that $R_x[\ell]$ is known. We take this significant step in the next sections.

III. PROBLEM FORMULATION

Consider an (R, α_n) mod-ADC as described in the previous section, with a fixed modulo range $\Delta = 2^R$, but an *adaptable*, possibly *time-varying* resolution parameter $\alpha_n \in \mathbb{R}^+$. The mod-ADC is fed with the input discrete-time signal $\{x_n \triangleq x(nT_s) \in$

$\mathbb{R}\}_{n \in \mathbb{N}^+}$, acquired by sampling the analog, continuous-time signal $x(t)$ every $T_s = f_s^{-1}$ seconds. We assume that x_n is a zero-mean stationary stochastic process with an *unknown* autocorrelation function $R_x[\ell] \triangleq \mathbb{E}[x_n x_{n-\ell}]$. The observed, distorted signal at the output of the mod-ADC reads

$$y_n = [\alpha_n x_n + z_n] \bmod 2^R, \quad \forall n \in \mathbb{N}^+, \quad (7)$$

where the quantization noise process $\{z_n \sim \text{Unif}((-1, 0])\}$ is i.i.d. Further, we redefine the unfolded quantized signal,

$$v_n \triangleq \alpha_n x_n + z_n, \quad \forall n \in \mathbb{N}^+, \quad (8)$$

which, in general, is no longer stationary. Nonetheless, when α_n is held fixed on a specific time interval, then v_n can be regarded as stationary on that particular interval.

The primary goal in this context is to estimate the input signal x_n as accurately as possible based on the observed sequence $\{y_n\}$ at the output of the mod-ADC using a causal system. However, since v_n is merely a scaled version of x_n contaminated by white noise (8), the problem essentially boils down to recovering v_n , and is stated concisely as follows.

Problem Statement: For a fixed number of bits R , design an adaptive mechanism for estimating $\{x_n\}$ from the output of the mod-ADC with the lowest possible MSE distortion, *without prior knowledge on* $R_x[\ell]$.

An interpretation of this problem statement is to design an update mechanism for maximizing the resolution parameter α_n , while still allowing for reliable recovery of v_n from $\{y_k\}_{k \leq n}$, and design the recovery mechanism.

As explained in Section II, provided v_n is *exactly* recovered w.h.p., i.e., $\hat{v}_n \stackrel{\text{w.h.p.}}{=} v_n$, the input signal is readily estimated as

$$\hat{x}_n \triangleq \frac{\hat{v}_n + \frac{1}{2}}{\alpha_n}, \quad (9)$$

where α_n is a known system parameter, and $\frac{1}{2}$ is to compensate for the quantization noise (non-zero) expectation $\mathbb{E}[z_n] = -\frac{1}{2}$.

IV. BLIND MODULO ADC CONVERSION

In this section, we present the blind mod-ADC algorithm, which simultaneously estimates the input signal x_n and performs online learning of the (possibly time-varying) SOSs of the unfolded quantized signal (8), necessary for estimation of x_n . We note that a key characterizing quantity of interest, to be used at some parts of the following derivation, is the ratio

$$M_n \triangleq \frac{\Delta}{\alpha_n} = \frac{2^R}{\alpha_n} \in \mathbb{R}^+, \quad (10)$$

dubbed the *effective modulo range*, rather than Δ or α_n individually. Although theoretically M_n could be adapted by fixing the resolution parameter and adapting the modulo range, due to practical considerations in the actual implementation of the modulo operation, we keep Δ fixed, and vary the resolution parameter α_n . This mechanism can be realized by changing the gain of the input x_n before feeding it to the mod-ADC.

The structure of the proposed blind mod-ADC is depicted in Fig. 3. Note that, in contrast to an informed mod-ADC (cf. Fig. 3 in [3]), here both the encoder and decoder are adaptive, and vary with time according to the statistical properties of the input signal. The price paid for the expected robustness we enjoy by using the blind mod-ADC is mainly in the form of an *adaptive* filter, rather than a pre-defined, constant one.

The underlying concept of our approach is the following. For a fixed resolution parameter α_n , given that at any time

instance n the unfolded signal v_n can be *exactly* recovered, we may estimate the optimal length- p FIR filter \mathbf{h}_{opt} , corresponding to the optimal LMMSE predictor of v_n based on the last consecutive p samples $\{v_{n-1}, \dots, v_{n-p}\}$. This can be done, e.g., using the celebrated LMS algorithm [19], which converges³ to \mathbf{h}_{opt} . Upon convergence, the resolution parameter α_n can be slightly increased, and as long as the prediction error of the linear causal prediction—currently no longer optimal—is sufficiently small, v_n could still be recovered using the same technique as in Algorithm 1. Now, fixing α_n again to its new value, the FIR filter can be adapted again to the optimal one using the LMS algorithm. The process is repeated until a certain level of effective modulo range is attained. This level, reflecting the desired trade-off between the MSE distortion D (5) and the probability of an overload event $\mathcal{E}_{\text{OL}_n}^*$ (4), will be later on discussed in detail.

Intuitively, and informally, only appropriate initial conditions and sufficiently smooth transitions from one resolution level to another are required for convergence of the above adaptive process. Conceptually, once these are fulfilled, we attain successful steady state operation of a blind mod-ADC (i.e., $R_x[\ell]$ unknown), in the desired effective modulo range.

Fortunately, with careful attention to more, important and relevant, details, this idea can be realized, and is rigorously described as our algorithm in the following subsections.

A. Phase 1: Initialization

We begin with a “small” initial value for the resolution parameter, α_0 (equivalently, $M_0 = 2^R/\alpha_0$), that ensures an essentially degenerated modulo operation, i.e., $\tilde{y}_n = v_n$, where

$$\tilde{y}_n \triangleq \left(\left[y_n + \frac{1}{2} 2^R \right] \bmod 2^R \right) - \frac{1}{2} 2^R, \quad (11)$$

such that \tilde{y}_n is the “modulo-shifted” version of y_n . Note that, since x_n is zero-mean, y_n actually undergoes a modulo operation quite often (roughly half of the time when $\text{Pr}(x_n < 0) = 0.5$). However, this is not an essential modulo due to a large amplitude of the input $\alpha_n x_n$, and is merely due to the fact that the dynamic range under consideration is $[0, 2^R)$, rather than $[-\frac{1}{2} 2^R, \frac{1}{2} 2^R)$. Nevertheless, we stick to this formulation as it more accurately describes the actual realization of our proposed method. For purposes that will become clear in the sequel, we further define for convenience

$$\bar{v}_n \triangleq \frac{v_n + \frac{1}{2}}{\alpha_n} = x_n + \frac{z_n + \frac{1}{2}}{\alpha_n}, \quad (12)$$

the “ α_n -standardized” version of v_n . Note that (12) still depends on the resolution parameter α_n . However, since

$$\mathbb{E}[\bar{v}_n] = \mathbb{E}[x_n] + \frac{\mathbb{E}[z_n + \frac{1}{2}]}{\alpha_n} = 0,$$

$$\mathbb{E}[\bar{v}_n^2] = \mathbb{E}[x_n^2] + \frac{\mathbb{E}[(z_n + \frac{1}{2})^2]}{\alpha_n^2} \triangleq \sigma_x^2 + \frac{1}{12\alpha_n^2}, \quad (13)$$

when α_n is sufficiently large, the variance of \bar{v}_n is dominated by σ_x^2 , and can be considered to be approximately independent of the system parameter α_n for certain needs. Of course, during the initialization phase, this is (still) not the case.

Assuming that $\tilde{y}_n = v_n$ as long as $\alpha_n = \alpha_0$ is fixed, the optimal length- p FIR filter for prediction of v_n (8) based on $\bar{\mathbf{v}}_n \triangleq [\bar{v}_{n-1} \dots \bar{v}_{n-p}]^T \in \mathbb{R}^{p \times 1}$ (12), which is merely a shifted-scaled version of \mathbf{v}_n , can be estimated with the LMS algorithm

³In the mean sense, under mild conditions stated explicitly in the sequel.

Algorithm 2: Blind Modulo Unfolding

$$\hat{v}_n, \hat{v}_n^p = \text{BlindModUnfold}(y_n, \hat{v}_n, \hat{h}_n, R).$$
Input: $y_n, \hat{v}_n, \hat{h}_n, R$ **Output:** \hat{v}_n, \hat{v}_n^p
1 Compute the linear predictor of v_n based on \bar{v}_n

$$\hat{v}_n^p \triangleq \hat{h}_n^T \hat{v}_n - \frac{1}{2}, \quad (14)$$

 where the shift is to compensate for $\mathbb{E}[v_n] = \frac{1}{2}$ (8);

2 Compute

$$w_n = [y_n - \hat{v}_n^p] \bmod 2^R, \quad (15)$$

$$\hat{e}_n^p = \left([w_n + \frac{1}{2}2^R] \bmod 2^R \right) - \frac{1}{2}2^R;$$

3 Return $\hat{v}_n \triangleq \hat{v}_n^p + \hat{e}_n^p, \hat{v}_n^p$.

using the following update equation [19],

$$\hat{h}_n = \hat{h}_{n-1} + \mu \cdot \bar{v}_n e_n^p. \quad (16)$$

 Here, \hat{h}_n is the FIR filter used in Algorithm 2 for the recovery of v_n , μ is the learning rate (or step size), and

$$e_n^p \triangleq v_n - \hat{v}_n^p \quad (17)$$

 is the prediction error of the linear predictor \hat{v}_n^p as in (14). It should be emphasized that, in practice, we never have access to the true error e_n^p , but only to \hat{e}_n^p , defined in (15). However, for simplicity of the exposition⁴, and as mentioned above, we assume that $\tilde{y}_n = v_n$, which means that $\hat{e}_n^p = e_n^p$, during the entire initialization phase, hence e_n^p appears in (16).

 In addition, rather than using $\{v_n\}$ (8), we use the “ α_n -standardized” process $\{\bar{v}_n\}$ as the observations in (14), since as the adaptive process evolves and α_n increases, the SOSs of \bar{v}_n gradually become less affected by α_n (13). This alleviates the estimation (/learning) of the optimal filter coefficients, as explained in detail in Appendix A.

 A discussion on the convergence of the LMS algorithm, as well as how to choose the appropriate step size μ which guarantees this convergence, will be given in Subsection IV-D. For now, assume that μ is chosen so as to ensure that [22],

$$\alpha_n = \alpha_0 : \lim_{n \rightarrow \infty} \mathbb{E}[\hat{h}_n] = \mathbf{h}_{\text{opt}}, \quad (18)$$

 where \mathbf{h}_{opt} is the optimal length- p filter corresponding to the oracle LMMSE predictor, a function of $\alpha_n = \alpha_0$ and $R_x[\ell]$.

As an intermediate summary for the initialization, we have:

Phase 1: Initialization. Input: \hat{h}_0, α_0 .

1. Fix $\alpha_n = \alpha_0$, and accumulate $p + 1$ samples $\{y_i\}_{i=1}^{p+1}$;
2. Compute $\{\tilde{y}_i = v_i\}_{i=1}^p$ as in (11), and \bar{v}_{p+1} as in (12);
3. Set $\hat{h}_{p+1} = \hat{h}_0$;
4. For $n = p + 1, p + 2, \dots$ do
 - 4.1. $\hat{v}_n, \hat{v}_n^p = \text{BlindModUnfold}(y_n, \bar{v}_n, \hat{h}_n, R)$;
 - 4.2. $\hat{h}_{n+1} = \hat{h}_n + \mu \cdot \bar{v}_n e_n^p$.

 After enough iterations, since we assume that α_0 is sufficiently small to ensure that $\hat{v}_n = v_n$ for every n during initialization, which gives us access to e_n^p (17), the filter \hat{h}_n will approximately converge to an unbiased estimate of \mathbf{h}_{opt} , as in (18). Assuming

⁴The initial resolution parameter α_0 can be chosen such that $\Pr(\hat{e}_n^p = e_n^p)$ is arbitrarily close to 1. It can even be exactly equal to 1 in case some (possibly partial) knowledge about the support of x_n is available. At any rate, we touch upon and handle this aspect more accurately in the next subsection.

 the learning rate μ is sufficiently small, the MSE of \hat{v}_n^p will approximately converge to the MSE of $\hat{v}_{\text{LMMSE},n}^p$ (with α_n replacing α , according to the definition (8)),

$$\exists N_0 : \forall n > N_0 : \mathbb{E}[(e_n^p)^2] \underset{\mu \ll 1}{\approx} \sigma_{\text{LMMSE},p}^2.$$

 Accordingly, once $\mathbb{E}[(e_n^p)^2]$ is close enough to $\sigma_{\text{LMMSE},p}^2$, by virtue of (6), an overload will not occur w.h.p., namely,

$$\Pr(\mathcal{E}_{\text{OL}_n}) \triangleq \Pr(|e_n^p| \geq \frac{1}{2}2^R) \approx \Pr(\mathcal{E}_{\text{OL}_n}^*), \quad (19)$$

 where $\mathcal{E}_{\text{OL}_n}^*$ (4) refers to the overload event of the informed mod-ADC. In other words, the no overload event $\mathcal{E}_{\text{OL}_n} \triangleq \{\hat{e}_n^p = e_n^p\}$ —the complement of the overload event $\mathcal{E}_{\text{OL}_n}$ —occurs w.h.p. At this point, we can increase the resolution parameter α_n , so as to decrease M_n , and use the quantizer’s output raw bits to a finer description of the input signal.

B. Phase 2: Updating the Resolution Parameter

 As explained above, in order to increase α_n , we must somehow detect that the filter \hat{h}_n already approximates the optimal one well enough, such that the prediction errors e_n^p are sufficiently small with respect to the dynamic range $\Delta = 2^R$. When this is the case, a small change in the resolution would not affect our ability to recover v_n w.h.p. from y_n and \hat{v}_n^p .

 To see this, assume that we increase the resolution parameter $\alpha_n = \alpha_0 + \epsilon_\alpha$, where $\epsilon_\alpha \in \mathbb{R}^+$ is a small increment, and accordingly also scale the respective filter coefficients by $\frac{\alpha_0 + \epsilon_\alpha}{\alpha_0}$.

 Now, \hat{h}_n is no longer optimal, since the second-order statistical properties of $\{\bar{v}_\ell\}_{\ell \geq n}$ are different than those of $\{\bar{v}_\ell\}_{\ell < n}$, based on which \hat{h}_n has been estimated thus far. However, if ϵ_α is sufficiently small, a straightforward “small-error” analysis yields that \hat{h}_n will now only slightly deviate from the (approximately) optimal filter to the new value of α_n , such that right after increasing the resolution,

$$\alpha_n = \alpha_0 + \epsilon_\alpha \implies \hat{h}_n = \hat{h}_n^0 + \epsilon_n,$$

 where \hat{h}_n^0 denotes an unbiased estimator of the optimal FIR filter corresponding to the LMMSE predictor of v_n for the *updated* resolution parameter $\alpha_0 + \epsilon_\alpha$, and $\epsilon_n \in \mathbb{R}^{p \times 1}$ is a vector of “small” biases, due to ϵ_α . Accordingly, the MSE of the (currently no longer optimal) linear predictor \hat{v}_n^p , conditioned on the p no overload past events⁵ $\mathcal{E}_{\text{OL}_n}^{(p)} \triangleq \bigcap_{k=1}^p \mathcal{E}_{\text{OL}_{n-k}}$ enabling the exact recovery of v_n , will be slightly increased⁶

$$\sigma_{p,n}^2 \triangleq \mathbb{E}[(e_n^p)^2 | \mathcal{E}_{\text{OL}_n}^{(p)}] = \sigma_{\text{LMMSE},p}^2 + \underbrace{\sigma_\epsilon^2}_{\text{Due to } \epsilon_\alpha}. \quad (20)$$

 Nonetheless, as long as $\sigma_{p,n} \ll \frac{1}{2}2^R$, such that $\mathcal{E}_{\text{OL}_n}$ still occurs w.h.p., v_n is still exactly recovered w.h.p. using Algorithm 2. Indeed, an important observation is that v_n can be recovered using Algorithm 2 even when a suboptimal linear predictor is used in (14). For successful operation, we *only* require that the linear predictor is accurate enough to ensure that $\mathcal{E}_{\text{OL}_n}$ occurs

⁵Strictly speaking, $\sigma_{\text{LMMSE},p}^2$ in (20) of the blind mod-ADC is not equal to the MSE of (3), since for informed mod-ADC we do not condition on $\mathcal{E}_{\text{OL}_n}^{(p)}$. However, under mild conditions, stated explicitly below, the difference is negligible, hence we use the same notation for simplicity.

⁶For simplicity, we assume here that prior to increasing the resolution $\sigma_{p,n}^2 = \sigma_{\text{LMMSE},p}^2$, such that σ_ϵ^2 is *only* due to ϵ_α . In practice, we have $\sigma_{p,n}^2 \approx \sigma_{\text{LMMSE},p}^2$, such that σ_ϵ^2 encapsulates estimation errors due to \hat{h}_n as well. Still, after changing α_n , σ_ϵ^2 will be dominated by errors due to ϵ_α .

w.h.p. Consequently, for short transition periods in which the optimal filter is learned, a suboptimal filter would suffice.

Hence, we conclude the following:

- ① If the resolution parameter is adapted in small increments, we are able to maintain sufficiently small prediction errors, and safely continue recovering v_n w.h.p.; and
- ② Before increasing the resolution, we desire to arrive at an intermediate steady state, wherein $\mathcal{E}_{\text{OL}_n}$ occurs w.h.p.

Since ϵ_α is a user-controlled system parameter, ① can be easily achieved. As for ②, since we are operating in a blind scenario, where the distribution of the input x_n is unknown, it is generally unclear how to ensure rarity of no overload. Therefore, for this purpose only, we take the simplifying, but useful, assumption that $e_n^p | \mathcal{E}_{\text{OL}_n}^{(p)} \sim \mathcal{N}(0, \sigma_{p,n}^2)$. Note, however, that this assumption is not strictly required in order to analytically justify the derivation which follows, and is merely to simplify the exposition. In this case, $\sigma_{p,n}$ is directly related to $\mathcal{E}_{\text{OL}_n}^{(p)}$, conditioned on $\mathcal{E}_{\text{OL}_n}^{(p)}$. Specifically, if $\Delta/2 = \kappa \cdot \sigma_{p,n}$ for some $\kappa \in \mathbb{R}^+$, then we have

$$\Pr \left(|e_n^p| > \frac{1}{2} 2^R \left| \mathcal{E}_{\text{OL}_n}^{(p)} \right. \right) = 2Q \left(\frac{\Delta}{2\sigma_{p,n}} \right) = 2Q(\kappa), \quad (21)$$

where $Q(x) \triangleq \int_x^\infty \frac{1}{\sqrt{2\pi}} e^{-t^2/2} dt$ is the Q -function. Put simply, if the linear predictor is good enough, such that half the modulo range $\Delta/2$ is κ times greater than its root MSE (RMSE), and κ is sufficiently large, $\mathcal{E}_{\text{OL}_n}^{(p)} | \mathcal{E}_{\text{OL}_n}^{(p)}$ occurs w.h.p. This provides the conditions to re-learn the filter corresponding to the LMMSE predictor of v_n with the updated resolution α_n .

In practice, though, since $R_x[\ell]$ is unknown, $\sigma_{p,n}^2$ is clearly not known as well. Nevertheless, since $\hat{v}_n \stackrel{\text{w.h.p.}}{=} v_n$ throughout the adaptive process, we can estimate $\sigma_{p,n}$ online by

$$\hat{\sigma}_{p,n}^2 \triangleq \frac{1}{L_s} \sum_{k=0}^{L_s-1} (\hat{v}_{n-k} - \hat{v}_{n-k}^p)^2, \quad \hat{\sigma}_{p,n} \triangleq \sqrt{\hat{\sigma}_{p,n}^2}, \quad (22)$$

where $L_s \in \mathbb{N}^+$ is a moving average window length, and is also set to be the minimal (discrete) time stabilization interval wherein α_n is kept fixed after its last update. More details on the parameters L_s and κ are given in Subsection IV-D. Thus, to achieve ②, we increase α_n only when $\mathbb{1}_{\uparrow\alpha,n} = 1$, where

$$\mathbb{1}_{\uparrow\alpha,n} \triangleq \begin{cases} 1, & \kappa \cdot \hat{\sigma}_{p,n} < \frac{\Delta}{2} \\ 0, & \kappa \cdot \hat{\sigma}_{p,n} > \frac{\Delta}{2} \end{cases}. \quad (23)$$

Whenever $\mathbb{1}_{\uparrow\alpha,n} = 0$, we infer that the prediction errors are not satisfactorily small. In these cases, we decrease the resolution so as to resort to a state where v_n is again recovered w.h.p. By this, we allow the LMS filter (16) to converge to the desired filter, and then safely increase the resolution again. Following our previous observation on the accuracy required by the linear predictor \hat{v}_n^p , at this point the desired filter is not necessarily the optimal one, but is merely one attaining $\mathbb{1}_{\uparrow\alpha,n} = 1$. Accordingly, it is certainly possible that α_n would be increased before \hat{h}_n would converge to the optimal filter (at least before steady state, as discussed in the next subsections).

Note that, conditioned on $\mathcal{E}_{\text{OL}_{n+1}}^{(L_s)}$, (22) is a consistent estimate (with respect to L_s) of $\sigma_{p,n}$ for a wide class of signals, even when the errors $e_n^p | \mathcal{E}_{\text{OL}_n}^{(p)}$ are non-Gaussian. Hence, this mechanism is generally robust, and relies on Gaussianity only for (21).

Naturally, this assumption implies that the expected stability would be obtained for sub-Gaussian⁷ errors as well.

Upon updating α_n , we also appropriately update the filter \hat{h}_n , since the input signal \bar{v}_n is scaled with α_n as well (12). Therefore, it is convenient to use *multiplicative* updates, rather than additive, to α_n and \hat{h}_n . For a fixed $\delta_\alpha \in (0, 1)$, we update

$$\text{Increase Resolution : } \alpha_n = \delta_\alpha^{-1} \alpha_{n-1}, \quad \hat{h}_n = \delta_\alpha^{-1} \hat{h}_{n-1},$$

$$\text{Decrease Resolution : } \alpha_n = \delta_\alpha \alpha_{n-1}, \quad \hat{h}_n = \delta_\alpha \hat{h}_{n-1}.$$

A summary for the resolution updating phase is as follows:

Phase 2: Resolution Update. Input: κ, δ_α .

Assumption: α_n was held fixed (at least) L_s samples.

1. $\alpha_{n+1} = \left[\mathbb{1}_{\uparrow\alpha,n} \cdot \frac{1}{\delta_\alpha} + (\mathbb{1}_{\uparrow\alpha,n} - 1) \cdot \delta_\alpha \right] \alpha_n$;
2. $\hat{h}_{n+1} = \left[\mathbb{1}_{\uparrow\alpha,n} \cdot \frac{1}{\delta_\alpha} + (\mathbb{1}_{\uparrow\alpha,n} - 1) \cdot \delta_\alpha \right] \hat{h}_n$.

It is now straightforward to generalize this adaptive process, since, conceptually, we now only need to repeatedly execute the properly interlaced Phases 1 and 2. In the repeated Phase 1, the “initial” values for \hat{h}_n, α_n would be the corresponding values of the previous time step. Additionally, \bar{v}_n will be replaced by $\hat{v}_n \triangleq [\hat{v}_{n-1} \cdots \hat{v}_{n-p}]^T \in \mathbb{R}^{p \times 1}$, whose entries

$$\hat{v}_n \triangleq \hat{x}_n = \frac{\hat{v}_n + \frac{1}{2}}{\alpha_n}, \quad (24)$$

are $\{\hat{v}_{n-i} \stackrel{\text{w.h.p.}}{=} \bar{v}_{n-i}\}_{i=1}^p$, assuming previous successful recoveries w.h.p. The repeated Phase 2 would then be executed after (at least) L_s time steps with the updated resolution.

Note that we intentionally use in (24), and hereafter, the (seemingly redundant) notation \hat{v}_n rather than \hat{x}_n , since in this context we are actually trying to perfectly recover \bar{v}_n , defined in (12), rather than to estimate x_n . This is to enable the proper operation of the LMS algorithm (16), whose input should be $\{\bar{v}_{n-i}\}_{i=1}^p$, and not $\{x_{n-i}\}_{i=1}^p$ [3], [7].

Ideally, alternating between these two phases would lead to convergence near the limit $\kappa \cdot \sigma_{p,n} = \frac{\Delta}{2}$, as in (23), up to small fluctuations due to the limited-resolution adaptation step δ_α and estimation errors in $\hat{\sigma}_{p,n}$. However, recall that $\mathcal{E}_{\text{OL}_n}^{(p)} | \mathcal{E}_{\text{OL}_n}^{(p)}$, which implies that $\hat{e}_n^p = e_n^p$, and in turn $\hat{v}_n = v_n$, is *only* w.h.p., and in practice this is certainly not true for all $n \in \mathbb{N}^+$. Indeed, whenever $\mathcal{E}_{\text{OL}_n} = \{\hat{v}_n \neq v_n\}$ occurs, an extremely fast and destructive error propagation process begins. To detect such errors and prevent the consequent error propagation, we propose the defense mechanism presented next.

C. Error Propagation Prevention

One natural way of coping, and eventually preventing the aforementioned error propagation is by splitting the problem into two parts. The first part is to detect that an error has occurred, namely that v_n has not been perfectly recovered. In other words, the event $\mathcal{E}_{\text{OL}_n} = \{\hat{v}_n \neq v_n\}$ has to be detected, and as quickly as possible. The second part is, given that $\mathcal{E}_{\text{OL}_n}$ has been detected, to mitigate the error effect so as to reclaim a high-resolution functioning mod-ADC steady state.

Provided that an error event has been detected, a simple, though conservative mitigation solution is to fully “re-open”

⁷The real-valued random variable u is called *sub-Gaussian* if $\exists c, \gamma \in \mathbb{R}^+$ such that $\forall t > 0 : \Pr\{|u| > t\} \leq c \cdot \exp\{-\gamma t^2\}$.

the effective modulo range (10) to its initial value M_0 for (at least) p time steps. By this, we effectively initialize the process and guarantee that no errors occur, at the expense of (locally) retreating to a low-resolution regime. This solution, however, is useful only if the detection of \mathcal{E}_{OL_n} can be handled very accurately, i.e., with a very low false-alarm rate. Otherwise, the average operational time percentage of the mod-ADC in a degenerate modulo state (corresponding to a large M_n) would be high, and there would be no gain in using a mod-ADC. Hence, we turn our attention to the detection of \mathcal{E}_{OL_n} .

Formally, our goal now is to derive an estimator $\hat{\mathbb{1}}_{\mathcal{E}_{OL_n}} \in \{0, 1\}$ of the oracle indicator

$$\mathbb{1}_{\mathcal{E}_{OL_n}} \triangleq \begin{cases} 1, & \mathcal{E}_{OL_n} \\ 0, & \mathcal{E}_{OL_n}^c \end{cases} = \begin{cases} 1, & \hat{v}_n \neq v_n \\ 0, & \hat{v}_n = v_n \end{cases}, \forall n \in \mathbb{N}^+.$$

Since this is required at every time step n , and assuming that with $M_n = M_0$ there are no overload events, this is essentially a change detection problem (e.g., [23]). In particular, since $\mathcal{E}_{OL_n} | \mathcal{E}_{OL_n}^{(p)}$ ($\approx \mathcal{E}_{OL_n}^*$, (19)) is rare (6), it can also be viewed as a special instance of the fraud detection problem [24], where \hat{v}_n is pretending to be v_n , while in fact it is not, viz., $\hat{v}_n \neq v_n$.

Fortunately, our specific problem has favorable properties that allow us to develop an accurate detector. In particular, observe that increasing α_n essentially “pushes” \bar{v}_n towards approximate wide-sense stationarity. Specifically, using (12),

$$\mathbb{E}[\bar{v}_n \bar{v}_{n-\ell}] = R_x[\ell] + \frac{1}{12\alpha_n^2} \cdot \mathbb{1}_{\ell=0}, \quad (25)$$

such that even if α_n changes over time, for sufficiently large values of α_n , the autocorrelation of \bar{v}_n —even when unknown—can be considered as being approximately a function of ℓ only. Furthermore, it is seen from (25) that the variance of \bar{v}_n (13) is the only source of non-stationarity.

Similarly to our comment above (21), in our blind scenario, information such as (25) is not necessarily sufficient to be able to design an accurate detector of the event \mathcal{E}_{OL_n} . Therefore, at this point we again invoke Gaussianity, and assume that $\{x_n\}$ is a Gaussian process with an autocorrelation function $R_x[\ell]$. Under this assumption, we have the following result, whose proof is given in Appendix B.

Proposition 1: Assume $\min_n \{\alpha_n\} = \alpha_0$. Under no overloads, i.e., $\mathcal{E}_{OL_n}^{(n)}$, in the absence of estimation errors in $\hat{\sigma}_{p,n}$, and with an infinite resolution step size $\delta_\alpha \rightarrow 1$, define

$$\lim_{n \rightarrow \infty} \alpha_n \triangleq \alpha_\infty, \quad \bar{v}_n^\infty \triangleq \frac{v_n + \frac{1}{2}}{\alpha_\infty}, \quad \mathbb{E}[(\bar{v}_n^\infty)^2] \triangleq \sigma_v^2, \quad (26)$$

as the asymptotic resolution parameter, the ideal steady state process and its variance, respectively. Then, if the “decaying” condition $R_x[\ell] \cdot \log(\ell) \xrightarrow{\ell \rightarrow \infty} 0$ holds, it follows that

$$\Pr\left(\|\bar{v}_n^\infty\|_{\mathcal{N}} > \sqrt{2\sigma_v^2 \log(N)}\right) \xrightarrow{N \rightarrow \infty} 0. \quad (27)$$

This means that, asymptotically, knowing only the variance of the process \bar{v}_n , and observing its magnitudes, is sufficient in order to detect almost surely an abnormality in the form of a large, improbable deviation exceeding the threshold in (27).

Now, observe that an overload event \mathcal{E}_{OL_n} inflicts a prediction error in \hat{v}_n , and in turn in \bar{v}_n , of the order of Δ . Clearly, this

creates a large “discontinuity”⁸, which is exactly the abnormality form we can identify w.h.p. according to (27). In light of all the above, we propose

$$\hat{\mathbb{1}}_{\mathcal{E}_{OL_n}} \triangleq \begin{cases} 1, & |\hat{v}_n| > \sqrt{2\hat{\sigma}_{\bar{v},n}^2 \log(n)}, \forall n \geq N_s, \\ 0, & \text{otherwise} \end{cases} \quad (28)$$

as the detector of an error event due to \mathcal{E}_{OL_n} , where

$$\hat{\sigma}_{\bar{v},n}^2 \triangleq \frac{1}{n-1} \sum_{k=1}^{n-1} \hat{v}_k^2, \quad \forall n \geq 2, \quad (29)$$

and $N_s \in \mathbb{N}^+$ is a fixed stabilization time-interval, wherein (28) is still not sufficiently accurate, and we enforce a simple, more conservative condition for the transition phase $n \leq N_s$. For example, one reasonable choice could be

$$\hat{\mathbb{1}}_{\mathcal{E}_{OL_n}} \triangleq \begin{cases} 1, & |\hat{v}_n| > \beta \cdot \hat{\sigma}_{\bar{v},n}, \forall n < N_s, \\ 0, & \text{otherwise} \end{cases}$$

where $\beta \in \mathbb{R}^+$ is some predefined number (e.g., $\beta = 5$). From practical considerations, since the threshold value in (28) increase logarithmically with n , a plausible solution would be to reset the time-index in this threshold every error event \mathcal{E}_{OL_n} .

Once we observe $\hat{\mathbb{1}}_{\mathcal{E}_{OL_n}} = 1$, we set $\alpha_{n+1} = \alpha_0$, and reset the process as described in the beginning of this subsection. The proposed defense mechanism is summarized as follows:

Error Propagation Defense Mechanism

Assumption: $n \geq p$.

1. Update $\hat{\sigma}_{\bar{v},n}^2 = \frac{n-2}{n-1} \cdot \hat{\sigma}_{\bar{v},n-1}^2 + \frac{1}{n-1} \cdot \hat{v}_{n-1}^2$;
2. If $\hat{\mathbb{1}}_{\mathcal{E}_{OL_n}} = 1$
 - 2.1. Reset $\alpha_{n+1} = \alpha_0$, and adapt $\hat{h}_{n+1} = \frac{\alpha_0}{\alpha_n} \cdot \hat{h}_n$.

We note that, conditioned on the no overload event $\mathcal{E}_{OL_n}^c$, and assuming (26) holds, (29) is consistent, i.e., $\hat{\sigma}_{\bar{v},n}^2 \xrightarrow{P} \sigma_v^2$, where \xrightarrow{P} denotes convergence in probability as $n \rightarrow \infty$. Thus, the decision rule in (28) becomes increasingly accurate as we approach steady state, indeed, a desirable outcome.

D. Key System Parameters and Corresponding Trade-Offs

First and foremost, convergence of the adaptive process described above is conditioned on the no overload event. Therefore, the parameter κ , setting the probability that the prediction errors e_n^p are kept inside $(-\frac{\Delta}{2}, \frac{\Delta}{2})$, must be sufficiently large, so as to ensure that (21) is sufficiently low. For example, choosing $\kappa = 7$ already gives $\Pr(\mathcal{E}_{OL_n} | \mathcal{E}_{OL_n}^{(p)}) \propto 10^{-12}$. Yet, as κ increases, the asymptotic resolution α_∞ (26) decreases, as already alluded from (23). A formal characterization of this asymptotic trade-off is provided in the next subsection.

Given that κ was chosen properly, we continue with the convergence and asymptotic analysis, conditioned on no overload. Specifically, we now focus on the learning rate μ . Assuming momentarily that α_n is fixed, based on the well-established

⁸This, of course, is not a discontinuity in the formal sense as defined for deterministic functions. Rather, we use this term here informally to refer to an improbable transition from one value to another, in a manner that is inconsistent with $R_x[\ell]$, governing the statistical nature of $\{\bar{v}_n\}$.

theory of the LMS algorithm [19], if we choose μ such that

$$0 < \mu < [\text{Tr}(\mathbb{E}[\widehat{\bar{v}}_n \widehat{\bar{v}}_n^T])]^{-1} = \left[p \cdot \left(\sigma_x^2 + \frac{1}{12\alpha_n^2} \right) \right]^{-1}, \quad (30)$$

then the FIR filter $\widehat{\mathbf{h}}_n$ would converge in the sense (18), namely it will randomly fluctuate about \mathbf{h}_{opt} , corresponding to the LMMSE predictor. Recall that conditioned on no overload, $\widehat{\bar{v}}_n = \bar{v}_n$, and when α_n is fixed, \bar{v}_n is stationary in the respective time interval, hence the diagonal elements of $\mathbb{E}[\bar{v}_n \bar{v}_n^T] \in \mathbb{R}^{p \times p}$ are all equal to the variance (13), and the right hand side of (30) follows. Now, since α_n is in fact time-varying, we would like to choose μ such that

$$0 < \mu < \frac{1}{p \cdot \left(\sigma_x^2 + \frac{1}{12 \min_n \{ \alpha_n^2 \}} \right)} = \frac{1}{p \cdot \left(\sigma_x^2 + \frac{1}{12\alpha_0^2} \right)}. \quad (31)$$

However, the upper bound (31) is unknown, since σ_x^2 is unknown. Therefore, we propose to choose $\mu = \epsilon_\mu / (p \cdot \widehat{\sigma}_{\bar{v}, p+1}^2)$, where $\epsilon_\mu \in \mathbb{R}^+$ is some small constant (e.g., 10^{-2}), and $\widehat{\sigma}_{\bar{v}, p+1}^2$ (29) can be computed during the initialization phase. Since α_0^{-2} is typically small, and thus $\widehat{\sigma}_{\bar{v}, p+1}^2$ is dominated by σ_x^2 , our empirical experience indicates that choosing ϵ_μ appropriately, so as to ensure the desired stability, is rather easy. Furthermore, a longer initialization phase could be performed to yield a more accurate estimate of the variance of \bar{v}_n . Lastly, and although not necessary, μ could be easily adapted throughout the process based on the online estimate (29).

Another system parameter is L_s , the minimal discrete-time interval in which the resolution α_n is held fixed before another resolution update is allowed. In the extreme case $L_s \rightarrow \infty$, we have the highest stability (α_n is fixed, $\widehat{\mathbf{h}}_n$ converges) but the slowest (zero) progress towards high resolution. In the other extreme case $L_s = 1$, α_n can be updated at all times, but the estimate (22), and therefore the detector (23), becomes extremely inaccurate. Therefore, L_s should be set so as to appropriately handle this trade-off. Since Algorithm 2 assumes that the previous p samples of \bar{v}_n are available (via $\widehat{\bar{v}}_n$), it is reasonable to choose L_s proportional to p (e.g., $\text{round}(p/2)$).

The resolution step size parameter $\delta_\alpha \in (0, 1)$ also balances a similar trade-off. As δ_α decreases, the convergence towards α_∞ is faster. However, the LMS would be required to cope with more abrupt changes in the variance of \bar{v}_n , harming the predictor \widehat{v}_n^p , and thus locally inflicting larger prediction errors e_n^p , which in turn promote overload events. On the other hand, as $\delta_\alpha \rightarrow 1^-$ approaches 1 (from below), the transition becomes smoother, allowing the LMS to adjust conveniently, and as explained above, decrease the probability of overload events. Naturally, this comes at the cost of slower convergence to α_∞ .

To conclude this section, we refer to the parameter p , the length of the FIR filter of the linear predictor \widehat{v}_n^p . Preferably, p should be chosen based on some prior knowledge related to the specific application the mod-ADC is used. In particular, if the effective support⁹ L_x of the unknown autocorrelation function $R_x[\ell]$ is known even approximately, then an educated choice would be $p = L_x$. Indeed, if the (one-sided) support of $R_x[\ell]$ is *precisely* $L_x + 1$, then the causal Wiener filter [25], i.e., the optimal (generally not FIR) filter of the causal LMMSE predictor, for predicting the process $\{v_n\}$ based on $\{\bar{v}_k\}_{k < n}$ is an FIR filter of length $p = L_x$. Note in passing that L_x can be estimated during initialization, since $\tilde{y}_n = v_n$ (w.h.p.).

⁹For some $\epsilon > 0$, the ϵ -effective support of $R_x[\ell]$ is the number $L_x \in \mathbb{N}^+$ for which $\forall |\ell| > L_x : |R_x[\ell]| < \epsilon$. Loosely speaking, we say that L_x is the effective support of $R_x[\ell]$ when $\forall |\ell| > L_x : |R_x[\ell]| \approx 0$.

E. The Asymptotic Performance of a Blind Mod-ADC

Let us assume that all the parameters have been chosen such that an overload does not occur. In this ideal (merely theoretical) case, if we assume further that $\widehat{\sigma}_{p,n}^2 = \sigma_{p,n}^2$ and that an infinitely fine step size $\delta_\alpha \rightarrow 1$ is used, the resolution of the blind mod-ADC converges to α_∞ (26). In this asymptotic state, we have the following equilibrium proposition, whose proof appears in Appendix C.

Proposition 2: Assume that an overload does not occur, $\widehat{\sigma}_{p,n}^2 = \sigma_{p,n}^2$ and $\delta_\alpha \rightarrow 1$. Then, as $n \rightarrow \infty$, we have

$$\bar{\sigma}_{p,\infty} = \frac{1}{2\kappa} \cdot \frac{\Delta}{\alpha_\infty} \triangleq \frac{M_\infty}{2\kappa}, \quad (32)$$

where

$$\bar{\sigma}_{p,n}^2 \triangleq \mathbb{E} \left[\left(\bar{v}_n - \widehat{\bar{v}}_n^p \right)^2 \middle| \mathcal{E}_{\text{OL}n+1}^{(n)} \right] \Rightarrow \bar{\sigma}_{p,\infty}^2 \triangleq \lim_{n \rightarrow \infty} \bar{\sigma}_{p,n}^2, \quad (33)$$

and the linear predictor of the “ α_n -standardized” process \bar{v}_n , which appears in (33), is defined as,

$$\widehat{\bar{v}}_n^p \triangleq \frac{\widehat{v}_n^p + \frac{1}{2}}{\alpha_n}, \quad \forall n \geq p + 1. \quad (34)$$

Of course, in practice, both $\widehat{\sigma}_{p,n}^2 \neq \sigma_{p,n}^2$ and the occurrence of an overload event at some point are with probability 1, and in any case δ_α is obviously finite. Nevertheless, Proposition 2 implies that even under the best theoretical conditions, for a particular set of system parameters, the highest resolution is limited. This motivates us to identify the point in time at which the system has reached its limiting capability, and stop the resolution adaptation—which yields stationarity henceforth—favoring stability and reducing the computational load. Clearly, the optimal scenario is when the adaptation-free mod-ADC is working constantly at the highest attainable resolution α_∞ .

In view of (32) of Proposition 2, we propose the following

Steady State Detector

$$\widehat{\mathbb{I}}_{M_\infty, n} \triangleq \begin{cases} 1, & \widehat{\sigma}_{p,n} > \frac{M_n}{2\kappa}, \\ 0, & \text{otherwise} \end{cases}, \quad \forall n \geq N_s, \quad (35)$$

where

$$\widehat{\sigma}_{p,n}^2 \triangleq \frac{1}{n} \sum_{k=1}^n \left(\widehat{v}_k - \widehat{v}_k^p \right)^2, \quad \forall n \in \mathbb{N}^+. \quad (36)$$

In words, when the estimated RMSE of the linear predictor \widehat{v}_n^p is κ times greater than half the effective modulo range, we infer that the mod-ADC has reached the limit of its capability, in terms of the highest attainable resolution for the given set of system parameters. Note the difference between (23) and (35), where the former uses $\widehat{\sigma}_{p,n}$ and the latter uses $\widehat{\sigma}_{p,n}$, respectively. As seen from its definition (22), $\widehat{\sigma}_{p,n}$ is a “short-term” memory estimate of the “local” standard deviation of v_n . In contrast, as seen from (36), $\widehat{\sigma}_{p,n}$ is a “long-term” memory estimate of the average standard deviation of \bar{v}_n , which converges to $\bar{\sigma}_{p,\infty}$ in the absence of an overload,

$$\mathbb{E} \left[\widehat{\sigma}_{p,n}^2 \middle| \mathcal{E}_{\text{OL}n+1}^{(n)} \right] \xrightarrow{n \rightarrow \infty} \bar{\sigma}_{p,\infty}^2,$$

as shown in Appendix D, with further discussion on (36).

Hence, as the adaptive process unfolds, the condition $\widehat{\sigma}_{p,n} > \frac{M_n}{2\kappa}$, which is a practical proxy for the ideal (merely theoretical) condition $\widehat{\sigma}_{p,n} = \frac{M_n}{2\kappa}$, and is essentially a decision rule for detecting the limit resolution α_∞ , becomes increasingly accurate. Although the ideal conditions hold only approximately in practice, as we show in Section V via simulations, the steady state detector (35) works quite well and is fairly accurate.

It is also instructive to express the asymptotic resolution α_∞ , via (32) and $\Delta = 2^R$, as,

$$\alpha_\infty = \frac{1}{\kappa} \cdot \left(\frac{1}{2}2^R\right) \cdot \frac{1}{\bar{\sigma}_{p,\infty}}, \quad (37)$$

which provides several insightful observations. First, and most obviously, increasing the number of bits R increases the asymptotic resolution. Second, the trade-off in choosing the confidence level parameter κ is now apparent. Indeed, increasing κ leads to an exponential decrease in overload probability (21), but at the same time decreases the asymptotic resolution (37). Third, the inverse RMSE $1/\bar{\sigma}_{p,\infty}$ reflects the *unknown* causal and linear predictability accuracy. That is, how accurately the current sample of \bar{v}_n can be predicted, using a linear causal predictor, based on the previous p samples \bar{v}_n . The lower $\bar{\sigma}_{p,\infty}$, the higher the predictability, and accordingly, the higher the asymptotic resolution α_∞ .

Interestingly, (37) also provides a fresh look at the result [3, Eq. 6]. Indeed, if we assume that the filter learned by the LMS asymptotically converged *exactly* to the optimal one, implying that $\sigma_{\text{LMMSE},p} = \alpha_\infty \cdot \bar{\sigma}_{p,\infty}$, then combined with (37) written as $M_\infty = 2\kappa\bar{\sigma}_{p,\infty}$, the bound [3, Eq. 6] reads,

$$\begin{aligned} \Pr(\mathcal{E}_{\text{OL}n}^*) &\leq 2 \exp\left\{-\frac{3}{2}2^{2(R-\frac{1}{2}\log_2(12\sigma_{\text{LMMSE},p}^2))}\right\} \\ &= 2 \exp\left\{-\frac{3}{2}2^{2\log_2\left(\frac{M_\infty}{\sqrt{12}\bar{\sigma}_{p,\infty}}\right)}\right\} \\ &= 2 \exp\left\{-\frac{3}{2}2^{2\log_2\left(\frac{\kappa}{\sqrt{3}}\right)}\right\} = 2 \exp\left\{-\frac{\kappa^2}{2}\right\}. \end{aligned} \quad (38)$$

That is, for this ideal case, the blind mod-ADC coincides with the oracle mod-ADC, and the overload probability decreases exponentially with κ^2 . This is in perfect agreement with (21).

Yet another way to see the consistency of the blind mod-ADC asymptotic performance with the that of the informed mod-ADC is via the asymptotic rate. By isolating R in (37), and using $\sigma_{\text{LMMSE},p}^2 = \alpha_\infty^2 \cdot \bar{\sigma}_{p,\infty}^2$ as above, we have

$$\begin{aligned} R &= \frac{1}{2} \log_2\left(\frac{\bar{\sigma}_{p,\infty}^2}{\frac{1}{12\alpha_\infty^2}}\right) + \log_2\left(\frac{\kappa}{\sqrt{3}}\right) \\ &\triangleq \frac{1}{2} \log_2(12\sigma_{\text{LMMSE},p}^2) + \delta_\kappa, \end{aligned} \quad (39)$$

which, again, is in perfect compliance with [3], (12) therein, such that κ controls the overload probability, and inevitably the excess rate δ_κ with respect to Shannon's lower bound [18].

Having described in detail all the individual components, the complete algorithm of the blind mod-ADC is given in Algorithm 3. An important observation is that our algorithm can work *without* incorporating the steady state detector. Thus, (35) is not a necessary component in order to ensure proper operation of the blind mod-ADC. Moreover, when working in highly dynamic environments, we might intentionally disable this detector, allowing the LMS to constantly adapt the filter \hat{h}_n according to the input, which can possibly be only "quasi-stationary" with

SOSs that change sufficiently slow (relative to the sampling period T_s) over time.

V. SIMULATION RESULTS

In this section, we present empirical results of two simulation experiments that demonstrate the successful operation of the proposed algorithm. These results corroborate our analytical derivations, and to the best of our knowledge, serve as the first empirical evidence for the implementation feasibility of a blind mod-ADC for scalar time series input signals. Additional illustrations for these simulations are given in Appendix E.

A. A Gaussian Input Signal

We consider the case where the input signal x_n is Gaussian. This is quite a common assumption; for example, digital communication signals are commonly modeled as Gaussian, see, e.g., [26]. Specifically, we generate the input as

$$x_n = \frac{1}{\sqrt{L_x}} \sum_{\ell=0}^{L_x-1} \xi_{n-\ell},$$

where $\{\xi_n\}$ is a zero-mean unit-variance Gaussian i.i.d. process. Accordingly, the autocorrelation function of x_n , assumed to be unknown, is given by $R_x[\ell] = (1 - \frac{|\ell|}{L_x}) \cdot \mathbb{1}_{|\ell| < L_x}$.

We simulate a quantizer with $R = 10$ bits, and generate the signal v_n according to (8). We then apply the 2^R -modulo operator to v_n , which yields the simulated mod-ADC output process y_n , as in (7). The chosen set of required system parameters, prescribed in the input to Algorithm 3, is given in Table II. Further, we set $L_x = 15$, and emphasize that we intentionally choose $p \neq L_x$, and specifically $p > L_x$. This simulates the more probable scenario, in which the support of $R_x[\ell]$ is unknown, hence the length of the FIR filter \hat{h}_n will not be matched to the length of the optimal LMMSE causal filter, which is of length $L_x - 1$ in this case. We also choose $L_s = p$ to demonstrate that calibration of the system parameters could be simple, and rather straightforward. Note also that we choose $\kappa = 4.5$, which gives $\Pr(|e_n^p| > \frac{1}{2}2^R | \mathcal{E}_{\text{OL}n}^{(p)}) = 2Q(4.5) \approx 6.7953 \times 10^{-6}$.

Fig. 4 presents the effective modulo range M_n (10) vs. discrete-time for an input of length $N = 10^4$ samples. Starting from $M_0 = 10.24$, more than 10 times the standard deviation of x_n , effectively guarantees that $v_n = \tilde{y}_n$, i.e., no folding during the initialization phase. This enables the LMS algorithm to learn the optimal filter. When $C_\alpha > L_s$, a resolution update is allowed, and α_n is increased, as can be seen conveniently in superimposed "close-up" of the convergence interval.

The adaptive process continues with updates of multiplicative step-sizes δ_α , and whenever required, is also decreased. Furthermore, it is seen that at some point, the steady state detector $\hat{\mathbb{1}}_{M_\infty,n}$ is turned on, indicating that the asymptotic resolution has been approximately attained. Indeed, the convergence is not exactly to M_∞ (equivalently to α_∞), since in practice $\widehat{\sigma}_{p,n}^2 \neq \sigma_{p,n}^2$ with probability 1 and the adaptations of α_n are of finite resolution. Nonetheless, as evident from Fig. 4, our asymptotic analysis, carried out under ideal theoretical conditions, provides a considerably accurate forecast.

We repeat the experiment in exactly the same setting, only now with $\kappa = 3.5$. Hence, we expect to observe an increased asymptotic resolution, at the cost of more frequent overload events. Fig. 5, presenting the corresponding plot as in Fig. 4, reflects exactly this trend. Nevertheless, in the case of an overload event, our error propagation defense mechanism comes into play, and

Algorithm 3: Blind Modulo ADC Encoding-Decoding.

Input: Signal: $\{x(t)|_{t=nT_s}\}$, System Parameters: $R, \alpha_0, p, \hat{\mathbf{h}}_0, \kappa, L_s, N_s, \epsilon_\mu, \delta_\alpha, \beta$

Output: $\{\hat{x}_n\}$

- 1 Set $\alpha_n = \alpha_0$, $C_\alpha = 0$, and $\hat{\mathbf{h}}_{p+1} = \hat{\mathbf{h}}_0$, and accumulate $p + 1$ sample $\{y_i\}_{i=1}^{p+1}$; ▷ Set initial parameters
- 2 Compute $\{\tilde{y}_i = v_i\}_{i=1}^p$ as in (11), $\bar{\mathbf{v}}_{p+1}$ as in (12), and set $\hat{\mathbf{v}}_{p+1} = \bar{\mathbf{v}}_{p+1}$; ▷ First p samples are not folded
- 3 Compute $\hat{\sigma}_{\bar{\mathbf{v}}, p+1}^2$ as in (29), and set $\mu = \epsilon_\mu / (p \cdot \hat{\sigma}_{\bar{\mathbf{v}}, p+1}^2)$ as in Subsection IV-D; ▷ Compute the LMS learning rate
- 4 Set $C_\alpha = 0$ and $\mathbb{1}_{M_\infty} = 0$; ▷ C_α : # iterations after adapting α_n , $\mathbb{1}_{M_\infty}$: M_∞ flag

for $n = p + 1, p + 2, \dots$ **do**

- 5 $\hat{v}_n, \hat{v}_n^p = \text{BlindModUnfold}(y_n, \hat{\mathbf{v}}_n, \hat{\mathbf{h}}_n, R)$;
- 6 Update $\hat{\sigma}_{\bar{\mathbf{v}}, n}^2 = \frac{n-2}{n-1} \cdot \hat{\sigma}_{\bar{\mathbf{v}}, n-1}^2 + \frac{1}{n-1} \cdot \hat{v}_{n-1}^2$;
- if** $\hat{\mathbb{1}}_{\mathcal{E}_{OL_n}} = 1$ **then**
 - 7 Reset $\alpha_{n+1} = \alpha_0$, $C_\alpha = 0$ and $\hat{\mathbf{h}}_{n+1} = \frac{\alpha_0}{\alpha_n} \cdot \hat{\mathbf{h}}_n$; ▷ Re-open modulo range
 - 8 Accumulate $(p + 1)$ new samples $\{y_i\}_{i=n+1}^{n+p+1}$, compute $\{\tilde{y}_i = v_i\}_{i=n+1}^{n+p}$; ▷ Re-initialization
 - 9 Output the respective estimates $\{\hat{x}_i = (v_i + \frac{1}{2}) / \alpha_0\}_{i=n+1}^{n+p}$ of the accumulated samples;
 - 10 Continue from $n = n + p + 1$;
- else**
 - 11 Output $\hat{x}_n = (\hat{v}_n + \frac{1}{2}) / \alpha_n$;
 - 12 Compute the estimated error $\hat{e}_n^p = \hat{v}_n - \hat{v}_n^p$, and the estimated MSE $\hat{\sigma}_{p,n}^2$ as in (22);
 - 13 Update $\hat{\mathbf{h}}_{n+1} = \hat{\mathbf{h}}_n + \mu \cdot \hat{\mathbf{v}}_n \hat{e}_n^p$, and increase $C_\alpha = C_\alpha + 1$; ▷ LMS learning step
 - if** $C_\alpha > L_s \cap \mathbb{1}_{M_\infty} = 0$ **then**
 - 14 $\alpha_{n+1} = \left[\mathbb{1}_{\uparrow \alpha, n} \cdot \frac{1}{\delta_\alpha} + (\mathbb{1}_{\uparrow \alpha, n} - 1) \cdot \delta_\alpha \right] \alpha_n$, $C_\alpha = 0$; ▷ Update the resolution
 - 15 $\hat{\mathbf{h}}_{n+1} = \left[\mathbb{1}_{\uparrow \alpha, n} \cdot \frac{1}{\delta_\alpha} + (\mathbb{1}_{\uparrow \alpha, n} - 1) \cdot \delta_\alpha \right] \hat{\mathbf{h}}_n$; ▷ Adapt the filter accordingly
 - 16 **if** $\hat{\mathbb{1}}_{M_\infty, n} = 1 \cap \mathbb{1}_{\uparrow \alpha, n} = 1$ **then** $\mathbb{1}_{M_\infty} = 1$; ▷ Steady state detection;

TABLE II
CHOSEN SYSTEM PARAMETERS FOR SIMULATION EXPERIMENT 1

| Parameter | Value | Parameter | Value |
|---|-----------------------------|-----------------|-----------|
| R ($\rightarrow \Delta = 2^R$) | 10 ($\rightarrow 1024$) | L_s | $p = 40$ |
| α_0 ($\rightarrow M_0 = 2^R / \alpha_0$) | 100 ($\rightarrow 10.24$) | N_s | 500 |
| p | 40 | ϵ_μ | 10^{-2} |
| $\hat{\mathbf{h}}_0 \in \mathbb{R}^{p \times 1}$ | $[1 \ 0 \ \dots \ 0]^T$ | δ_α | 0.9 |
| κ | 4.5 | β | 5 |

maintains proper continuous operation, as observed from Fig. 5. Thus, the blind mod-ADC automatically balances the trade-off between effective quantization and continuous operation, which is highly important in practice. These results corroborate our analysis leading to the detector (28).

B. A Bandlimited Input With Narrowband Interferers

In this experiment we consider a non-Gaussian signal of interest (SOI), and the presence of narrowband interferences. Specifically, here the SOI \tilde{x}_n is generated by applying a non-ideal, minimum-order filter with a stopband attenuation of 60 dB, to the driving noise $\{\xi_n\}$, which is drawn from the Rademacher distribution, namely $\Pr(\xi_n = 1) = \Pr(\xi_n = -1) = \frac{1}{2}$. We then normalize the output, such that \tilde{x}_n is a zero-mean unit-variance process. In addition, we consider the presence of narrowband interference signals, a particularly relevant scenario in the context of communication systems.

Thus, the input to the blind mod-ADC in this experiment, which is of length $N = 10^5$ samples, is given by

$$x_n = \tilde{x}_n + \sum_{i=1}^3 \underbrace{g_i \cdot \sin(\phi_i + \omega_i n) \cdot \mathbb{1}_{n > \tau_i}}_{\triangleq \zeta_n^i} = \tilde{x}_n + \sum_{i=1}^3 \zeta_n^i, \quad (40)$$




Fig. 4. Experiment 1: The effective modulo range M_n vs. discrete-time, with $\kappa = 4.5$. Since on this time interval an overload did not occur, the resolution α_n approximately converged to the asymptotic resolution α_∞ , as predicted by our analytical asymptotic analysis. Upon convergence, stability is also evident, and this is due to the successful operation of the steady state detector (35).

where \tilde{x}_n simulates the SOI, and $\{\zeta_n^i\}_{i=1}^3$ simulate three narrowband interference signals. Here, for the interference ζ_i , the parameters g_i, ϕ_i, ω_i and τ_i are the unknown gain, phase, carrier (angular) frequency and transmission start time, respectively. We draw $\phi_i \sim \text{Unif}((0, 2\pi))$ independently, and set the rest of the parameters as reported in Table III.¹⁰ We also fix $\hat{\mathbb{1}}_{M_\infty, n}$ to zero, to demonstrate the unecessity of the steady state detector for a successful operation.

¹⁰Note that ω_3 is not a rational multiplication of π , hence ζ_n^3 is not a periodic signal.




Fig. 5. Experiment 1: M_n vs. discrete-time, but with $\kappa = 3.5$ (instead of $\kappa = 4.5$). For a lower value of κ , overload events are more frequent. Nonetheless, our algorithm successfully detects these events, and automatically decreases the resolution to maintain proper continuous operation of the blind mod-ADC.

TABLE III
CHOSEN INTERFERENCE PARAMETERS FOR SIMULATION EXPERIMENT 2

| Parameter | $i = 1$ | $i = 2$ | $i = 3$ |
|------------|-----------------|-----------------|-----------------|
| g_i | 2 | 2 | 2 |
| ω_i | $\pi/4$ | $4\pi/5$ | $\sqrt{2}\pi/3$ |
| τ_i | 2×10^4 | 4×10^4 | 7×10^4 |

The spectrogram of the input x_n is presented in Fig. 2(a). During the starting period ($n < 2 \times 10^4$), the blind mod-ADC operates in the presence of the SOI \tilde{x}_n only. Once $n > \tau_1$, the first interference is added, and the scenario becomes harder after τ_2 and τ_3 , where the SOI is impaired by the interferers.

Fig. 6 presents the evolution of the effective modulo range M_n in (discrete) time. Whenever an interference starts transmitting, the SOSs of the input signal abruptly changes, and as a result the prediction error e_n^p increases dramatically, thus causing an overload event. We recall that an overload, by definition (see (19)), *does not* mean that the amplitude of the input signal exceeds the modulo range $\Delta = 2^R$, as in a standard ADC. Rather, an overload event occurs when the magnitude of the prediction error e_n^p exceeds half the modulo range. In these cases, our overload detector detects these large prediction errors, and the effective range is re-opened. This way, a stream of low-resolution, though *unfolded* samples are produced, allowing the LMS algorithm to re-learn the new (and different) optimal filter. This process happens right after τ_1 , τ_2 and τ_3 , but operation in high-resolution is gained anew.

We also report the average number of failures in perfectly recovering v_n , $\widehat{\Pr}(v_n \neq \hat{v}_n) \triangleq \frac{1}{N} \sum_{\ell=1}^N \mathbb{1}_{v_n \neq \hat{v}_n} = 7.8 \cdot 10^{-4}$. Evidently, the blind mod-ADC provides highly reliable recovery of the unfolded signal, and in turn, allows for highly accurate estimation of the input x_n via (9). In this specific scenario, since the input (40) is perfectly recovered (almost everywhere), the narrowband interferers can be easily detected and filtered out (e.g., using notch filters). This digital solution, which is now simple thanks to the modulo-ADC, could not have been achieved by a standard, probably saturated ADC.

The results reported in this Section were verified by multiple runs, and were consistently observed for different realizations.




Fig. 6. The effective modulo range M_n vs. discrete-time. Our adaptive algorithm quickly recovers from the abrupt changes due to the interferences, re-learns the suitable filter, and returns to a high-resolution operational mode.

VI. CONCLUDING REMARKS

In the context of analog-to-digital conversion, we present an algorithmic framework, allowing for a stable and reliable operation of a mod-ADC without access to prior knowledge of the input signal's SOSs. We put forth the key design parameters, and discuss the corresponding trade-offs. In addition, we derive the asymptotic resolution of the proposed blind mod-ADC, and link our current results with the performance of the oracle mod-ADC presented in [3]. We demonstrate by simulations the successful operation of the proposed solution, which corroborates its underlying theoretical infrastructure. Furthermore, we demonstrate the advantage in using a mod-ADC in an environment of multiple interference signals.

As ADCs are being used in a host of applications, more often than not when perfect knowledge on the input signal is not available (if at all), the robustness of such devices is imperative, and almost crucial. The ability of operating blindly under dynamic conditions is essential for practical purposes, and constitutes a key advantage in effective sensing. Therefore, this work is yet another important step towards realization of mod-ADCs, shrinking the gap between sensing performance in practice and the respective theoretical limits.

Based on the principles presented in this work, in [27] our results are extended to the spatiotemporal case, where blind mod-ADCs for vector processes are developed and analyzed. The results in [27] can serve as a starting point for developing additional mod-ADC-based architectures with accompanying algorithms for various applications in array processing.

There are a number of important research directions associated with this work that remain to be explored. One that is central relates to the hardware implementation of the basic modulo ADC unit. Based on a given implementation design, addressing the nonidealities and inherent physical limitations/impairments via efficient (digital) algorithmic solutions is key to the success of using the modulo concept in sensing.

Another important issue pertains to the blind mod-ADC's reaction upon detection of an overload event. As observed from Fig. 2(b), in these cases the instantaneous squared error increases significantly. While standard ADCs with a (causal) AGC generally also suffer from such spikes in their reconstruction

error process, it is possible that this phenomenon can be further mitigated (for example, with additional resources).

Other, more theoretical avenues related to the proposed algorithm are a non-asymptotic performance analysis of the overload detector (28), and of the typical convergence time to the asymptotic resolution (37). Such analyses could yield valuable characterization of the resolution adaptation evolution process in terms of key system parameters, which is instructive for a comprehensive evaluation of the mod-ADC architecture.

APPENDIX A

SOSs OF THE PROCESSES v_n AND \bar{v}_n

As explained in Subsection IV-A, when α_n increases, the SOSs of $\{\bar{v}_n\}$ gradually become less affected by α_n , which is not true for $\{v_n\}$. To see this more clearly, observe that the autocorrelation function of $\{\bar{v}_n\}$ is given by

$$R_{\bar{v}}[\ell] \triangleq \mathbb{E}[\bar{v}_n \bar{v}_{n-\ell}] = R_x[\ell] + \mathbb{1}_{\ell=0} \cdot \frac{1}{12\alpha_n^2},$$

where we have used (i) the definition (12); (ii) the fact that x_n and z_n are statistically independent; and (iii) z_n is an i.i.d. process. In contrast, the autocovariance of $\{v_n\}$ is given by

$$\begin{aligned} R_v[\ell] &\triangleq \mathbb{E}\left[\left(v_n - \frac{1}{2}\right)\left(v_{n-\ell} - \frac{1}{2}\right)\right] \\ &= \alpha_n \alpha_{n-\ell} R_x[\ell] + \mathbb{1}_{\ell=0} \cdot \frac{1}{12}, \end{aligned}$$

and we examine the autocovariance (rather than the autocorrelation) since $\mathbb{E}[v_n] = \frac{1}{2}$. Evidently, for any non-zero lag $\ell \neq 0$, $R_{\bar{v}}[\ell] = R_x[\ell]$, and in particular $R_{\bar{v}}[\ell]$ is independent of α_n for $\ell \in \{1, \dots, p\}$. Clearly, this is not the case for $R_v[\ell]$. Moreover, the variance $R_{\bar{v}}[0]$, also given in (13), approaches σ_x^2 as α_n increases. Since $R_{\bar{v}}[\ell]$ is less sensitive than $R_v[\ell]$ to adaptations in α_n , by using the recovered values of $\{\bar{v}_n\}$ rather than $\{v_n\}$ as the observations in (14) (as opposed to (3)), we alleviate the estimation (/learning) of the optimal filter coefficients, which depend on the SOSs of the observations, throughout the adaptive process.

APPENDIX B

PROOF OF PROPOSITION 1

Proof: It is known that [28], if

$$R_x[\ell] \cdot \log(\ell) \xrightarrow{\ell \rightarrow \infty} 0,$$

then

$$\Pr\left(\|x_n\|_{\ell_N^\infty} \leq \sqrt{2\sigma_x^2 \log(N)}\right) \xrightarrow{N \rightarrow \infty} 1.$$

Since $\{z_n \in (-1, 0]\}$ is a process with bounded support, assuming that $\min_n \{\alpha_n\} = \alpha_0$, we also have

$$\Pr\left(\|\bar{v}_n\|_{\ell_N^\infty} \leq \sqrt{2\left(\sigma_x^2 + \frac{1}{12\alpha_0^2}\right) \log(N)}\right) \xrightarrow{N \rightarrow \infty} 1.$$

Recall, however, that $\{\alpha_n\}$ is typically an increasing sequence on average, and conditioned on no overload events $\mathcal{E}_{\text{OL}_{n+1}}^{(n)}$, converges (up to small fluctuations) to the value for which $\kappa \cdot \sigma_{p,n} = \frac{\Delta}{2}$, as explained in Subsection IV-B. Thus, under $\bigcap_{k=1}^n \{\mathbb{1}_{\mathcal{E}_{\text{OL}_k}} = 0\}$, in the absence of estimation errors in $\hat{\sigma}_{p,n}$ and with an infinite resolution step size $\delta_\alpha \rightarrow 1$,

$$\lim_{n \rightarrow \infty} \alpha_n = \alpha_\infty \Rightarrow \lim_{n \rightarrow \infty} \mathbb{E}[\bar{v}_n^2] = \sigma_x^2 + \frac{1}{12\alpha_\infty^2} = \sigma_{\bar{v}}^2.$$

Hence, for the ideal steady state process \bar{v}_n^∞ , we have

$$\Pr\left(\|\bar{v}_n^\infty\|_{\ell_N^\infty} \leq \sqrt{2\sigma_{\bar{v}}^2 \log(N)}\right) \xrightarrow{N \rightarrow \infty} 1,$$

or, equivalently,

$$\Pr\left(\|\bar{v}_n^\infty\|_{\ell_N^\infty} > \sqrt{2\sigma_{\bar{v}}^2 \log(N)}\right) \xrightarrow{N \rightarrow \infty} 0. \quad \blacksquare$$

APPENDIX C

PROOF OF PROPOSITION 2

Proof: In the asymptotic state under the assumptions stated in the proposition, we have the equilibrium

$$\kappa \cdot \tilde{\sigma}_{p,\infty} = \frac{\Delta}{2}, \quad (41)$$

where

$$\tilde{\sigma}_{p,\infty}^2 \triangleq \lim_{n \rightarrow \infty} \tilde{\sigma}_{p,n}^2, \quad \tilde{\sigma}_{p,n}^2 \triangleq \mathbb{E}\left[(e_n^p)^2 \middle| \mathcal{E}_{\text{OL}_{n+1}}^{(n)}\right], \quad (42)$$

and notice the difference between $\tilde{\sigma}_{p,n}^2$ in (42) and $\sigma_{p,n}^2$ in (20). Now, recall that \hat{v}_n^p is a function of \bar{v}_n (14), which, asymptotically, is a function of α_∞ . Hence $\tilde{\sigma}_{p,\infty}^2$ is also a function of α_∞ . Thus, under the ideal (only theoretical) conditions stated in the proposition, the highest resolution attainable for a particular fixed set of system parameters (e.g., $\kappa, p, \Delta = 2^R$) is governed by the equilibrium (41).

Now, observe that the linear predictor of the “ α_n -standardized” process \bar{v}_n , defined in (34), has the following conditional MSE,

$$\bar{\sigma}_{p,n}^2 = \mathbb{E}\left[\left(\bar{v}_n - \hat{v}_n^p\right)^2 \middle| \mathcal{E}_{\text{OL}_{n+1}}^{(n)}\right] = \frac{\tilde{\sigma}_{p,n}^2}{\alpha_n^2}.$$

Therefore, under the ideal (theoretical) conditions stated in the proposition, asymptotically,

$$\begin{aligned} \bar{\sigma}_{p,\infty}^2 &\triangleq \lim_{n \rightarrow \infty} \bar{\sigma}_{p,n}^2 = \lim_{n \rightarrow \infty} \frac{\tilde{\sigma}_{p,n}^2}{\alpha_n^2} = \frac{\tilde{\sigma}_{p,\infty}^2}{\alpha_\infty^2} \\ &\implies \tilde{\sigma}_{p,\infty} = \alpha_\infty \cdot \bar{\sigma}_{p,\infty}. \end{aligned} \quad (43)$$

Substituting (43) into (41) gives the equivalent equilibrium

$$\bar{\sigma}_{p,\infty} = \frac{1}{2\kappa} \cdot \frac{\Delta}{\alpha_\infty} \triangleq \frac{M_\infty}{2\kappa}. \quad \blacksquare$$

APPENDIX D

THE ASYMPTOTIC RMSE ESTIMATOR

Further analytical justification of (35) is gained by

$$\begin{aligned} \mathbb{E}\left[\hat{\sigma}_{p,n}^2 \middle| \mathcal{E}_{\text{OL}_{n+1}}^{(n)}\right] &\stackrel{\textcircled{1}}{=} \frac{1}{n} \sum_{k=1}^n \mathbb{E}\left[\left(\hat{v}_k - \hat{v}_k^p\right)^2 \middle| \mathcal{E}_{\text{OL}_{n+1}}^{(n)}\right] \\ &\stackrel{\textcircled{2}}{=} \frac{1}{n} \sum_{k=1}^n \mathbb{E}\left[\left(\bar{v}_k - \hat{v}_k^p\right)^2 \middle| \mathcal{E}_{\text{OL}_{n+1}}^{(n)}\right] \\ &\stackrel{\textcircled{3}}{=} \frac{1}{n} \sum_{k=1}^n \frac{1}{\alpha_k^2} \cdot \mathbb{E}\left[(e_k^p)^2 \middle| \mathcal{E}_{\text{OL}_{n+1}}^{(n)}\right] \\ &\stackrel{\textcircled{4}}{\xrightarrow{n \rightarrow \infty}} \frac{\tilde{\sigma}_{p,\infty}^2}{\alpha_\infty^2} = \bar{\sigma}_{p,\infty}^2, \end{aligned}$$

where we have used:

- ① Linearity of the expectation;
- ② Under $\mathcal{E}_{\text{OL}_k}$, $\hat{v}_k = v_k \implies \hat{v}_k = \bar{v}_k$;




Fig. 7. Experiment 1 with $\kappa = 4.5$: The estimated error process \hat{e}_n^p vs. discrete-time. It is seen that $|\hat{e}_n^p| < \frac{1}{2}2^R$ for the entire time interval. Since here $\hat{e}_n^p = e_n^p$, an overload event \mathcal{E}_{OL_n} did not occur, and perfect recovery of v_n is attained.

- ③ By definition (12), (34), (17), $\bar{v}_k - \hat{v}_k^p = e_k^p/\alpha_k$; and
- ④ Under the same ideal conditions described at the outset of Proposition 2, $\exists N_{\alpha_\infty} \in \mathbb{N}^+ : \alpha_n = \alpha_\infty, \forall n > N_{\alpha_\infty}$; See the discussion below regarding this analysis.

Although $\hat{\sigma}_{p,n}^2$, defined in (36), is perhaps the most intuitive estimator of $\sigma_{p,n}^2$, an exact analysis of its asymptotic properties is far from trivial. Indeed, since it is a random process whose statistical properties are implicitly determined by the resolution update, error propagation prevention and steady state detection rules (steps 14, 7 and 16 in Algorithm 3, respectively), it is even non-stationary to begin with.

However, to further justify our proposed steady state detector (35), which is based on $\hat{\sigma}_{p,n}^2$ (36), it suffices to consider a simplified scenario, in which an overload event never occurs. Of course, this happens with probability zero when considering an infinitely long observation of the error process $\{e_n^p\}$, since we assume $e_n^p | \mathcal{E}_{OL_n}^{(p)} \sim \mathcal{N}(0, \sigma_{p,n}^2)$. Nevertheless, such an analysis is informative for finite, but sufficiently long realizations, in which our proposed adaptive mechanism for the blind mod-ADC converges to steady state, i.e., $\hat{\mathbb{1}}_{M_\infty, n} = 1$, which occurs w.h.p. with proper selection of the system parameters.

The simulation results presented in Section V corroborate this argument, and further justifies this approach for a simplified, yet informative analysis, resulting in (37), which is consistent with the analysis of the informed mod-ADC presented in [3], as evident from (38) and (39).

APPENDIX E

SIMULATION RESULTS: ADDITIONAL ILLUSTRATIONS

We provide additional illustrations of the results of the two simulation experiments presented in Section V. The *estimated* prediction errors \hat{e}_n^p of the linear predictor \hat{v}_n^p are presented in Fig. 7. Recall that in order to unfold \tilde{y}_n , these prediction errors (15) are necessary, and $\hat{e}_n^p = e_n^p$ holds only when there is no overload. As reflected from Fig. 7, this is exactly the case, since $\mathcal{E}_{OL_n} = \{|e_n^p| < \frac{1}{2}2^R\} = \{\hat{e}_n^p = e_n^p\}$. Accordingly, in this experiment the averaged squared error is $\frac{1}{N} \sum_{n=1}^N (v_n - \hat{v}_n)^2 \approx$




Fig. 8. Results of simulation experiment 1, but with $\kappa = 3.5$ instead of $\kappa = 4.5$: The estimated error process \hat{e}_n^p vs. discrete-time. For a lower value of κ , overload events are more frequent. Nonetheless, our algorithm successfully detects these events, and automatically lowers the resolution in order to maintain proper continuous operation of the blind mod-ADC.

5.067×10^{-27} , which is clearly due to machine accuracy limitations, thus implying perfect recovery of the signal v_n , from which x_n can be readily estimated.

For the repeated experiment with a smaller value of $\kappa = 3.5$, it is seen in Fig. 7, presenting the corresponding plot to Fig. 8 that three overload event have occurred, successfully detected, and mitigated.

ACKNOWLEDGMENT

This material is based upon work supported by the United States (US) Air Force under Air Force Contract No. FA8702-15-D-0001. Any opinions, findings, conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the US Air Force.

REFERENCES

- [1] B. Wang and K. R. Liu, "Advances in cognitive radio networks: A survey," *IEEE J. Sel. Topics Signal Process.*, vol. 5, no. 1, pp. 5–23, Feb. 2011.
- [2] D. Cabric and R. W. Brodersen, "Physical layer design issues unique to cognitive radio systems," in *Proc. IEEE 16th Int. Symp. Pers., Indoor, Mobile Radio Commun.*, vol. 2, 2005, pp. 759–763.
- [3] O. Ordentlich, G. Tabak, P. K. Hanumolu, A. C. Singer, and G. W. Wornell, "A modulo-based architecture for analog-to-digital conversion," *IEEE J. Sel. Topics Signal Process.*, vol. 12, no. 5, pp. 825–840, Oct. 2018.
- [4] D. N. Green, "Global stability analysis of automatic gain control circuits," *IEEE Trans. Circuits Syst.*, vol. CAS-30, no. 2, pp. 78–83, Feb. 1983.
- [5] F. Sun, J. Singh, and U. Madhoo, "Automatic gain control for ADC-limited communication," in *Proc. GLOBECOM*, 2010, pp. 1–5.
- [6] T. Ericson and V. Ramamoorthy, "Modulo-PCM: A new source coding scheme," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, vol. 4, 1979, pp. 419–422.
- [7] R. Zamir, Y. Kochman, and U. Erez, "Achieving the Gaussian rate-distortion function by prediction," *IEEE Trans. Inf. Theory*, vol. 54, no. 7, pp. 3354–3364, Jul. 2008.
- [8] O. Ordentlich and U. Erez, "Integer-forcing source coding," *IEEE Trans. Inf. Theory*, vol. 63, no. 2, pp. 1253–1269, Feb. 2017.
- [9] E. Romanov and O. Ordentlich, "Blind unwrapping of modulo reduced Gaussian vectors: Recovering MSBs from LSBs," *IEEE Trans. Inf. Theory*, vol. 67, no. 3, pp. 1897–1919, Mar. 2021.
- [10] A. Bhandari, F. Kraemer, and R. Raskar, "On unlimited sampling," in *Proc. Int. Conf. Sampling Theory Appl.*, 2017, pp. 31–35.
- [11] A. Bhandari, F. Kraemer, and R. Raskar, "Unlimited sampling of sparse sinusoidal mixtures," in *Proc. IEEE Int. Symp. Inf. Theory*, 2018, pp. 336–340.
- [12] A. Bhandari, F. Kraemer, and R. Raskar, "On unlimited sampling and reconstruction," *IEEE Trans. Signal Process.*, vol. 69, pp. 3827–3839, 2021.

- [13] A. Bhandari and F. Krahmer, "HDR imaging from quantization noise," in *Proc. IEEE Int. Conf. Image Process.*, 2020, pp. 101–105.
- [14] E. Romanov and O. Ordentlich, "Above the Nyquist rate, modulo folding does not hurt," *IEEE Signal Process. Lett.*, vol. 26, no. 8, pp. 1167–1171, Aug. 2019.
- [15] A. Bhandari and F. Krahmer, "On identifiability in unlimited sampling," in *Proc. Int. Conf. Sampling Theory Appl.*, 2019, pp. 1–4.
- [16] O. Graf, A. Bhandari, and F. Krahmer, "One-bit unlimited sampling," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, 2019, pp. 5102–5106.
- [17] A. Bhandari, F. Krahmer, and T. Poskitt, "Unlimited sampling from theory to practice: Fourier-prony recovery and prototype ADC," *IEEE Trans. Signal Process.*, vol. 70, pp. 1131–1141, 2021.
- [18] T. Berger, *Rate Distortion Theory: A Mathematical Basis for Data Compression*. Englewood Cliffs, NJ, USA: Prentice-Hall, 1971.
- [19] S. S. Haykin and B. Widrow, *Least-Mean-Square Adaptive Filters*, vol. 31. Hoboken, NJ, USA: Wiley Online Library, 2003.
- [20] R. M. Gray and D. L. Neuhoff, "Quantization," *IEEE Trans. Inf. Theory*, vol. 44, no. 6, pp. 2325–2383, Oct. 1998.
- [21] R. Zamir, *Lattice Coding for Signals and Networks: A Structured Coding Approach to Quantization, Modulation, and Multiuser Information Theory*. Cambridge, U.K.: Cambridge Univ. Press, 2014.
- [22] A. Feuer and E. Weinstein, "Convergence analysis of LMS filters with uncorrelated Gaussian data," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. ASSP-33, no. 1, pp. 222–230, Feb. 1985.
- [23] M. Basseville, "Detecting changes in signals and systems—a survey," *Automatica*, vol. 24, no. 3, pp. 309–326, 1988.
- [24] T. Fawcett and F. Provost, "Adaptive fraud detection," *Data Mining Knowl. Discov.*, vol. 1, no. 3, pp. 291–316, 1997.
- [25] E. Kamen and J. Su, "The Wiener filter," in *Introduction to Optimal Estimation*. London, U.K.: Springer, 1999, pp. 101–147.
- [26] P. Banelli and S. Cacopardi, "Theoretical analysis and performance of OFDM signals in nonlinear AWGN channels," *IEEE Trans. Commun.*, vol. 48, no. 3, pp. 430–441, Mar. 2000.
- [27] A. Weiss, E. Huang, O. Ordentlich, and G. W. Wornell, "Blind modulo analog-to-digital conversion of vector processes," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, 2022, pp. 5617–5621.
- [28] M. R. Leadbetter, G. Lindgren, and H. Rootzén, *Extremes and Related Properties of Random Sequences and Processes*. New York, NY, USA: Springer, 2012.



Amir Weiss (Member, IEEE) received the B.Sc., (*magna cum laude*), M.Sc., and Ph.D. degrees in electrical engineering from Tel Aviv University (TAU), Tel-Aviv, Israel, in 2013, 2015 and 2020, respectively. From 2019 to 2020, he was a Postdoctoral Fellow with the Department of Computer Science and Applied Mathematics, Weizmann Institute of Science, Rehovot, Israel, and is currently a Postdoctoral Associate with the Research Laboratory of Electronics, Massachusetts Institute of Technology, Cambridge, MA, USA. His research interests mainly include statistical

and digital signal processing, estimation theory, and machine learning. He has held a Researcher Position with Elbit Systems, EW and SIGINT Elisra Ltd., Holon, Israel, from 2013 to 2020, specializing in detection and estimation of RADAR and SONAR signals.

He was the recipient of the scholarship for excellent M.Sc. students from the Faculty of Engineering, TAU in 2015, and The David and Paulina Trotsky foundation award for outstanding Ph.D. students in 2019, and the 2021 ICASSP Outstanding Paper Award. He was also the recipient of the Nadav Levanon Studies Prize for graduate students in 2016, the scientific publication prize (three times) in 2017, 2018 and 2020, and the David Burshtein scientific publication prize in 2019, all from the The Yitzhak and Chaya Weinstein Research Institute for Signal Processing.



Everest Huang (Member, IEEE) received the S.B. degree in physics, the S.B. and M.Eng. and the Ph.D. degrees in electrical engineering from the Massachusetts Institute of Technology, Cambridge, MA, USA, in 1996, 1998, and 2006, respectively. He is currently a Technical Staff Member with Advanced Satcom Systems and Operations Group, MIT Lincoln Laboratory, Lexington, MA, USA. His research interests mainly include the development and analysis of robust wireless communications systems, with focus on the physical, and link layers.



Or Ordentlich (Member, IEEE) received the B.Sc. (*cum laude*), M.Sc. (*summa cum laude*), and the Ph.D. degrees in electrical engineering from Tel Aviv University, Tel Aviv, Israel, in 2010, 2011, and 2016, respectively. From 2015 to 2017, he was a Postdoctoral Fellow with the Laboratory for Information and Decision Systems, Massachusetts Institute of Technology, Cambridge, MA, USA, and the Department of Electrical and Computer Engineering, Boston University, Boston, MA, USA. He is currently an Associate Professor with the School of Computer Science and

Engineering, the Hebrew University of Jerusalem, Jerusalem, Israel. He is serving as an Associate Editor for the IEEE TRANSACTION ON INFORMATION THEORY since 2021.



Gregory W. Wornell (Fellow, IEEE) received the B.A.Sc. degree in electrical engineering and computer science from the University of British Columbia, Vancouver, BC, Canada, and the S.M. and Ph.D. degrees in electrical engineering and computer science from the Massachusetts Institute of Technology (MIT), Cambridge, MA, USA, in 1985, 1987 and 1991, respectively.

Since 1991, he has been on the faculty with MIT, where he is currently the Sumitomo Professor of Engineering with the Department of Electrical Engineering and Computer Science (EECS) and the Schwarzman College of Computing, and area Co-Chair of the EECS Doctoral Program. At MIT, he leads the Signals, Information, and Algorithms Laboratory, and is affiliated with the Research Laboratory of Electronics, the Computer Science and Artificial Intelligence Laboratory, and the Institute for Data, Systems and Society. He has held visiting appointments with the Department of Electrical Engineering and Computer Science, University of California, Berkeley, CA, USA, during 1999–2000, Hewlett-Packard Laboratories, Palo Alto, CA, USA, in 1999, and AT&T Bell Laboratories, Murray Hill, NJ, USA, during 1992–1993.

His research interests include signal processing, information theory, statistical inference, artificial intelligence, information security, architectures for sensing, learning, computing, communication, storage, systems for computational imaging, vision, perception, aspects of computational biology and neuroscience, and the design of wireless networks. He is involved with the Information Theory and Signal Processing Societies of the IEEE in a variety of capacities, and maintains a number of close industrial relationships and activities. He has won a number of awards for both his research and teaching, including the 2019 IEEE Leon K. Kirchmayer Graduate Teaching Award.