LOW-COMPLEXITY DIGITAL ENCODING STRATEGIES FOR WIRELESS SENSOR NETWORKS

Haralabos C. Papadopoulos, Gregory W. Wornell, and Alan V. Oppenheim

Research Laboratory of Electronics
Massachusetts Institute of Technology
Cambridge, MA 02139

ABSTRACT

Low-complexity schemes for digital encoding of a noise-corrupted signal and associated signal estimators are presented. This problem arises in wireless distributed sensor networks where an environmental signal of interest is to be estimated at a central site from low-bandwidth digitized information received from collections of remote sensors. We show that the use of a properly designed and often easily implemented additive control input before signal quantization can significantly enhance overall system performance. In particular, efficient estimators can be constructed and used with optimized pseudo-noise, deterministic, and feedback-based control inputs, resulting in a hierarchy of practical systems with very attractive performance-complexity characteristics.

1. INTRODUCTION

In a number of existing and future wireless sensor networks, sensor dynamic range and resolution are often severely limited due to either physical limitations in sensor design, or power and bandwidth constraints in the communication link back to the central site. In such cases, quantization is an integral part of the sensor model and can be viewed as a digital encoding of the environmental signal being acquired. In this paper, we develop low-complexity schemes that perform digital encoding of the environmental signal of interest at each sensor prior to transmission to the central site, and also present efficient estimators of the environmental signal from these digitized encodings. Depending upon bandwidth and power resources, the central site may or may not broadcast information back to these sensors; both scenarios will be considered.

In order to overcome the dynamic range and finite-precision constraints due to quantization—or, equivalently, obtain an effective digital encoding—we consider the use of a control input added to the information-bearing signal before signal quantization. We focus on the single-sensor case; the associated block diagram is shown in Fig. 1, where $A[n]$ denotes the information-bearing signal, $v[n]$ represents sensor noise, $w[n]$ is a control input, and $y[n]$ denotes the quantized signal based on which $A[n]$ will be estimated. Multi-sensor extensions of the single-sensor systems we present can be easily developed [2].

We focus on the static case of the estimation problem depicted in Fig. 1 where $A[n] = A$, *i.e.*, we examine the problem of estimating an unknown parameter A from quantized noisy observations.

Figure 1: Signal estimation from quantized noisy observations in the context of an additive control input.

This case reveals several key features of signal estimation based on observations from a system comprising a control input and a quantizer. Extensions to time-varying inputs are developed in [3].

Several basic variations of the estimation problem in Fig. 1 can arise in practice, which differ in the control information that is available for estimation and the associated freedom in the control input selection. For pseudo-noise control inputs whose statistical characterization alone is exploited at the receiver, we show that there is an optimal power level for minimizing the mean-square estimation error (MSE). The existence of a non-zero optimal pseudo-noise power level reveals strong connections to the phenomenon of stochastic resonance, which is encountered in a number of physical nonlinear systems [1]. Performance can be further enhanced if detailed knowledge of the control waveform is exploited at the receiver. In this scenario, we develop methods for judiciously selecting the control input from a suitable class of periodic waveforms for any given system. Finally, if feedback from the quantized output to the control input is available, we show that, when combined with suitably designed receivers, these signal quantizers come within a small loss of the quantizer-free performance.

2. SYSTEM MODEL

As outlined in Sec. 1, we consider the problem of estimating an unknown parameter A from observation of

$$y[n] = F(A + v[n] + w[n]) \quad n = 1, 2, \cdots, N, \quad (1)$$

where the sensor noise $v[n]$ is an independent identically-distributed (IID) process, $w[n]$ is a control input, and the function $F(\cdot)$ is an M-level quantizer, with the quantized output $y[n]$ taking M distinct values $Y_1, Y_2, \cdots, Y_M$, *i.e.*,

$$F(x) = \begin{cases} Y_i & \text{if } X_{i-1} \le x < X_i, \text{ for } 2 \le i \le M \\ Y_1 & \text{otherwise} \end{cases} \quad (2)$$

where $X_0 = -\infty$ and $X_M = \infty$. We also define the sequence

$$s[n] \triangleq A + v[n] = A + \sigma_v \, \tilde{v}[n] \,. \quad (3)$$

We will often be interested in system performance for a family of sensor noises parameterized by σ_v in (3), arising from scaling an IID sequence $\tilde{v}[n]$. We use p_z $(\cdot)$ to denote the probability density function (PDF) of any sample of an IID sequence $z[n]$, and C_z $(\cdot)$ to denote one minus the associated cumulative distribution, *i.e.*,

$$C_z\left(x\right) = \int_x^\infty p_z\left(t\right) dt .$$

We refer to an IID process as *admissible* if the associated PDF is non-zero and smooth (*i.e.*, C^1) almost everywhere. Throughout the paper, we assume that all noise processes are admissible, including $v[n]$ and also $w[n]$, when $w[n]$ is viewed as a pseudo-noise process. Also, when referring to a Gaussian process we assume it is IID and zero-mean.

3. PERFORMANCE LIMITS

In this section we quantify the performance degradation that results from estimating A based on observation of $y[n]$ instead of $s[n]$. We first introduce the concept of *information loss*, which we use as a figure of merit to design quantizer systems and evaluate the associated estimators. We then present a brief preview of performance limits based on this notion for a number of important scenarios and finally consider these performance limits in detail.

We define the information loss for a quantizer system as the ratio of the Cramér-Rao bounds for unbiased estimates of the parameter A obtained via $y[n]$ and $s[n]$, respectively, *i.e.*,

$$\mathcal{L}(A) \triangleq \frac{\mathcal{B}\left(A;\ \mathbf{y}^N\right)}{\mathcal{B}\left(A;\ \mathbf{s}^N\right)} , \tag{4}$$

where $\mathcal{B}\left(A;\ \mathbf{y}^N\right)$ is the Cramér-Rao bound for estimating A from

$$\mathbf{y}^N \triangleq \begin{bmatrix} y[1] & y[2] & \cdots & y[N] \end{bmatrix}^T , \tag{5}$$

and where $\mathcal{B}\left(A;\ \mathbf{s}^N\right)$ and $\mathbf{s}^N$ are defined similarly. When viewed in dB the information loss represents the additional MSE in dB that arises from observing $y[n]$ instead of $s[n]$ in the context of efficient estimation of A. From this perspective, better systems achieve smaller information loss over the parameter range of interest.

Taking into account the inherent dynamic range limitations of these quantizers, we assume that the unknown parameter takes values in the range $(-\Delta,\ \Delta)$, with Δ assumed to be known.

Worst-case performance is used to characterize the overall system; we define the worst-case Cramér-Rao bound and worst-case information loss via

$$\mathcal{B}_{\max}\left(\Delta\right) \triangleq \sup_{|A|<\Delta} \mathcal{B}\left(A;\ \mathbf{y}^N\right) , \tag{6}$$

and

$$\mathcal{L}_{\max}(\Delta) \triangleq \sup_{|A|<\Delta} \mathcal{L}(A) , \tag{7}$$

respectively. Both $\mathcal{B}_{\max}$ and $\mathcal{L}_{\max}$ are functions of other system parameters, such as σ_v and $F(\cdot)$, the dependence on which is suppressed for convenience in the above definitions.

As a consequence of the linear model (3) we obtain

$$\mathcal{B}\left(A;\ \mathbf{s}^N\right) = \sigma_v^2\, \mathcal{B}\left(0;\ \tilde{s}\right)/N , \tag{8}$$

Control Input	Order of information loss growth	
	Gaussian case	General case
Control-free case	$e^{\chi^2/2}$	$> \chi^2$
Pseudo-noise	χ^2	χ^2
Known input	χ	χ
Feedback-based input	1	1

Table 1: Order of growth of information loss as a function of $\chi = \Delta/\sigma_v$ for large χ and for any M-level quantizer. The Gaussian case refers to Gaussian sensor noise of variance σ_v^2. The general case includes any (admissible) sensor noise $v[n] = \sigma_v\,\tilde{v}[n]$.

where $\mathcal{B}\left(0;\ \tilde{s}\right)$ is the Cramér-Rao bound for estimating A from a single sample of the IID sequence $\tilde{s}[n]$. Hence, since $\mathcal{B}\left(A;\ \mathbf{s}^N\right)$ is independent of A, both $\mathcal{B}_{\max}\left(\Delta\right)$ and $\mathcal{L}_{\max}(\Delta)$ can be used interchangeably to assess the performance of signal quantizers.

Table 1 summarizes the performance limits for a number of important scenarios. In any of these scenarios the worst-case information loss can be conveniently characterized in terms of the ratio $\chi = \Delta/\sigma_v$, which we may view as a measure of peak-signal-to-noise ratio (SNR). Specifically, for pseudo-noise control inputs with properly chosen power levels the worst-case information loss grows only quadratically with χ, while it always grows faster than quadratically in the control-free case for any (admissible) sensor noise. For scenarios where the control input is known for estimation, the associated worst-case information loss can be made to grow as slow as χ. Finally, if feedback from the quantized output to the control input is available and properly used, a fixed small information loss, which is independent of χ, can be achieved.

3.1. Pseudo-noise Control Inputs

We next consider control inputs $w[n] = \sigma_w\tilde{w}[n]$ that are sample paths of an IID process independent of $v[n]$, and determine the performance limits in estimating A from $\mathbf{y}^N$ by exploiting the statistical characterization of $w[n]$ at the receiver. In particular, we wish to select σ_w to minimize the worst-case information loss.

Since $\alpha[n] = v[n] + w[n]$ is IID, we have $\mathcal{B}\left(A;\ \mathbf{y}^N\right) = \mathcal{B}\left(A;\ y\right)/N$, where $\mathcal{B}\left(A;\ y\right)$ is the Cramér-Rao bound for estimating A based on a single sample of the IID sequence $y[n]$. By taking the second partial derivative of the associated log-likelihood function with respect to A followed by an expectation, we obtain

$$\mathcal{B}\left(A;\ y\right) = \left(\sum_{i=1}^M \frac{\left[p_\alpha\left(X_{i-1} - A\right) - p_\alpha\left(X_i - A\right)\right]^2}{C_\alpha\left(X_{i-1} - A\right) - C_\alpha\left(X_i - A\right)}\right)^{-1} . \tag{9}$$

Eqn. (9) also provides the Cramér-Rao bound for estimates of A from $y[n]$ in the control-free case where α is replaced by v.

Proper use of pseudo-noise can improve performance over the control-free system for any (admissible) $p_{\tilde{v}}$ $(\cdot)$ and for any M-level quantizer. Let $\mathcal{B}_{\max}\left(\Delta; \sigma_v, \sigma_w\right)$ denote (6) for a given Δ, σ_v and σ_w. Since $\tilde{v}[n]$ is admissible, the bound (9) is continuous in σ_v, and so is $\mathcal{B}_{\max}\left(\Delta; \sigma_v, \sigma_w\right)$. Thus, given any fixed $\sigma_w > 0$ and Δ, for small enough σ_v we have

$$\mathcal{B}_{\max}\left(\Delta; \sigma_v, \sigma_w\right) \approx \mathcal{B}_{\max}\left(\Delta; 0, \sigma_w\right) . \tag{10}$$

Substitution of (10) and (8) in (7) reveals that $\mathcal{L}_{\max}^{\mathrm{pn}}(\chi) \sim \chi^2$ is achievable for large χ . Furthermore, since $\mathcal{B}_{\max}\left(\Delta; \sigma_v, \sigma_w\right)$ is

also continuous in σ_w, for any $F(\cdot)$ with fixed $M < \infty$

$$\inf_{\sigma_w \in [0, \infty)} \mathcal{B}_{\max}(\Delta; 0, \sigma_w) > 0 , \qquad (11)$$

which in conjunction with (8) and (10) implies that $\mathcal{L}_{\max}$ can not be made to grow slower than χ^2 for pseudo-noise control inputs.

In the case that the sensor noise level is fixed, from (10)–(11) the optimal worst-case information loss rate is achieved by $\sigma_w = \lambda \Delta$ for any $\lambda > 0$. For comparison, the control-free worst-case information loss grows faster than χ^2 for large χ in any sensor noise [2]. Remarkably, pseudo-noise control inputs with properly selected power levels improve performance over the control-free systems at high peak SNR for any sensor noise.

In the special case that $v[n]$ is Gaussian with variance σ_v^2, and $M = 2$ (i.e., $F(x) = \text{sgn}(x)$), the control-free worst-case information loss is given by

$$\mathcal{L}_{\max}^{\text{free}}(\chi) = 2\pi Q(\chi) Q(-\chi) e^{\chi^2} , \qquad (12)$$

where $Q(x) = \int_{t=x}^{\infty} \exp(-t^2/2)/\sqrt{2\pi} \, dt$. For comparison, if $w[n]$ is Gaussian, the worst-case information loss is optimized with $\sigma_w \approx \pi/2 \Delta$ for large χ [2], and is given by

$$\mathcal{L}_{\max}^{\text{pn}}(\chi) \approx \frac{8}{\pi} Q\left(\frac{\pi}{2}\right) Q\left(-\frac{\pi}{2}\right) e^{\frac{\pi^2}{4}} \chi^2 . \qquad (13)$$

3.2. Known Control Inputs

We next outline the performance limits for scenarios where the estimator can exploit detailed knowledge of a suitably designed control waveform. In particular, we present control input selection strategies that achieve the minimum possible growth rate of the worst-case information loss as a function of χ.

The Cramér-Rao bound for unbiased estimates of A based on $\mathbf{y}^N$ and given knowledge of the associated N samples of $w[n]$ is denoted by $\mathcal{B}(A; \mathbf{y}^N, \mathbf{w}^N)$ and satisfies

$$\mathcal{B}\left(A; \mathbf{y}^N, \mathbf{w}^N\right) = \left[\sum_{n=1}^{N} [\mathcal{B}(A + w[n]; y)]^{-1}\right]^{-1} , \qquad (14)$$

where $\mathcal{B}(A; y)$ is given by (9), with α replaced by v. In [2] it is shown that for any known control input selection method the worst-case information loss grows at least as fast as χ, for any sensor noise and any M.

Classes of periodic waveforms parameterized by the period K are appealing candidates for known control inputs, since they are easy to construct and can be chosen so that the worst-case information loss grows at the minimum possible rate. In order to achieve the minimum possible growth rate it suffices to select $w[n]$ from a properly constructed K-periodic class for which there is an one-to-one correspondence between each element in the class and K. If N is a multiple of K, optimal selection of the control input in this case amounts to selecting the period K that minimizes (14)

$$K_{\text{opt}}(\Delta, \sigma_v) \triangleq \arg\min_K \sup_{|A| < \Delta} \frac{K}{\sum_{n=1}^{K} [\mathcal{B}(A + w[n]; y)]^{-1}} \qquad (15)$$

where $\mathcal{B}(A; y)$ is given by (9) with α replaced by v.

The construction of the K-periodic class in the $M = 2$ case is based on the observation that in the control-free scenario the worst-case information loss grows with Δ for fixed σ_v, i.e., the information loss is typically largest for parameter values that are

furthest from the quantizer threshold. To optimize over the worst-case performance, we construct the K-periodic waveform $w[n]$ so as to minimize the largest distance between any A in $(-\Delta, \Delta)$ and the closest *effective* quantizer threshold. Specifically, we consider K-periodic sawtooth control inputs, namely,

$$w[n] = \delta_w \cdot \left(-\frac{K-1}{2} + n \bmod K\right) , \qquad (16)$$

where the effective spacing between thresholds is given by $\delta_w = 2\Delta/(K-1)$. The net effect of $w[n]$ from (16) and the symmetric two-level quantizer is equivalent to a two-level quantizer with a periodically time-varying threshold that comes within at least $\delta_w/2$ of any possible parameter value once every K samples.

In the context of the class (16), strategies that select K so as to keep a fixed sawtooth spacing δ_w achieve the minimum possible growth rate. In particular, if we select K in (16) via

$$K = \lceil \lambda \chi + 1 \rceil \qquad (17)$$

where $\lambda > 0$, the worst-case information loss grows linearly with χ [2]. In general, there is an optimal λ for any $p_{\tilde{v}}(\cdot)$, resulting in an optimal normalized sawtooth spacing. Specifically, consider the normalized sawtooth spacing, namely, $d(\chi; K) = \delta_w/\sigma_v = 2\chi/(K-1)$, and let $d_{\text{opt}}(\chi)$ denote the normalized spacing associated with the optimal period $K_{\text{opt}}(\chi)$ from (15). A method for computing the asymptotic optimal normalized spacing $d_\infty = \lim_{\chi \to \infty} d_{\text{opt}}(\chi)$ associated with a particular $p_{\tilde{v}}(\cdot)$ is also outlined in [2]. For instance, if the sensor noise is Gaussian with variance σ_v^2, we have $d_\infty \approx 2.5851$, while for large χ [2]

$$\mathcal{L}_{\max}^{\text{per}}(\chi) \approx 1.4754 \left(\frac{2\chi}{d_\infty} + 1\right) . \qquad (18)$$

K-periodic sawtooth waveforms similar to (16) can be used as known control inputs in quantizer systems with $M > 2$, and can be chosen so that the worst-case information loss grows as χ. In this case, the existence of multiple quantizer thresholds allows for reduction of the dynamic range that each $w[n]$ needs to span [2].

3.3. Control Inputs in the Presence of Feedback

In this section we develop performance bounds for estimates of A based on $\mathbf{y}^N$, where the control input $w[n]$ is a function of all past quantized observations. In particular, we show that the worst-case information loss for any feedback-based control input strategy is lower bounded by the *minimum* possible information loss for the same system with $w[n] = 0$. As (14) reveals, for any $|A| < \Delta$ we can obtain information loss equal to $\mathcal{L}(A_0)$ by selecting $w[n] = A_0 - A$. In particular, if there exists a real A_* for which $\mathcal{B}(A; y) \geq \mathcal{B}(A_*; y)$ for all real A (where $\mathcal{B}(A; y)$ is given by (9) with α replaced by v), then using (14) and (8) in (4) we obtain

$$\mathcal{L}\left(A; \mathbf{w}^N\right) \geq \mathcal{L}(A; A_* - A) = \mathcal{L}(A_*) , \qquad (19)$$

with equality achieved for $w[n] = A_* - A$, $1 \leq n \leq N$, where $\mathcal{L}(A)$ is given by (4), and $\mathcal{B}(A; y)$ by (9) with α replaced by v.

The minimum information loss from (19) decreases as the number of quantization levels increases; as expected, $\mathcal{L}(A_*)$ tends to zero as M approaches infinity for any sensor noise [2].

For the symmetric two-level quantizer in Gaussian noise, use of (9) for $\sigma_\alpha = \sigma_v$ in (4) reveals that $A_* = 0$ and

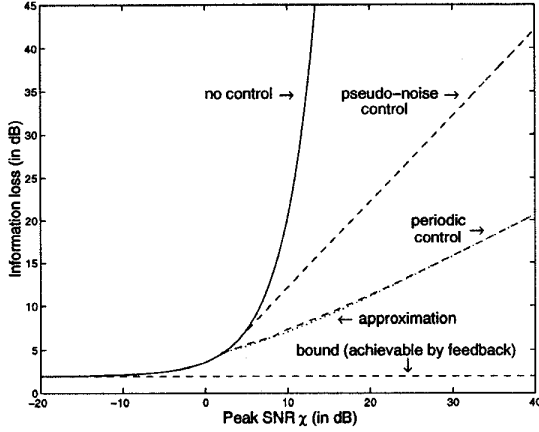$$\mathcal{L}\left(A; \mathbf{w}^N\right) \geq \mathcal{L}(0) = \frac{\pi}{2} . \qquad (20)$$

3275

Figure 2: Worst-Case information loss over $|A| < \Delta$ for a two-level quantizer in Gaussian noise of variance σ_v^2, for various control strategies. The dotted curve depicts (18). The lower dashed line depicts the minimum possible information loss ($\approx$ 2 dB).

Fig. 2 depicts the worst-case information loss of the two-level quantizer in Gaussian noise for the control input scenarios that we have examined. As the figure reflects, the performance of the control-free system (solid curve) degrades rapidly as χ increases. The benefits of pseudo-noise control inputs (upper dashed curve) at high peak SNR are clearly evident, and known periodic control inputs provide additional performance benefits (middle dashed curve) over pseudo-noise control inputs. Finally, in the presence of feedback from the quantized output to the control input, the performance is lower bounded by the minimum possible information loss of approximately 2 dB from (20), which is independent of χ.

4. EFFICIENT ESTIMATION

We next outline control input selection methods and associated estimators which achieve the performance limits developed in Sec. 3, for the Gaussian noise scenario in the context of pseudo-noise and feedback-based control inputs. Estimators for known control inputs as well as extensions to the nonGaussian case are developed in [2]. A natural measure of performance for a specific system comprising a control input, a quantizer and an estimator is the *MSE loss*, which we define as the ratio of the actual MSE of this estimator of A based on observation of $\mathbf{y}^N$, divided by $\mathcal{B}\left(A; \mathbf{s}^N\right)$. Analogously to $\mathcal{L}_{\max}$ in (4), the worst-case MSE loss of an estimator is defined as the supremum of the MSE loss over $|A| < \Delta$.

4.1. Pseudo-noise Control Inputs

For pseudo-noise control inputs, the maximum-likelihood (ML) estimator of A based on $\mathbf{y}^N$ over $|A| \leq \Delta$ satisfies

$$\hat{A}_{\mathrm{ML}}\left(\mathbf{y}^N; \Delta\right) = \arg\max_{|\theta| \leq \Delta} \ln P\left(\mathbf{y}^N; \theta\right) , \qquad (21)$$

where $\ln P\left(\mathbf{y}^N; \theta\right)$ is the associated log-likelihood function.

If $F(\cdot)$ is the symmetric two-level quantizer and since $\alpha[n]$ is Gaussian, (21) can be found in closed form, by setting to zero the

partial derivative of $\ln P\left(\mathbf{y}^N; A\right)$ with respect to A, viz.,

$$\hat{A}_{\mathrm{ML}}\left(\mathbf{y}^N; \Delta\right) = -\mathcal{I}_\Delta\left(\sigma_\alpha Q^{-1}\left(\frac{\mathcal{K}_1\left(\mathbf{y}^N\right)}{N}\right)\right) , \qquad (22)$$

where $Q^{-1}(\cdot)$ is the inverse of $Q(\cdot)$, $\mathcal{K}_{Y_i}\left(\mathbf{y}^N\right)$ denotes the number of elements in $\mathbf{y}^N$ that are equal to Y_i, and

$$\mathcal{I}_\Delta(x) = \begin{cases} x & |x| \leq \Delta \\ \Delta\,\mathrm{sgn}(x) & \text{otherwise} \end{cases} . \qquad (23)$$

Although the ML estimate (22) is asymptotically efficient in the sense that its MSE converges to the associated Cramér-Rao bound (9) for large enough N, convergence in not uniform in A [2].

For systems with $M > 2$, (21) can be obtained via an EM algorithm; by using $x[n] = A + \alpha[n]$ as the complete data we obtain [2]

$$\hat{A}_{\mathrm{EM}}^{(k+1)} = \mathcal{I}_\Delta\left(\hat{A}_{\mathrm{EM}}^{(k)} + \sum_{m=1}^M \frac{\sigma_\alpha \mathcal{K}_{Y_m}(\mathbf{y}^N)}{\sqrt{2\pi}N} \frac{\hat{b}_{m-1}^{(k)} - \hat{b}_m^{(k)}}{Q\left(\hat{z}_{m-1}^{(k)}\right) - Q\left(\hat{z}_m^{(k)}\right)}\right) \qquad (24)$$

initialized with $\hat{A}_{\mathrm{EM}}^{(0)} = 0$, where $\hat{z}_m^{(k)} = (X_m - \hat{A}_{\mathrm{EM}}^{(k)})/\sigma_\alpha$, and $\hat{b}_m^{(k)} = \exp\left(-[\hat{z}_m^{(k)}]^2/2\right)$. Provided that the likelihood function does not possess multiple local minima, we have $\hat{A}_{\mathrm{ML}}\left(\mathbf{y}^N; \Delta\right) = \lim_{k\to\infty} \hat{A}_{\mathrm{EM}}^{(k)}$. Empirical evidence suggests that $\lim_{k\to\infty} \hat{A}_{\mathrm{EM}}^{(k)}$ is also asymptotically efficient.

4.2. Control Inputs in the Presence of Feedback

The analysis of Sec. 3.3 suggests viable control input selection methods based on past quantized observations. For instance, for the system with $M = 2$ we may select $w[n] = -\hat{A}[n - 1]$, where $\hat{A}[n]$ is any consistent estimate of A based on $\mathbf{y}^n$, such as the ML estimate. In this case, the ML estimate $\hat{A}_{\mathrm{ML}}[n]$ of A based on $\mathbf{y}^n$ can be obtained using the following EM algorithm [2],

$$\hat{A}_{\mathrm{EM}}^{(k+1)}[n] = \mathcal{I}_\Delta\left(\hat{A}_{\mathrm{EM}}^{(k)}[n] + \sum_{m=1}^n \frac{\sigma_v y[m]}{\sqrt{2\pi}n} \frac{\exp\left(-\frac{\left(\hat{z}^{(k)}[m;n]\right)^2}{2}\right)}{Q(y[m]\,\hat{z}^{(k)}[m;n])}\right) \qquad (25)$$

with $\hat{A}_{\mathrm{EM}}^{(0)}[n] = \hat{A}_{\mathrm{ML}}[n - 1]$ and $\hat{A}_{\mathrm{ML}}[0] = 0$, where $\hat{A}_{\mathrm{ML}}[n] = \lim_{k\to\infty} \hat{A}_{\mathrm{EM}}^{(k)}[n]$, and $\hat{z}^{(k)}[m; n] = (\hat{A}_{\mathrm{ML}}[m - 1] - \hat{A}_{\mathrm{EM}}^{(k)}[n])/\sigma_v$. Empirical evidence suggests that (25) achieves the bound (20) for moderate N. Algorithms that achieve this bound, but require significantly fewer computations than (25), are developed in [2], as are extensions for systems with $M > 2$ and any sensor noise.

5. REFERENCES

[1] R. Benzi, A. Sutera, and A. Vulpiani, "The mechanism of stochastic resonance," *J. Phys.*, vol. A14, pp. L453–L457, 1981.

[2] H. C. Papadopoulos, G. W. Wornell, and A. V. Oppenheim, "Signal encoding from noisy measurements using quantizers with dynamic bias control," preprint.

[3] H. C. Papadopoulos, *Analysis and Applications of a Class of Signal Quantizers*. PhD thesis, M.I.T., Cambridge, MA. In preparation.

3276