Title
Predictive Coding and Control

Journal
IEEE Transactions on Control of Network Systems, 6(2)

ISSN
2372-2533

Authors
Huang, Chun-Chia
Amini, Behrooz
Bitmead, Robert R

Publication Date
2019

DOI
10.1109/tcns.2018.2882190

Peer reviewed

Predictive Coding and Control

Chun-Chia Huang, Behrooz Amini, *Student Member, IEEE* and Robert R Bitmead, *Fellow, IEEE*

***Abstract*—This paper deals with feedback control over a single fixed-rate channel using predictive coding at the transmitter side. The central thrust is to demonstrate that optimal control based on predictive coding plus fixed memoryless quantization at the transmitter, designed to improve the efficiency of the channel usage and exemplified (or perhaps extremized) by the transmission of the quantized innovations signal, in general requires the construction of the joint density of both the plant and predictor states at the receiver side and inherits a plant stability requirement, which is examined. The Bayesian filter is developed. This recursive filter's state density is used to compute the optimal feedback control. This is in contrast to the less complicated propagation solely of the predictor state, which would suffice in the linear quadratic optimal control problem – a feature that is elucidated. A linear non-quadratic optimal control example is provided to illustrate the approach and its benefits over control based on the recovered predictor state density or control without predictive coding. In each of these competing cases, a lower complexity receiver architecture is possible but at the expense of closed-loop control performance.**

***Index Terms*—Quantization, state estimation, Bayesian filter, innovations, optimal control.**

I. Introduction

PREDICTIVE quantization [1] is used in communication systems to whiten the transmitted digital signal and remove redundancy, thereby improving coding performance. In delay-free coding environments, a prediction of the source signal is computed and then subtracted from the signal to yield a prediction error, which is then quantized. Compared with the original signal, the prediction error is both closer to white, i.e. less correlated over time, and possesses a smaller variance, which aids in scaling the quantizer range for improved effectiveness. Such systems form the basis of familiar schemes such as ITU-T G.721/722/726 Adaptive Differential Pulse Coded Modulation (ADPCM) standards [2] and Delta Modulation [1]. The ADPCM schema is depicted in Figure 1. The quantizer $\mathcal{Q}$ is fixed and the adaptive predictor and gain serve to whiten and limit dynamic range fluctuations of the transmitted error signal, thereby improving distortion between transmitter and receiver. The decoder/receiver mimics the encoder/transmitter to undo its operations and recover an approximation, $\hat{s}_t$, of the signal s_t. ADPCM has been proven in service in telecommunications systems since at least 1984 and has spawned a number of variants as commercial lossy speech compressors. ADPCM has been interpreted as a disturbance rejection feedback control system in [3]. It provides a kindred example in the paper, but without its attendant gain adaptation, which could bring it closer to [4], nor its limitation to using the receiver-side signal at the transmitter. It manifests similar stability requirements.

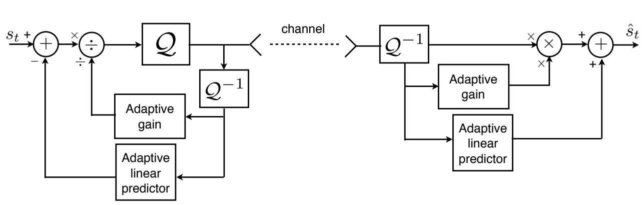

Fig. 1. ITU-G.722 Adaptive Differential Pulse-Coded Modulation schema.

The authors are with the Department of Mechanical and Aerospace Engineering, University of California, San Diego, CA 92093-0411, USA, e-mail: hhunhia@gmail.com;bamini,rbitmead@eng.ucsd.edu.

For network control systems, these methods can be applied in the link between plant and controller to achieve more efficient use of the available link bit-rate and, thereby, improved control performance because the more effective coding leads to more accurate reconstruction of the transmitted signal at the receiver. It is this reconstruction and, in particular, plant state density estimation, which is the focus of this paper. The Bayesian filter is used to calculate the joint and marginal densities of the plant and predictor states conditioned on the received data. The general set-up is depicted in Figure 2 and will be made precise shortly.

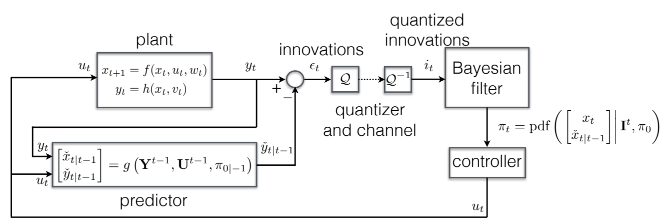

Fig. 2. Predictive quantization based feedback control set-up.

While the structure and analysis with the Bayesian filter pertains for general nonlinear systems, the most edifying and best studied case concerns linear plants with Gaussian noise paired with the Kalman filter as the state and output predictor. In this case, the innovations sequence is Gaussian, zero mean and white. Quantized innovations state estimation has been studied in this case in [5], [6], [7], [8], both from the perspective of state estimation and of control, notably LQG control. We too shall specialize to the linear Gaussian case, since the analysis is both relatively direct and most informative, because the comparator unquantized controls are so well known. However, we shall take an immediate departure from the linear predictive coding approach of signal processing by seeking at the receiver to compute the precise conditional state density, $p(x_t|\mathbf{I}^t, \pi_0)$, rather than to capture the transmitter side prediction or output, $\check{x}_{t|t-1}$ or $\check{y}_{t|t-1}$. Indeed, part of the message is that the plant state density is in general the important aspect for control; a point which we illustrate with an example.

Pertinent prior literature

Quantization in control under communications constraints is a longstanding subject with emphases on both stabilization and performance [8], [9], [10], [4]. Studies include adaptive quantization of the control signal and of the plant output measurement sequence and include simultaneous coding and quantization. A subset of papers studies *dynamic quantization* [11], [12] in which the coding is restricted to a finite-dimensional linear system. Almost universally, the setting is linear systems control with a quadratic criterion function. From a signal processing and telecommunications perspective, predictive coding [1] is familiar and involves pre-whitening of the signal before quantization. In a Kalman filtering framework, such methods equate to quantization of the innovations sequence, which inherently is white, Gaussian with minimum variance. Thus, the innovations is optimally coded and so we seek a static memoryless quantizer as in ADPCM. This is a restriction that we make on the coder.

Beginning with Fischer [13] the connection between quantization and LQ control performance was studied with the optimal quantizer being time-varying, although with an asymptotically time-invariant dequantizer, with level boundaries depending on the current state estimate. For Fischer, the quantizer operates on the computed optimal control signal. Fu [14] and Yüksel [15] extend these results to causal coding quantizers and Fu identifies some technical errors in [13]. The focus is on fixed-rate quantizer design and the existence or otherwise of a separation theorem in this case. The work [16] treats a related output feedback control problem with variable bit rate coding and random channel delay, deriving bounds for the limiting average codeword length given a specified bound on the controlled state covariance.

Papers [17], [18] consider the *cost-rate tradeoff* in linear quadratic control. This problem seeks the lowest average bitrate, $\mathbb{R}(b)$, channel required to achieve a specified LQ performance b. Stavrou et al. [19] treat a related Kalman filtering problem and seek the minimal data rate required to achieve a specific distortion or mean squared error between the plant state and the receiver-side Kalman filter. All three papers start from a vector autoregressive plant model with fully measured state at the transmitter. They restrict attention to zero-delay coding schemes which is appropriate for feedback control. Each of these papers arrives at a coding scheme based on Shannon entropy coding of a quantized innovations, but a different innovations from here.

Each of [17], [18], [19] and ourselves has a single bandlimited forward channel and a high-fidelity return channel, used in [19] for communication of the receiver state estimate and in [17], [18] and here to communicate the control signal. Kostina and Hassibi [17] and Tanaka et al. [18] require stabilizability of the plant system's $[A, B]$ pair and adjust the coding to accommodate the plant feedback stabilization as part of their calculation. This is evident in their inherent satisfaction of Tatikonda's and Mitter's [20] and Nair's and Evans' [4] lower bound on the bit rate based on the unstable eigenvalues of A. Here, because we quantize directly the output innovations process constructed at the transmitter, we must limit the analysis to stable plants. This is a property proven in Corollary 1 and explained as it arises in ADPCM. Tanaka et al. [18] and Stavrou et al. [19] incorporate the communication of the receiver's state estimate of the plant state to compute their state innovations, which is then quantized in the coder. We note too that the encoding strategy in [17] also is based on quantizing then entropy coding the state innovations between the true plant state and the receiver's estimate and can accommodate partial state measurements with Gaussian noise. Initial values, channel noise and quantization error force the state estimates at the transmitter and receiver to differ. The effect of this is seen in the additional stability condition in Corollary 1. To improve performance and simplify analysis the quantizers can be subtractively dithered in each of these works.

Uniform subtractive dithered quantizers of finite support have been analyzed in [21], [22], where they demonstrated that such quantizers and predictive coding arise in achieving a Gaussian nonanticipative rate distortion function with a specified mean square filtering error overbound, provided quantizer overload is avoided. In this context, they treat full-state transmission over n-parallel AWGN channels with feedback of the state prediction from the decoder. They also propose an approach to mitigate the effects of overload. This formulation differs from ours in the full-state communication and in the feedback of the receiver prediction. The computation of the innovations process before encoding, however, is similar and demonstrated to be close to optimal for their constrained zero-delay coding problem. For us, we take the predictive coder with uniform subtractive dithered quantizers in each channel as the starting point. Papers [21], [22] provide a justification of this as a sensible starting point. The work of [23] connects some of these coding aspects to the feedback control problem.

Contributions

- The results commence from the noisy plant output measurement rather than the full state and extend to nonlinear systems with non-quadratic optimal control. The treatment includes quantized LQG and computed examples.
- They expose the role of the Bayesian filter and state conditional density, rather than moment, reconstruction.
- They are based on or limited to predictive coding at the transmitter using quantized output innovations by a memoryless, fixed-rate quantizer. This solution is necessarily zero-delay and fixed bit rate, in comparison with the other entropy coded approaches which are variable rate.
- Our problem focus is to optimize the plant performance given the communications structure based on predictive coding. This is in contrast to [17], [18], [19] where the control performance is specified and communications required to achieve this is then designed.

A number of papers are dedicated to reconstruction of the conditional density of either the plant state or the predictor state using methods allied with Kalman filtering [5], Bayesian filtering [6] and particle filtering [7]. These are closest in focus to the current work, although they are limited to the consideration of linear systems. Once the transmitter-side prediction and quantization scheme is decided, the problem that we consider is the reconstruction by Bayesian filter at the receiver of the filtered conditional density of the plant state, as opposed to the density of the predictor state. We do this in a fully nonlinear context and then specialize to the linear problem. We provide theory and demonstrate by example the control performance benefits of using: the plant state density, the filtered density versus the predicted density, and the quantized prediction error versus quantized output signals. The contribution is to provide a unifying nonlinear framework in which to treat the predicted and filtered state conditional density reconstruction and to explore its connection with other approaches from the linear context based on reconstruction of the predictor state conditional density or its mean value.

Notation

We denote probability density functions (pdfs) by $p(\cdot)$. Gaussian pdfs of mean μ and covariance P are denoted $\mathcal{N}(\mu, P)$. The initial pdf of the transmitter state will be denoted $\pi_{0|-1}$. The data available at the receiver at time t is $\mathbf{I}^t = \{\pi_{0|-1}, i_0, \ldots, i_t\}$. By the same token, the data available at the transmitter is $\mathbf{E}^t = \{\pi_{0|-1}, \epsilon_0, \ldots, \epsilon_t\}$. We presume that, at time t, the input signal, u_t, computed at the receiver/controller is available also to the transmitter side.

II. NONLINEAR PREDICTIVE QUANTIZATION – TRANSMITTER SIDE

We consider separately the general case of a nonlinear plant at the transmitter side and its specialization to a linear Gaussian system. The Bayesian filter construction applies to both but the linear formulation allows us to draw on well understood ideas from Kalman filtering. For comparison and brevity, we present them side by side. Although the quantized linear innovations problem has been more widely studied.

A. Nonlinear plant & predictor

The nonlinear stochastic plant system is described by

$$x_{t+1} = f_t(x_t, u_t, w_t), \quad x_0, \tag{1}$$

$$y_t = h_t(x_t, v_t). \tag{2}$$

Here, state $x_t \in \mathbb{R}^{n_x}$, input $u_t \in \mathbb{R}^{n_u}$, output $y_t \in \mathbb{R}^{n_y}$, process noise $w_t \in \mathbb{R}^{n_w}$, measurement noise $v_t \in \mathbb{R}^{n_v}$. Noise sequences $\{w_t\}$ and $\{v_t\}$ are mutually independent, zero-mean and white with known densities. The plant initial condition, x_0, has known density, $\pi_{0|-1}$, and is independent from w_t and v_t for all t.

The measured output and control signals at the transmitter, u_t and y_t, are the inputs to a finite-dimensional predictor

$$\xi_{t+1} = \bar{g}_t(\xi_t, u_t, y_t), \quad \xi_0, \tag{3}$$

$$\check{y}_t = j_t(\xi_t). \tag{4}$$

The prediction, in turn, is combined with y_t to produce a prediction error or *innovations* signal.

$$\epsilon_t = y_t - \check{y}_t. \tag{5}$$

Using (4)-(5), then (3) becomes

$$\xi_{t+1} = g_t(\xi_t, u_t, \epsilon_t), \quad \xi_0, \tag{6}$$

since y_t can be reconstructed from ϵ_t and $\check{y}_t$.

B. Linear Gaussian plant & predictor

The linear plant system is described by

$$x_{t+1} = Ax_t + Bu_t + w_t, \tag{7}$$

$$y_t = Cx_t + v_t, \tag{8}$$

where $\{w_t\}$ and $\{v_t\}$ are mutually independent, white noises of known densities and also zero-mean Gaussian with covariances Q and R respectively. The state estimator and predictor is the Kalman predictor with state $\check{x}_t$ and innovations

$$\epsilon_t = \begin{bmatrix} C & -C \end{bmatrix} \begin{bmatrix} x_t \\ \check{x}_t \end{bmatrix} + v_t. \tag{9}$$

The predictor recursion is

$$\begin{aligned} \check{x}_{t+1} &= A\check{x}_t + Bu_t + L_t\epsilon_t, &(10)\\ &= (A - L_tC)\check{x}_t + Bu_t + L_ty_t, \\ &= (A - L_tC)\check{x}_t + L_tCx_t + Bu_t + L_tv_t, &(11)\\ \check{y}_t &= C\check{x}_t, \\ L_t &= A\Sigma_{t|t-1}C^T(C\Sigma_{t|t-1}C^T + R)^{-1}, \end{aligned}$$

where

$$\Sigma_{t|t-1} = E\{[x_t - \check{x}_{t-1}][x_t - \check{x}_{t-1}]^T|\mathbf{E}^{t-1}\},$$

$$\Sigma_{t|t} = E\{[x_t - \check{x}_t][x_t - \check{x}_t]^T|\mathbf{E}^{t-1}\}.$$

Here the covariance matrix, $\Sigma_{t|t-1}$, is given by the Riccati difference equation commencing from $\Sigma_{0|-1}$ [24]. The prediction error or *innovations* is given by (9).

C. Quantization

The 'quantizer' (or combined quantizer-dequantizer pair) in Figures 1 and 2, is a known single-valued function of the same dimension, n_y, as its input. Typically (and in our calculations below), the quantization function, $\mathcal{Q}$, maps real intervals to unique fixed digital signal values in each channel and the dequantization function, $\mathcal{Q}^{-1}$, maps these received values to unique points in their corresponding intervals.

$$i_t = \mathbf{Q}(\epsilon_t) = \mathcal{Q}^{-1}\left(\mathcal{Q}[\epsilon_t]\right). \tag{12}$$

Gersho and Gray [1] describe many quantizer designs for communication systems, covering both scalar and vector quantization including optimization for properties such as minimal distortion and the Lloyd-Max quantizer, which is adapted to signals with Gaussian distributions.

Assumption 1: The quantizer-dequantizer function, $\mathbf{Q}(\cdot)$, is known, finite range and memoryless.

Our formulation makes no other specific assumptions about the quantizer. Although, our calculations later for linear systems are based on a uniform (linear) mid-rise quantizer and its mid-point 'inverse'. We arbitrarily absorb the quantizer into the transmitter side, since we only care about the quantizer-dequantizer pair in the signal domain without regard to the specifics of the channel representation. Although, it is straightforward to incorporate other features into the full formulation, such as channel noise or the ADPCM structure as above.

More generally, the quantizer is a codec, a coder-decoder pair. Since we focus on control, we limit our study to delay-free coding. We consider memoryless coding for simplicity in order not to have to incorporate the codec states and to exploit the whiteness of the innovations. The important feature of quantized innovations signal, $\{i_t\}$, compared with the transmitter-side innovations, $\{\epsilon_t\}$, is that it has reduced information content, measured by entropy or other metrics, and this therefore diminishes its utility as a means to compute the optimal control. Understanding the cost of quantization borne by the control performance in this setting is the aim of this paper.

The following result, a specialization of the Data Processing Inequality [25], captures this relationship.

Lemma 1: Denote the following σ-algebras: $\mathcal{I}_t = \sigma(\mathbf{I}^t)$ and $\mathcal{E}_t = \sigma(\mathbf{E}^t)$. Then, since $i_t = \mathbf{Q}(\epsilon_t)$ for known $\mathbf{Q}(\cdot)$,

$$\mathcal{I}_t \subseteq \mathcal{E}_t. \tag{13}$$

Subtractive dithered quantizer for the linear case

A subtractive dithered quantizer, $\mathcal{Q}_d(\cdot)$, consists of a fixed, finite-range quantizer function $\mathfrak{Q}(\cdot)$ with a predetermined dither signal, $\{d_t\}$, which is known to both transmitter and receiver. It is defined

$$\mathcal{Q}_d(\epsilon_t) \triangleq \mathfrak{Q}(\epsilon_t + d_t) - d_t. \tag{14}$$

A *linear* quantizer is one with equally spaced steps with the center points of the steps mapping the input value to the same output value. A *midrise* quantizer has discontinuity at the origin of the input. We have this result following [26].

Theorem 1: Suppose quantizer $\mathbf{Q}(\cdot)$ is a subtractive, dithered, b-bit-per-channel, midrise, symmetric, linear quantizer with saturation values $\pm\zeta$. Suppose, further, that the subtractive dither signal is white, independent in each channel, and either uniformly distributed $\mathcal{U}(-\zeta/2^b, \zeta/2^b)$ or triangularly distributed, $d_t \sim \text{tr}(-\zeta/2^{b-1}, \zeta/2^{b-1})$, and known exactly to the transmitter and receiver. If the signal $\epsilon_t + d_t \in [-\zeta, \zeta]$, then the quantization noise

$$\psi_t \triangleq i_t - \epsilon_t = \mathcal{Q}_d(\epsilon_t) - \epsilon_t, \tag{15}$$

is white, independent from $\{\epsilon_t\}$, and uniformly distributed $\psi_t \sim \mathcal{U}[-\zeta/2^b, \zeta/2^b]$. That is,

$$\mathrm{E}(\psi_t) = 0, \;\; \mathrm{E}(\psi_t^2) = \frac{\zeta^2}{3 \times 2^{2b}} \triangleq \Psi. \tag{16}$$

D. Transmitter assumptions

Assumption 2:

1) $x_t \in \mathbb{R}^{n_x}$, $u_t \in \mathbb{R}^{n_u}$, $y_t \in \mathbb{R}^{n_y}$, $\xi_t \in \mathbb{R}^{n_\xi}$.
2) Control signal u_t is known to both transmitter and receiver at time $t+1$.
3) $\{w_t\}$ and $\{v_t\}$ are mutually independent, white noises of known densities and, in the linear case, also zero-mean Gaussian with known covariances.
4) The transmitter-side initial states, x_0 and ξ_0 (respectively $\check{x}_0$), have known joint density, $\pi_{0|-1}$, independent from $\{w_t, v_t\}$. In the linear case,
$$\pi_{0|-1} = \mathcal{N}\left(\begin{bmatrix}\hat{x}_{0|-1}\\ \hat{x}_{0|-1}\end{bmatrix}, \begin{bmatrix}\Sigma_{0|-1} & 0\\ 0 & 0\end{bmatrix}\right).$$
The receiver has knowledge of these densities.
5) In the nonlinear case, if the conditional mean of the plant state is computed at the transmitter, it is a function of ξ_t.
$$\check{x}_t = \mathrm{E}(x_t|\mathbf{Y}^{t-1}) = \mathrm{E}(x_t|\mathbf{E}^{t-1}) = \ell_t(\xi_t). \tag{17}$$
6) The function $g_t(\cdot,\cdot,\cdot)$ in (6) causes the predictor state update to be uniformly incrementally input-to-state stable [27]. In the linear case, A in (7) and (10) has all eigenvalues strictly inside the unit circle. The origin of this stability condition, at least in the linear case, is examined in detail in Subsection III-A.

III. Quantized Innovations Bayesian Filtering – Receiver Side

The nonlinear signal model for the sequence, $\{i_t\}$, arriving at the receiver comprises:

- For the nonlinear case,
  - Using (2), (4) and (5), $\epsilon_t = h_t(x_t, v_t) - j_t(\xi_t)$. Then (6) and (7) yield the combined state recursion
$$\begin{aligned} z_{t+1} &\triangleq \begin{bmatrix} x_{t+1}\\ \xi_{t+1}\end{bmatrix} = \begin{bmatrix} f_t(x_t,u_t,w_t)\\ g_t(\xi_t,u_t,\epsilon_t)\end{bmatrix},\\ &= \begin{bmatrix} f_t(x_t,u_t,w_t)\\ g_t(\xi_t,u_t,h_t(x_t,v_t)-j(\xi_t))\end{bmatrix},\\ &\triangleq \mathfrak{f}_t(z_t,u_t,w_t,v_t). \end{aligned} \tag{18}$$
  - Output equation
$$\begin{aligned} i_t &= \mathbf{Q}[h_t(x_t,v_t) - j_t(\xi_t)],\\ &\triangleq \mathfrak{h}_t(z_t,v_t). \end{aligned} \tag{19}$$
- Specializing to the linear Gaussian case with
$$F_t = \begin{bmatrix} A & 0\\ L_tC & A-L_tC\end{bmatrix}, n_t = \begin{bmatrix} w_t\\ L_tv_t\end{bmatrix}, H = \begin{bmatrix} C & -C\end{bmatrix},$$
  - state equation
$$z_{t+1} = \begin{bmatrix} x_{t+1}\\ \check{x}_{t+1}\end{bmatrix}, \tag{20}$$
$$\begin{aligned} &= \begin{bmatrix} A & 0\\ L_tC & A-L_tC\end{bmatrix}\begin{bmatrix} x_t\\ \check{x}_t\end{bmatrix} + \begin{bmatrix} w_t\\ L_tv_t\end{bmatrix},\\ &= F_tz_t + n_t, \end{aligned} \tag{21}$$
  - output equation
$$\begin{aligned} i_t &= \mathbf{Q}\left(\begin{bmatrix} C & -C\end{bmatrix}\begin{bmatrix} x_t\\ \check{x}_t\end{bmatrix} + v_t\right),\\ &= \mathbf{Q}\left(Hz_t + v_t\right). \end{aligned} \tag{22}$$

A. Open-Loop System Stability Condition

The linear predictive decoder immediately highlights a stability requirement on the source system (7) in order that the receiver-side innovations filter also be stable. As the predictive codec is envisaged as part of a feedback control scheme, this imposes a restriction on the class of plants to which such a scheme might be applicable.

Lemma 2: Consider the transmitter-side linear, time-varying system (20), with joint state space $\mathbb{R}^{2n_x}$, together with its predictively-coded innovations output signal, $\{\epsilon_t\}$ from (9). The subspace

$$\text{Span}\left\{\begin{bmatrix}I_{n_x}\\ I_{n_x}\end{bmatrix}\right\} \subset \mathbb{R}^{2n_x},$$

is unobservable and is associated with the eigenvalues of A.

Proof: Evidently and no matter the value of L_t,

$$\begin{bmatrix}C & -C\end{bmatrix}\begin{bmatrix}I_{n_x}\\ I_{n_x}\end{bmatrix} = 0_{n_y\times n_x},$$
$$\begin{bmatrix}A & 0\\ L_tC & A-L_tC\end{bmatrix}\begin{bmatrix}I_{n_x}\\ I_{n_x}\end{bmatrix} = \begin{bmatrix}I_{n_x}\\ I_{n_x}\end{bmatrix}\times A.$$

Corollary 1: Consider the transmitter-side predictive coding system (7)-(9),(15) with the receiver calculating its state estimate using the innovations filter,

$$\hat{x}_{t+1} = A\hat{x}_t + Bu_t + L_t i_t. \tag{23}$$

If system matrix A has any eigenvalues outside or on the open unit disk and $\check{x}_t - \hat{x}_t$ possesses non-zero component in the direction of this eigenvector for some t, then the error between estimates, $\check{x}_t$ at the transmitter and $\hat{x}_t$ at the receiver, grows unbounded with time.

For stable linear systems at the transmitter, which might consist of the joint stable plant and its stable predictor, the linear analysis of the closed-loop fails when the fixed-range quantizer overflows or saturates. Then, the assumptions of Theorem 1 fail and the quantization error ceases to exhibit the independence properties. Naturally, the Gaussian property of the system noises guarantees both eventual overflow and non-infinitesimal probability of overflow at any time. The probability of saturation of a fixed-quantized signal in these circumstances has been studied using Markov methods by (a subset of) the authors in [28] for both intermittent and quantized data. The time of first overflow is called the escape time there. A feature of that analysis is that for many systems, the escape time can be very large, depending on system parameters including feedback gain K and saturation level ζ.

If A has all eigenvalues in the open unit disk and escape time has yet to occur, then using (10), (15), (23), we have

$$\check{x}_{t+1} - \hat{x}_{t+1} = A(\check{x}_t - \hat{x}_t) - L_t\psi_t,$$

and, letting t grow while vainly betting on no escape,

$$\mathrm{E}\left[\hat{x}_t\right] \to \mathrm{E}\left[\check{x}_t\right] = \mathrm{E}\left[x_t\right],$$
$$\text{cov}\left[\check{x}_t - \hat{x}_t\right] \to \sum_{j=0}^{\infty}\left\{A^j L\Psi L^T A^{j^T}\right\}.$$

Here, L is the limiting value of the Kalman gain, L_t. These results are independent of the feedback control law other than central dependence of the escape time itself on K.

These results of Lemma 2 and and the discussion following for the underlying linear time-invariant system (7) carry over directly to linear time-varying systems by the same argument [29]. For linear systems with nonlinear measurements, one may appeal to Curry [30], who shows that the innovations is independent of the control signal, and Lemma 2 to argue that the unobservability problem persists for these systems. For more general nonlinear systems, it is less clear how instability of the plant might be manifested in the error between transmitter-side and receiver-side estimates. Although, Assumption 2.6 would be needed to analyze the estimate errors locally.

The clear admonition of Corollary 1 is to apply predictive coding solely to the control of stable systems. The reconstruction of the state estimate from the receiver innovations otherwise is unstable. This occurs because there is no output injection of $\hat{x}_t$ into the computation of ϵ_t and thus i_t. To our knowledge, this was first observed in [31].

We also note that the practically implemented G.722 ADPCM standard [2], in the definition of the adaptive predictor in its Section 3.6, includes specific pole-parameter restrictions to enforce stability of the prediction model at both the transmitter and receiver; it limits the number of poles to two and projects the parameters to ensure stability. Thus, we offer four observations.

(i) Predictive coding does not appear suited to the control of unstable systems. We believe this to be a novel observation and a reflection of the nature of predictive coding itself.
(ii) The practical success of ADPCM indicates that predictive coding has something to offer in control of stable plants.
(iii) The computed feedback control performance, in the example presented in Section V for a stable linear system close to instability, is significantly improved (reduced by 56%) using predictive coding for a finite bit-rate channel over that in which the output signal itself is quantized.
(iv) The nonlinear example in Section V also demonstrates significant improvement – in this case with a maximization criterion and by a factor of 15 – of the innovations based approach versus quantization of the output signal.

B. Bayesian filter

The Bayesian filter uses the sequence of measurements, $\{u_t, i_t\}$, to compute recursively the joint conditional density of the transmitter-side state

$$\pi_t = p\left(z_t | \mathbf{I}^t\right) = p\left(\begin{bmatrix}x_t\\ \xi_t\end{bmatrix}\middle| \mathbf{I}^t\right).$$

For the general nonlinear system (18)-(19),

$$z_{t+1} = \mathfrak{f}_t(z_t, u_t, w_t, v_t),$$
$$i_t = \mathfrak{h}_t(z_t, v_t),$$

the Bayesian filter recursion is [32]

$$p(z_t|\mathbf{I}^t) = \frac{p(i_t|z_t,\mathbf{I}^{t-1})p(z_t|\mathbf{I}^{t-1})}{\int_{z_t} p(i_t|z_t,\mathbf{I}^{t-1})p(z_t|\mathbf{I}^{t-1})\,dz_t},$$
$$= \frac{p(i_t|z_t,\mathbf{I}^{t-1})p(z_t|\mathbf{I}^{t-1})}{p(i_t|\mathbf{I}^{t-1})}, \tag{24}$$

$$p(z_{t+1}|\mathbf{I}^t) = \int_{z_t} p(z_{t+1}|z_t,\mathbf{I}^t)p(z_t|\mathbf{I}^t)\,dz_t. \tag{25}$$

We have been careful to include explicitly the conditioning on $\mathbf{I}^t$ in both integrands, since this plays a role in the case of correlated process and measurement noises, as here.

The recursion commences from $\pi_{0|-1} = p(z_0)$ and consists of two parts:

- measurement update (24) with $p(i_t|z_t,\mathbf{I}^{t-1})$ derived from the output equation (2) via the function $\mathfrak{h}_t(\cdot,\cdot)$ and the density of v_t.
- time update (25) with $p(z_{t+1}|z_t,\mathbf{I}^t)$ reflecting $\mathfrak{f}_t(\cdot,\cdot,\cdot,\cdot)$ in (18) and the joint densities of w_t, v_t and i_t.

For linear Gaussian systems without quantization, the Bayesian and Kalman filters coincide, although the Kalman filter more efficiently computes just the sufficient statistics of these conditional densities: the mean and covariance.

C. Reduced-order Bayesian filter

We note that the (full-order) Bayesian filter (24)-(25) yields the joint density π_t. The marginal densities, $p(x_t|\mathbf{I}^t)$ and $p(\xi_t|\mathbf{I}^t)$, are simply computed from π_t by integration. If, however, only the predictor state density, $p(\xi_t|\mathbf{I}^t)$, is desired, this can more easily be calculated by applying the Bayesian filter to state equation (6) with measurement equation (12).

$$\begin{aligned}\xi_{t+1} &= g_t(\xi_t,u_t,\epsilon_t), \quad \xi_0,\\ i_t &= \mathbf{Q}(\epsilon_t).\end{aligned}$$

In the quantized linear Gaussian case, this corresponds to using (10) and (12) to compute the density $p(\check{x}_t|\mathbf{I}^t)$ without the attendant calculation of $p(x_t|\mathbf{I}^t)$. This results in a reduced-order Bayesian filter which yields solely the conditional density of $\check{x}_t$. Such receiver-side reconstruction of the predictor state is the mainstay of predictive coding in signal processing [1]. Such ideas underpin some approaches to quantized innovations Kalman and Bayesian filtering [5], [6], [7] and Delta Modulation.

D. Computational issues

The Bayesian filter is numerically demanding. Notably, the integration in time update (25) presents a challenge to computation, since it involves performing a $2n_x$-dimensional integral at each sample point in a $2n_x$-dimensional space, yielding an operation count of $\mathcal{O}\left(16n_x^4\right)$. By contrast, the measurement update (24) is relatively benign at $\mathcal{O}(4n_x^2)$. Increasing the number of sample points per dimension rapidly causes problems. This is exacerbated by densities in x_t and $\check{x}_t$ being poorly conditioned, such as can occur with singular densities for $\check{x}_t$ and with very fine quantization. No special numerical 'tricks' were applied in the computations in this paper. The Particle filter may be applied to implement approximately the Bayesian filter using resampling ideas to manage calculations. This comes with its own set of problems, issues and fixes [33]. In our examples, we compute the Bayesian filter on a fixed grid rather than by particles.

The distinction between computation of $p(x_t|\mathbf{I}^t)$ and $p(\check{x}_t|\mathbf{I}^t)$ rests solely with the willingness to devote resources to computation at the receiver. They both operate on the same data. In a control setting, this is also connected to the admissibility of accepting greater computational delay at the receiver, which itself might preclude any advantage versus the delay in accepting a prediction-based control signal.

It is certainly worth remarking that the Bayesian filter calculations, notably central recursion (25), lend themselves to highly parallelized implementation, which suggests using GPUs or other processor architectures to achieve speedup [34].

E. Density properties

We have the following general results for the nonlinear and linear cases.

Theorem 2: If the conditional mean state estimate is computed at the transmitter, so that

$$\check{x}_t = \mathrm{E}(x_t|\mathbf{E}^{t-1}),$$

then the two receiver-side conditional means coincide. That is,

$$\mathrm{E}(x_t|\mathbf{I}^{t-1}) = \mathrm{E}(\check{x}_t|\mathbf{I}^{t-1}). \tag{26}$$

Proof: Lemma 1 shows that $\sigma(\mathbf{I}^{t-1}) = \mathcal{I}_{t-1} \subseteq \mathcal{E}_{t-1} = \sigma(\mathbf{E}^{t-1})$. The *smoothing property* of conditional expectation [35] then establishes that

$$\mathrm{E}[\check{x}_t|\mathbf{I}^{t-1}] = \mathrm{E}\left[\mathrm{E}[x_t|\mathbf{E}^{t-1}]|\mathbf{I}^{t-1}\right] = \mathrm{E}[x_t|\mathbf{I}^{t-1}].$$

Theorem 3: In the general nonlinear case, if the innovations sequence, $\{\epsilon_t\}$, is white, then

$$p(\check{x}_t|\mathbf{I}^t) = p(\check{x}_t|\mathbf{I}^{t-1}).$$

Proof: The whiteness property of $\{\epsilon_t\}$ implies that the received signal, $\{i_t\}$, also is white and, since $\check{x}_t$ is computed causally from $\mathbf{E}^{t-1}$, that $\check{x}_t$ is independent from ϵ_t and, therefore, i_t.

• Theorem 2 states that, should the objective be to calculate the conditional mean of the plant state at the receiver, then one might use the reduced-order Bayesian filter to achieve this.
• Theorem 3 establishes that in the case where the prediction errors are white, the conditional $\check{x}_t$ density at the receiver (and transmitter) will update only at the time-update stage.
• We appreciate that, while the transmitter side recursion (3) is driven by the innovations, ϵ_t, derived directly from x_t, the receiver side Bayesian filter driven by the quantized innovations requires stability of $g_t(\cdot,\cdot,\cdot)$ as in Assumption 2.6 in order that its predictor state estimate not diverge too greatly from that at the transmitter. This, underlying predictor stability requirement is inherent in all works in this field and reflects the estimate convergence condition for two state estimators both driven by the innovations of one of the estimators.

IV. CONTROLLER

The sequence of quantized innovations, $\{i_t\}$, arrives at the receiver and is used to generate the feedback control signal, u_t, as depicted Figure 2. The Bayesian filter is applied to the received sequence to yield conditional densities $p(x_t|\mathbf{I}^{t-1})$ and $p(\check{x}_t|\mathbf{I}^t)$. We have the following result from stochastic optimal control.

Theorem 4 (Kumar & Varaiya [36], Bertsekas [37]): For any choice of optimization criterion admitting a bounded value function, the optimal causal output feedback control for system (18)-(19) is

$$u_t^{\text{opt}} = \mathfrak{k}_t(\pi_t),$$

where $\pi_t = p(z_t|\mathbf{I}^t)$ and feedback policy $\mathfrak{k}_t(\pi_t)$ is found by solving the stochastic dynamic programming equation based on the associated objective function.

The result follows from the Markovian property of (18) and involves two computationally challenging aspects; the

Bayesian filter for π_t and the solution of the stochastic dynamic programming equation. Inherently, this latter piece is the harder and requires duality of the controller. For our predictive coding setup, since the plant state x_t is not dependent on ξ_t, we can say more.

Corollary 2: For system (1)-(2), with objective function dependent solely on future $\{x_t, u_t\}$, the optimal causal output feedback control solution is

$$u_t^{\text{opt}} = k_t\left(p(x_t|\mathbf{I}^t)\right).$$

That is to say, the predictor state, ξ_t, and its conditional density are not explicitly part of the optimal control solution.

The control signal computation is based on the conditional x_t density from the Bayesian filter. Our central aim is to describe the Bayesian filter for estimating the joint state,

$$z_t = \begin{bmatrix} x_t \\ \check{x}_t \end{bmatrix},$$

representing the predictively coded transmitter side.

Theorem 5 (Curry [30], Section 5.4, pp. 75-78, Appendix D, pp. 114-116): For the linear state system (7) with: memoryless nonlinear measurement

$$y_t = \varphi_t(x_t, v_t),$$

independent, white but not necessarily Gaussian noise processes $\{w_t\}$ and $\{v_t\}$, and quadratic objective function

$$J_t = \mathrm{E}\left(\sum_{k=t}^{N+1} x_k^T Q_k x_k + u_k^T R_k u_k \,\middle|\, \mathbf{Y}^t, \mathbf{U}^{t-1}, \pi_{0|-1}\right), \quad (27)$$

the optimal output feedback control is given by

$$u_t^\star = K_t\,\mathrm{E}(x_t|\mathbf{Y}^t, \mathbf{U}^{t-1}, \pi_{0|-1}),$$

where, K_t is the LQ optimal feedback gain computed from the control Riccati equation.

Corollary 3: For linear system (7) with quantized innovations measurement (12), quadratic objective function (27), and one time-sample delay in the controller, the optimal output feedback control is given by

$$u_t^\star = K_t\,\mathrm{E}(\check{x}_t|\mathbf{I}^{t-1}, \mathbf{U}^{t-1}, \pi_{0|-1}). \quad (28)$$

Proof: Mita [38] establishes the optimality of the LQ-optimal feedback gain with the predictive state estimate. This translates directly to Curry's result. For linear systems, the innovations sequence is white and one may then appeal to Theorem 2 to establish that $\mathrm{E}(x_t|\mathbf{I}^{t-1}, \mathbf{U}^{t-1}, \pi_{0|-1}) = \mathrm{E}(\check{x}_t|\mathbf{I}^{t-1}, \mathbf{U}^{t-1}, \pi_{0|-1})$ and the result follows.

• Theorem 5 shows that, despite the nonlinear measurements, the optimal LQ control is to feed back the filtered conditional mean of the state. This would suggest using the receiver to compute this quantity and then to calculate the control.
• Appealing to the results of [38], we see that, for a single delay controller, the same calculation holds but with the predicted state estimate at the receiver, which is simpler to compute.
• Corollary 3 uses Theorem 2 to replace the conditional mean of x_t by that of $\check{x}_t$, which incurs substantially fewer computations for its estimation.
• It is worth noting that the filtered conditional mean of x_t is different from its predicted conditional mean, even though Theorem 3 shows that they coincide for $\check{x}_t$ when the innovations is white, as in the linear case.
• This theorem and corollary are specialized to linear systems with quadratic criteria. We shall see shortly an example, where the optimal controller depends on the complete density of x_t and not just on $\check{x}_t$. In this case, the feedback controller based on $p(x_t|\mathbf{I}^{t-1})$ outperforms that based on $p(\check{x}_t|\mathbf{I}^{t-1})$.

V. Quantized Linear Innovations Filtering

We now specialize the development to the case of quantized linear Gaussian innovations, the Bayesian filter for which is derived in the Appendix. In this section, we do not use a subtractive dithered quantizer and compute the full Bayesian filter. In the following section, we apply the subtractive dithered quantizer and avail ourselves of the whiteness and uniform density of the quantization noise.

The joint state z_t is defined in (20) and evolves according to the linear dynamics in (21), which defines system matrix F_t. The quantized innovations signal, i_t, is described by (22), which defines output matrix H. The Bayesian filter generates: the conditional density, $p(z_t|\mathbf{I}^t)$, of this $2n_x$-dimensional state, the marginal densities of which yield $p(x_t|\mathbf{I}^t)$, to be used for the optimal controller; and, $p(\check{x}_t|\mathbf{I}^t) = p(\check{x}_t|\mathbf{I}^{t-1})$ according to Theorem 3.

By the same token, we also consider the n_x-dimensional Bayesian filter for the innovations representation of the transmitter-side state estimator,

$$\begin{aligned} \check{x}_{t+1} &= A\check{x}_t + Bu_t + L_t\epsilon_t, \\ i_t &= \mathbf{Q}(\epsilon_t), \end{aligned} \quad (29)$$

to construct directly $p(\check{x}_t|\mathbf{I}^t)$, without the attendant complication of producing $p(x_t|\mathbf{I}^t)$, nor indeed of performing the measurement update step. The conditional density $p(\check{x}_t|\mathbf{I}^t)$ is identical whether produced via the full-order Bayesian filter or it reduced-order counterpart.

Example system
We consider a scalar example quantized innovations system with: values $A = 0.99$, $B = 1$, $C = 1$, $Q = \text{cov}(w_t) = 0.1$, $R = \text{cov}(v_t) = 0.1$; Kalman filter initialization $\hat{x}_{0|-1} = 0$, $\Sigma_{0|-1} = 1.3$. Depending on the signal, ϵ_t or y_t, being quantized, the corresponding steady-state standard deviation, σ_ϵ or σ_y, is computed and the 3-bit/8-level, linear, symmetric, midrise quantizer is used with saturation value at $\zeta_\epsilon = 5\sigma_\epsilon$ or $\zeta_y = 5\sigma_y$ respectively, where σ refers to the stationary variance.

The quantized innovations Bayesian filter, the quantized output Bayesian filter, and the unquantized Kalman filter were computed for a number of steps. The resulting predicted and filtered densities are displayed in Figures 3 and 4. The densities were propagated at 71 sample points in the range [-2, 2].

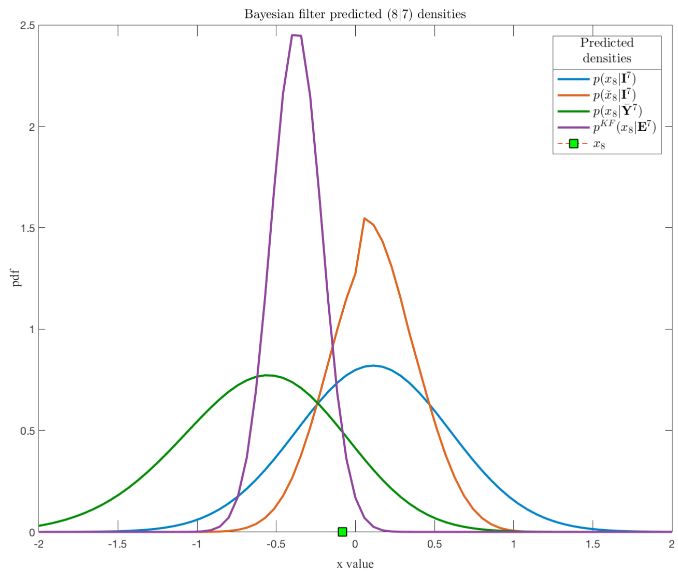

Fig. 3. Predicted (8|7) density functions for: quantized innovations Bayesian filter $p(x_8|\mathbf{I}^7)$ and $p(\breve{x}_8|\mathbf{I}^7)$, quantized output Bayesian filter $p(x_8|\bar{\mathbf{Y}}^7)$, transmitter-side Kalman predictor $p^{KF}(x_8|\mathbf{E}^7)$. Actual plant state x_8 depicted by a green square.

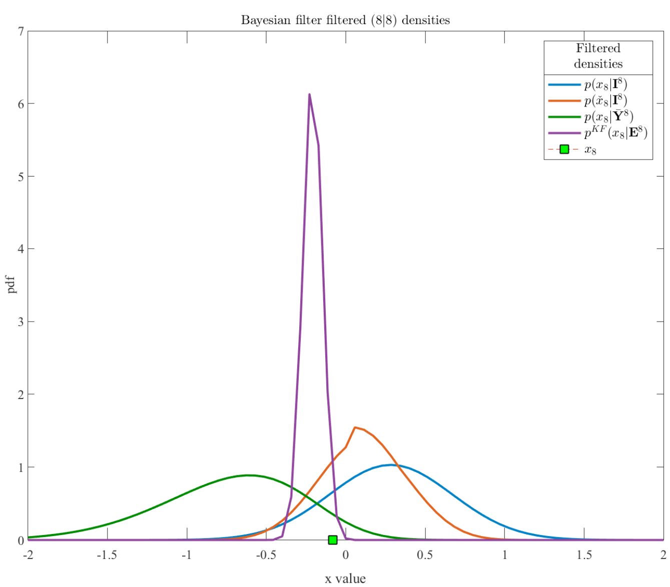

Fig. 4. Filtered (8|8) density functions for: quantized innovations Bayesian filter $p(x_8|\mathbf{I}^8)$ and $p(\breve{x}_8|\mathbf{I}^8)$, quantized output Bayesian filter $p(x_8|\bar{\mathbf{Y}}^8)$, transmitter-side Kalman predictor $p^{KF}(x_8|\mathbf{E}^8)$. Actual plant state x_8 depicted by a green square. Note change of vertical scale versus Figure 3.

We offer the following observations.

- The predicted densities $p(x_8|\mathbf{I}^7)$ and $p(\breve{x}_8|\mathbf{I}^7)$ are different, although their mean values are the same, as guaranteed by Theorem 2.
- The filtered densities $p(x_8|\mathbf{I}^8)$ and $p(\breve{x}_8|\mathbf{I}^8)$ are different, as are their mean values.
- The filtered density $p(\breve{x}_8|\mathbf{I}^8)$ is identical to the predicted density $p(\breve{x}_8|\mathbf{I}^7)$, as guaranteed by Theorem 3.
- The conditional densities of x_t based on quantized innovations are different from those based on quantized output y_t.

VI. Linear Innovations with Dithered Quantizer

We now replace the standard quantizer by a subtractive dithered quantizer, as described in Theorem 1, where now the quantization noise is assumed white and uniformly distributed with zero mean and covariance $\frac{\zeta^2}{3\times 2^{2b}}$. The Kalman filter provides optimal second-order estimation in this case.

We have the above linear transmitter systems (18) and (10) with transmitted data

$$i_t = \epsilon_t + \psi_t, \tag{30}$$
$$= Hz_t + v_t + \psi_t. \tag{31}$$

Immediately, one has the receiver-side Kalman filter recursion from (21)-(30) with usual accommodation of correlated process and measurement noises. Denote

$$\mathcal{F}_t = F_t\bar{P}_tH^T + \begin{bmatrix} 0 \\ L_tR \end{bmatrix}.$$

Then the recursion for the receiver's conditional mean and conditional covariance is:

$$\mu_{t+1} = F_t\mu_t + \mathcal{F}_t(H\bar{P}_tH^T + R + \Psi)^{-1}i_t, \tag{32}$$
$$\bar{P}_{t+1} = F_t\bar{P}_tF_t^T - \mathcal{F}_t(H\bar{P}_tH^T + R + \Psi)^{-1}\mathcal{F}_t^T + \begin{bmatrix} Q & 0 \\ 0 & L_tRL_t^T \end{bmatrix} \tag{33}$$

with initial condition

$$\mu_0 = \begin{bmatrix} \hat{x}_{0|-1} \\ \hat{x}_{0|-1} \end{bmatrix}, \qquad \bar{P}_0 = \begin{bmatrix} P_{0|-1} & 0 \\ 0 & 0 \end{bmatrix}.$$

The conditional joint density at the receiver has

$$\mathrm{E}(z_{t+1}|\mathbf{I}^t) = \mu_t,$$
$$\mathrm{cov}(z_{t+1}|\mathbf{I}^t) = \bar{P}_t.$$

Observations

- The first two conditional moments of $p(\breve{x}_t|\mathbf{I}^{t-1})$ are computed directly from signal model (10) and (30).

$$\hat{\breve{x}}_{t+1} = \mathrm{E}[\breve{x}_{t+1}|\mathbf{I}^t, \pi_{0|-1}],$$
$$= A\hat{\breve{x}}_t + Bu_t + L_ti_t, \tag{34}$$
$$M_{t+1} = \mathrm{cov}(\breve{x}_{t+1}),$$
$$= AM_tA^T + L_t\Sigma_{t|t-1}[I - (\Sigma_{t|t-1} + \Psi)^{-1}\Sigma_{t|t-1}]L_t^T. \tag{35}$$

The resultant conditional density is unique no matter the method of computation.

- The detailed recursion (34)-(35) for $\hat{\breve{x}}_t$ shows that the estimate only adjusts at the time-update step of the Kalman filter, since $p(\breve{x}_t|\mathbf{I}^t) = p(\breve{x}_t|\mathbf{I}^{t-1})$, which in turn is due to the independence of $\breve{x}_t$ and i_t from (34). This is a manifestation of Theorem 3.
- The two dimension-n_x components of the conditional mean are equal per Theorem 2.

$$\mathrm{E}(x_t|\mathbf{I}^{t-1}) = \mathrm{E}(\breve{x}_t|\mathbf{I}^{t-1}).$$

- While the conditional means of x_t and $\breve{x}_t$ at the receiver are identical in this case, their covariances, and thus their complete conditional densities are different.

• For the case with zero channel or quantization noise, ψ_t,

$$\bar{P}_t = \text{blockdiag}(\Sigma_{t|t-1} \;\; 0).$$

In this case, $\check{x}_t$ is reconstructed perfectly at the receiver, since it is a deterministic function of the innovations.

• The stability of matrix A is required for the convergence of conditional mean M_t in (35). This follows Corollary 1.

• We have deliberately ignored the condition $\epsilon_t \in [-\zeta, \zeta]$ from Theorem 1, required to ensure that ψ_t be white. The saturation levels on the quantizer are presumed chosen to enforce this with high probability.

• Additive white channel noise could be included into the analysis in a fashion identical to the quantization noise, subject to the saturation condition.

VII. Comparative optimal control examples

We use the scalar example system from Section V above and consider a sequence of control problems applied to the subtractive dithered quantized system. The aim is to identify circumstances where the control benefits accrue with the availability of the filtered state density.

A. LQG control with dithered quantizer

For the example system presented earlier and LQ cost function

$$J = \lim_{N\to\infty} \mathrm{E}\left\{\frac{1}{N}\sum_{t=0}^{N-1} x_t^T Q_c x_t + u_t^T R_c u_t\right\},$$

with $Q_c = 5$ and $R_c = 0.7$, the performance of three controllers was computed using the LQ-optimal feedback gain and the various conditional mean state estimates.

I. filtered state estimate from quantized innovations $u_t = -K\mathrm{E}(x_t|\mathbf{I}^t)$,
[This is the optimal control by Theorem 5.]

II. predicted state estimate from quantized innovations $u_t = -K\mathrm{E}(x_t|\mathbf{I}^{t-1}) = -K\mathrm{E}(\check{x}_t|\mathbf{I}^{t-1})$,

III. filtered state estimate from quantized outputs $u_t = -K\mathrm{E}(x_t|\bar{\mathbf{Y}}^t)$.

The achieved LQG costs are given in this table.

Control law	value J
$u_t = -K\mathrm{E}(x_t\|\mathbf{I}^t)$	0.9311
$u_t = -K\mathrm{E}(\check{x}_t\|\mathbf{I}^{t-1})$	1.4126
$u_t = -K\mathrm{E}(x_t\|\bar{\mathbf{Y}}^t)$,	1.6712

These quantifications indicate the following.

• The optimal control relies on the use of the filtered density and there is a substantial performance penalty to using the predictive density for this example. This needs to be balanced against the computational cost of operating the full-order Bayesian filter at the receiver.

• There is, for this example, a substantial performance benefit accruing the efficient use of the communications channel through the transmission of quantized innovations signals versus quantized outputs. Again, this comes at a complexity cost in computation at the receiver. But it shows that a control improvement can be realized via careful signal coding. This is generally well understood [14], [15] but is quantified by the example here.

B. Non-LQ optimal control

For the same system, define the performance function

$$\eta_t = \begin{cases} x_t, & \text{if } x_t < 1,\\ 0, & \text{else.}\end{cases}$$

The one-step-ahead control objective is

$$u_t = \arg\max_{u_t} \mathrm{E}\left[\eta_{t+1}|\mathbf{I}^t\right].$$

The solution for the optimal control, given by (42) in the Appendix, is

$$u^{\text{opt}} = 1 - \mathrm{E}(x_{t+1}^{\text{NF}}|\mathbf{I}^t) - \text{xsolv},$$

where xsolv is the solution of (43), an algebraic equation involving solely the Bayesian filter predicted density function $p^{\text{NF}}(x_{t+1}^{\text{NF}}|\mathbf{I}^t)$ for the unforced, $u_t = 0$, state. That is, the optimal control depends on the entire density of the state and not just upon its first moment. The formula (45) for the optimal cost in this case indicates dependence on the covariance.

We present three examples of optimal control of η_t based on the densities: $p(x_{t+1}|\mathbf{I}^t)$, $p(\check{x}_{t+1}|\mathbf{I}^t)$ and $p(x_t|\bar{\mathbf{Y}}^t)$. As seen from Figure 3, these densities differ and each is associated with a different value of the control parameters xsolv. Accordingly, their control performances differ, even though their conditional means might coincide.

Density	xsolv	$\mathrm{E}[\eta_{t+1}\|\mathbf{I}^t]$
$p(x_{t+1}\|\mathbf{I}^t)$	-0.0673	0.1623
$p(\check{x}_{t+1}\|\mathbf{I}^t)$	0.0587	0.1587
$p(x_t\|\bar{\mathbf{Y}}^t)$	-2.8118	0.0108

This reinforces the control performance value of the use of the filtered quantized innovations state density. The xsolv value computed from the $\check{x}$ density is inappropriate, leading to diminished performance. For the quantized output density, the increased variance in the density due to inefficient coding degrades control performance.

VIII. Conclusion & extensions

We have explored the application of the Bayesian filter for control based on predictively coded signals. The predictive coding brings efficiency in the use of the channel bits, which leads to improved state estimation at the receiver and, in turn, to a more accurate state density for control calculation. We have paid particular attention to the generation of the filtered state conditional density, the *information state*, at the controller and identified the inherent performance difference from the predicted state conditional density.

In addition to new theoretical results concerning the state estimation task with predictive coding, the demonstration of computed examples illustrates the feasible but high computational cost of these methods. We analyzed the control problem with dithered quantization which permits precise evaluation of control performance in several cases. Extensions of the computational examples are possible.

- Computation with fully nonlinear and time-varying state and measurement equations, as illustrated in Figure 2, requires some finesse in the following manner.
  - The transmitter-side predictor yielding ξ_t and $\check{y}_t$ needs to be based itself on a nonlinear filter, perhaps even a Bayesian filter.
  - The conditional densities $p(i_t|z_t)$ in (24) need to incorporate the nonlinearity $h_t(\cdot,\cdot,\cdot)$ in an appropriate fashion in addition to the inclusion of the quantizer.
  - The conditional densities $p(z_{t+1}|z_t)$ in (25) need to include the nonlinearity $f_t(\cdot,\cdot,\cdot)$.
  - The innovations sequence no longer need be white. Even in the linear non-Gaussian case, it is uncorrelated but not necessarily white.

  These are standard issues with the application of the Bayesian filter.
- Incorporation of further channel defects such as dropped packets, additive noise, delays are simple extensions of the Bayesian filter. We have already commented on additive channel noise above.
- Practical issues arise when implementing the Bayesian filter. Here, because we have chosen a stationary problem, we have been able to compute the conditional densities on a static grid in the z_t-space. More generally, the Bayesian filter is realized via the Particle filter [33]. This requires some skill.

ACKNOWLEDGEMENT

The authors are keen to acknowledge the technically sound and very helpful comments and guidance from the reviewers.

APPENDIX

A. Proof of Theorem 1

Theorem QTSD of [26], Section 19.8, pp. 506-512, states that, for a linear quantizer with quantization interval q, provided the characteristic function, $\Phi_d(\cdot)$, of the subtractive dither signal satisfies

$$\Phi_d\left(l\frac{2\pi}{q}\right) = 0, \quad \text{for } l = \pm 1,\, \pm 2,\, \ldots, \tag{36}$$

and

$$\epsilon + d \in [-\zeta, \zeta], \tag{37}$$

then the quantization error, $\epsilon - \mathcal{Q}(\epsilon)$, will be independent of the input signal, ϵ, and uniformly distributed $\mathcal{U}(-q/2, q/2)$.

For the linear quantizer of range 2ζ and 2^b levels, $q = \zeta/2^{b-1}$. The characteristic function of a $\mathcal{U}[-a,a]$ density is $\Phi_{d,\text{unif}}(\omega) = \operatorname{sinc} a\omega$. Taking $a = q/2$, $\Phi_{d,\text{unif}}(\omega) = \operatorname{sinc} q\omega/2$, which satisfies (36) above.

The pdf of the sum of two independent $\mathcal{U}(-q/2, q/2)$ random variables is the convolution of the uniform pdfs and is triangularly distributed $\text{tr}(-q, q)$. By the properties of the Fourier transform, $\Phi_{d,\text{tr}}(\omega) = \Phi^2_{d,\text{unif}}(\omega)$ and (36) is satisfied.

We note in passing that Theorem QTSD does not explicitly state the saturation condition (37) on the additively dithered signal. Without it, the theorem fails.

B. Derivation of the Bayesian filters for quantized linear systems

In Section VII, we explore the optimal control performance of three candidate approaches to state conditional density reconstruction at the receiver.

1) Full $2n_x^{\text{th}}$-order quantized innovations Bayesian filter, reconstruction of the conditional state density $p(x_t|\mathbf{I}^t)$, and computation of the optimal control using this density.
2) Simplified n_x^{th}-order quantized innovations Bayesian filter, reconstruction of the conditional state estimate density $p(\check{x}_t|\mathbf{I}^t)$, and computation of the optimal control using this density.
3) The n_x^{th}-order Bayesian filter operating directly on the quantized output signal, $\phi_t = \mathbf{Q}(y_t)$, reconstruction of the conditional density $p(x_t|\mathbf{\Phi}^t)$, and computation of the optimal control using this density.

We now present the detailed Bayesian filter for each case.

C. Bayesian filter for quantized innovations

Measurement update

We begin the Bayesian filter recursion from the predicted density $p(z_t|\mathbf{I}^{t-1})$ with the current measurement i_t in hand. This i_t corresponds to $\epsilon_t \in (\epsilon_{\text{lower}_t}, \epsilon_{\text{upper}_t}]$. Then, from (21)-(22),

$$\begin{aligned} p(i_t|z_t, \mathbf{I}^{t-1}) &= p(i_t|z_t), \\ &= \int_{\epsilon_{\text{lower}_t}}^{\epsilon_{\text{upper}_t}} p(v_t = \epsilon_t - Cx_t + C\check{x}_t)\, dv_t, \\ &= \texttt{mvncdf}(\epsilon_{\text{lower}_t}, \epsilon_{\text{upper}_t}, Hz_t, R). \end{aligned} \tag{38}$$

The MATLAB function `mvncdf` computes the multivariate normal cumulative distribution function between lower and upper limits with given mean and covariance. This is then used in (24) to yield the filtered joint conditional density $p(z_t|\mathbf{I}^t)$.

Time update

For the time update step (25), the system equations (18) and (11) yield

$$\begin{aligned} w_t &= x_{t+1} - Ax_t - Bu_t, \\ L_t v_t &= \check{x}_{t+1} - L_t Cx_t - (A - L_t C)\check{x}_t - Bu_t. \end{aligned}$$

Whence, the conditional density

$$\begin{aligned} p(z_{t+1}|z_t, \mathbf{I}^t) &= \frac{p(z_{t+1}, i_t|z_t, \mathbf{I}^{t-1})}{p(i_t|z_t, \mathbf{I}^{t-1})}, \\ &= \frac{p(z_{t+1}, i_t|z_t, \mathbf{I}^{t-1})}{p(i_t|z_t)}, \end{aligned} \tag{39}$$

since the innovations and quantized innovations, $i_t = \mathbf{Q}(\epsilon_t)$, are white. Denominator $p(i_t|z_t)$ is given by (38). The numerator comprises three terms

$$p(z_{t+1}, i_t|z_t, \mathbf{I}^{t-1}) = W \times V \times T, \tag{40}$$

with

$$\begin{aligned} W &= p(w_t = x_{t+1} - Ax_t - Bu_t),\\ &= \texttt{mvnpdf}(x_{t+1} - Ax_t - Bu_t, 0, Q),\\ V &= p(L_t v_t = \check{x}_{t+1} - L_tCx_t - (A - L_tC)\check{x}_t - Bu_t),\\ &= \texttt{mvnpdf}(\check{x}_{t+1} - L_tCx_t - (A - L_tC)\check{x}_t - Bu_t, 0, L_tRL_t^T),\\ T &= \mathbf{1}(Cx_t - C\check{x}_t + v_t \in (\epsilon_{\text{lower}_t}, \epsilon_{\text{upper}_t}]). \end{aligned}$$

Here, MATLAB function `mvnpdf` is the multivariate normal probability density function and $\mathbf{1}(\cdot)$ is the set indicator function. Relations (38) and (40) comprise the parts of (39) of the time update step (25) of the Bayesian filter.

D. Bayesian filter for state-estimate density calculation

In place of the $2n_x$-dimension Bayesian filter (24)-(25) using relations (38)-(40), we may appeal to (10) and (12) as the basis of an n_x-dimensional Bayesian filter for $p(\check{x}_t|\mathbf{I}^t)$. The system equations are

$$\begin{aligned} \check{x}_{t+1} &= A\check{x}_t + Bu_t + L_t\epsilon_t,\\ i_t &= \mathbf{Q}(\epsilon_t),\\ \check{x}_0 &= \hat{x}_{0|-1}. \end{aligned}$$

The driving noise process, $\{\epsilon_t\}$, is white and Gaussian with the following density

$$\epsilon_t \sim \mathcal{N}(0, C\Sigma_{t|t-1}C + R).$$

Further, this whiteness together with the $\check{x}_t$ update (11) ensures that $\check{x}_t$ is independent from ϵ_t. Thus,

$$\begin{aligned} p(i_t|\check{x}_t) &= p(i_t)\\ &= \texttt{mvncdf}(\epsilon_{\text{lower}_t}, \epsilon_{\text{upper}_t}, 0, C\Sigma_{t|t-1}C^T + R). \end{aligned} \tag{41}$$

Also, similarly to earlier,

$$\begin{aligned} p(\check{x}_{t+1}|\check{x}_t, \mathbf{I}^t) &= \frac{p(\check{x}_{t+1}, i_t|\check{x}_t, \mathbf{I}^{t-1})}{p(i_t|\check{x}_t, \mathbf{I}^{t-1})},\\ &= \frac{p(\check{x}_{t+1}, i_t|\check{x}_t, \mathbf{I}^{t-1})}{p(i_t)}. \end{aligned}$$

E. Bayesian filter for quantized outputs

This now proceeds directly from (1)-(2). Central quantities,

$$\begin{aligned} p(\bar{y}_t|x_t, \bar{\mathbf{Y}}^{t-1}) &= p(\bar{y}_t|x_t),\\ p(x_{t+1}|x_t, \bar{\mathbf{Y}}^t) &= p(x_{t+1}|x_t), \end{aligned}$$

are fully described by, respectively: $h_t(\cdot,\cdot)$ and the density of v_t; and $f_t(\cdot,\cdot,\cdot)$ and the density of w_t.

For the linear systems case,

$$\begin{aligned} p(\bar{y}_t|x_t) &= \texttt{mvncdf}(\epsilon_{\text{lower}_t}, \epsilon_{\text{upper}_t}, Hx_t, R),\\ p(x_{t+1}|x_t) &= \texttt{mvnpdf}(x_{t+1} - Ax_t - Bu_t, 0, Q). \end{aligned}$$

These expressions extend simply for nonlinear system equations involving solely additive noises.

F. Optimal control and value $E[\eta_{t+1}|\mathbf{I}^t]$ for system (7)

The predicted state density generated by the Bayesian filter is $p^{\text{NF}}(x^{\text{NF}}_{t+1}|\mathbf{I}^t)$, the unforced, i.e. $u_t = 0$, state since u_t has yet to be determined. Eventually, $x_{t+1} = x^{\text{NF}}_{t+1} + u_t$. Denote conditional mean $\mu_{t+1} = \text{E}(x^{\text{NF}}_{t+1}|\mathbf{I}^t)$ and define the centered unforced state $x^{\text{c}}_{t+1} = x^{\text{NF}}_{t+1} - \mu_{t+1}$. Then the forced state is decscribed by $x_{t+1} = x^{\text{c}}_{t+1} + \mu_{t+1} + u_t$. Thus,

$$\begin{aligned} E[\eta_{t+1}|\mathbf{I}^t] &= \int_{-\infty}^{1} x_{t+1}\, p(x_{t+1}|\mathbf{I}^t)\, dx_{t+1},\\ &= \int_{-\infty}^{1-\mu_{t+1}-u_t} (x^{\text{c}}_{t+1} + \mu_{t+1} + u_t)\, p^{\text{c}}(x^{\text{c}}_{t+1}|\mathbf{I}^t)\, dx^{\text{c}}_{t+1}\\ &= \int_{-\infty}^{1-\mu_{t+1}-u_t} x^{\text{c}}_{t+1}\, p^{\text{c}}(x^{\text{c}}_{t+1}|\mathbf{I}^t)\, dx^{\text{c}}_{t+1}\\ &\quad + (\mu_{t+1} + u_t) \int_{-\infty}^{1-\mu_{t+1}-u_t} p^{\text{c}}(x^{\text{c}}_{t+1}|\mathbf{I}^t)\, dx^{\text{c}}_{t+1} \end{aligned}$$

Differentiating with respect to u_t,

$$\begin{aligned} \frac{dE[\eta_{t+1}|\mathbf{I}^t]}{du_t} &= (\mu_{t+1} + u_t - 1)\, p^{\text{c}}(1 - \mu_{t+1} - u_t|\mathbf{I}^t)\\ &\quad - (\mu_{t+1} + u_t)\, p^{\text{c}}(1 - \mu_{t+1} - u_t|\mathbf{I}^t)\\ &\quad + \int_{-\infty}^{1-\mu_{t+1}-u_t} p^{\text{c}}(x^{\text{c}}_{t+1}|\mathbf{I}^t)\, dx^{\text{c}}_{t+1}. \end{aligned}$$

Setting this derivative to zero yields the optimal control

$$\begin{aligned} u^{\text{opt}} &= 1 - \mu_{t+1} - \text{xsolv},\\ &= 1 - \text{E}(x^{\text{NF}}_{t+1}|\mathbf{I}^t) - \text{xsolv}, \end{aligned} \tag{42}$$

Where xsolv satisfies

$$p^{\text{NF}}(\text{xsolv}|I_{k-1}) = \int_{-\infty}^{\text{xsolv}} p^{\text{NF}}(x^{\text{NF}}_{t+1}|\mathbf{I}^t) dx^{\text{NF}}_{t+1}. \tag{43}$$

Recall that $p^{\text{NF}}(x^{\text{NF}}_{t+1}|\mathbf{I}^t)$ is the state predicted density produced by the Bayesian filter. Thus xsolv is the point where the probability density function crosses the cumulative distribution function for x^{NF}_{t+1}.

The optimal value function or performance is given by

$$\begin{aligned} \text{E}(\eta_{t+1}|\mathbf{I}^t) &= \int_{-\infty}^{\text{xsolv}} x^{\text{c}}_{t+1}\, p^{\text{c}}(x^{\text{c}}_{t+1}|\mathbf{I}^t)\, dx^{\text{c}}_{t+1}\\ &\quad + (1 - \text{xsolv}) \int_{-\infty}^{\text{xsolv}} p^{\text{c}}(x^{\text{c}}_{t+1}|\mathbf{I}^t)\, dx^{\text{c}}_{t+1}. \end{aligned} \tag{44}$$

If the Bayesian filter predicted state density, $p^{\text{NF}}(x^{\text{NF}}_{t+1}|\mathbf{I}^t)$, is Gaussian $\mathcal{N}(\mu_{t+1}, \sigma^2)$ then

$$\begin{aligned} \text{E}(\eta_{t+1}|\mathbf{I}^t) &= -\frac{\sigma}{\sqrt{2\pi}} \exp\left[-\frac{1}{2\sigma^2}(\text{xsolv})^2\right]\\ &\quad + (1 - \text{xsolv})\, \text{normcdf}(\text{xsolv}, 0, \sigma). \end{aligned} \tag{45}$$

REFERENCES

[1] A. Gersho and R. Gray, *Vector Quantization and Signal Compression*. New York, NY: Kluwer Academic Publishers, 1992.

[2] Telecommunication Standardization Sector, *7 kHz audio-coding within 64 kbit/s*, International Telecommunication Union, September 2012. [Online]. Available: https://www.itu.int/rec/T-REC-G.722-201209-I/en

[3] H. Jones, R. Bitmead, and S. Crisafulli, "Feedback control applied to speech coding," in *IEEE Conference on Decision and Control*, Kobe, Japan, 1996, pp. 1881–1885.

[4] G. Nair and R. Evans, "Stabilizability of stochastic linear systems with finite feedback data rates," *SIAM Journal on Control and Optimization*, vol. 43, no. 2, pp. 413–436, 2004.
[5] K. You, L. Xie, S. Sun, and W. Xiao, "Multiple-level quantized innovation Kalman filter," *Proc. 17th International Federation of Automatic Control World Congress, Seoul, Korea*, 2008.
[6] A. Ribeiro, G. Giannakis, and S. Roumeliotis, "SOI-KF: Distributed Kalman filtering with low-cost communications using the sign of innovations," *IEEE Transactions on Signal Processing*, vol. 54, pp. 4782–4795, 2006.
[7] R. Sukhavasi and B. Hassibi, "The Kalman-like particle filter: Optimal estimation with quantized innovations/measurements," *IEEE Transactions on Signal Processing*, vol. 61, no. 1, pp. 131–136, January 2013.
[8] V. Borkar and S. Mitter, "LQG control with communication constraints," in *Communication, Computation, Control and Signal Processing*, A. Paulraj, V. Roychowdhury, and C. Schaper, Eds. New York NY: Kluwer Academic Publishers, 1997, ch. 21, pp. 365–373.
[9] D. Delchamps, "Extracting state information from quantized output record," *Systems & Control Letters*, vol. 13, pp. 365–372, 1989.
[10] W. S. Wong and R. Brockett, "Systems with finite communication bandwidth constraints. i. state estimation problems," *IEEE Transactions on Automatic Control*, vol. 42, no. 9, pp. 1294 –1299, September 1997.
[11] S.-I. Azuma and T. Sugie, "Optimal dynamic quantizers for discrete-valued input control," *Automatica*, vol. 44, pp. 396–406, 2008.
[12] H. Okajima, K. Sawada, and M. Matsunaga, "Dynamic quantizer design under communication rate constraints," *IEEE Transactions on Automatic Control*, vol. 61, no. 10, pp. 3190–3196, 2016.
[13] T. Fischer, "Optimal quantized control," *IEEE Trans Automatic Control*, vol. 27, no. 4, pp. 996–998, 1982.
[14] M. Fu, "Lack of separation principle for quantized Linear Quadratic Gaussian control," *IEEE Trasactions on Automatic Control*, vol. 57, no. 9, pp. 2385–2390, 2012.
[15] S. Yüksel, "Jointly optimal LQG quantization and control policies for multi-dimensional systems," *IEEE Transactions on Automatic Control*, vol. 59, no. 6, pp. 1612–1617, 2014.
[16] J. Zhang and C. Wang, "On the rate-cost of Gaussian linear control systems with random communication delays," in *IEEE International Symposium on Information Theory*, Vail CO, 2018, pp. 2441–2445.
[17] V. Kostina and B. Hassibi, "Rate-cost tradeoffs in control," in *Fifty-fourth Annual Allerton Conference*, Allerton House IL, September 2016, pp. 1157–1164.
[18] T. Tanaka, K. Johansson, T. Oechtering, H. Sandberg, and M. Skoglund, "Rate of prefix-free codes for LQG control systems," in *IEEE International Symposium on Information Theory*, Barcelona, Spain, 2016, pp. 2399–2403.
[19] P. A.Stravou, J. Østergaard, C. D. Charalambous, and M. Derpich, "An upper bound to zero-delay rate distortion via Kalman filtering for vector Gaussian sources," in *IEEE Information Theory Workshop*, Kaohsiung, Taiwan, 2017.
[20] S. Tatikonda and S. Mitter, "Control under communication constraints," *IEEE Transactions on Automatic Control*, vol. 49, no. 7, pp. 1056 – 1068, July 2004.
[21] P. Stavrou and J. Østergaard, "Fixed-rate zero-delay source coding for stationary vector-valued Gauss-Markov sources,," in *Data Compression Conference*, Snowbird UT, 2018, pp. 257–266.
[22] P. Stavrou, J. Østergaard, and C. Charalambous, "Zero-delay rate distortion via filtering for vector-valued Gaussian sources,," *IEEE Journal of Selected Areas in Signal Processing*, vol. 12, no. 5, pp. 841–856, October 2018.
[23] A. Matveev and A. Savkin, *Estimation and Control over Communication Networks*. Basel: Birkhäuser, 2009.
[24] B. Anderson and J. Moore, *Optimal Filtering*. Englewood Cliffs, New Jersey: Prentice-Hall, 1979.
[25] T. Cover and J. Thomas, *Elements of Information Theory*. New York: John Wiley and Sons Inc., 2006.
[26] B. Widrow and I. Kollár, *Quantization Noise: Roundoff Error in Digital Computation, Signal Processing, Control, and Communications*. Cambridge Univ. Press, 2008.
[27] D. Angeli, "A Lyapunov approach to incremental stability properties," *IEEE Transactions on Automatic Control*, vol. 47, no. 3, pp. 410–421, 2002.
[28] C.-C. Huang and R. R. Bitmead, "Escape time formulation of state estimation and stabilization with quantized intermittent communication," *Automatica*, vol. 61, no. 11, pp. 201–210, November 2015.
[29] B. Anderson and J. Moore, "Detectability and stabilizability of time-varying discrete-time linear systems," *SIAM J. Control Optimization*, vol. 19, no. 1, pp. 20–32, 1981.
[30] R. E. Curry, *Estimation and Control with Quantized Measurements*, ser. Research Monograph. Cambridge MA: MIT Press, 1970, vol. 60.
[31] S. Crisafulli, "Adaptive speech coding via feedback techniques," Ph.D. dissertation, Australian National University, Canberra ACT Australia, 1992.
[32] D. Simon, *Optimal State Estimation: Kalman, H_∞, and nonlinear approaches*. Hoboken NJ: John Wiley & Sons, 2006.
[33] A. Doucet, J. de Freitas, and N. Gordon, *Sequential Monte Carlo Methods in Practice*. New York, NY: Springer-Verlag, 2001.
[34] J. Ferreira, J. Lobo, and J. Dias, "Bayesian real-time perception algorithms on GPU," *Journal of Real-Time Image Processing*, vol. 6, no. 3, pp. 171–186, 2011.
[35] M. Loève, *Probability Theory vols 1 & 2*. Berlin: Springer Verlag, 1977.
[36] P. R. Kumar and P. Varaiya, *Stochastic Systems: Estimation, Identification, and Adaptive Control*. Englewood Cliffs, NJ: Prentice-Hall, 1986.
[37] D. P. Bertsekas, *Dynamic Programming - Deterministic and Stochastic Models*. Englewood Cliffs, NJ: Prentice Hall, 1987.
[38] T. Mita, "Optimal digital feeedback control systems counting computation time of control laws," *IEEE Transactions on Automatic Control*, vol. 30, no. 6, pp. 542–548, 1985.