

Online variable-length source coding for minimum bitrate LQG control

Travis C. Cuvelier, Takashi Tanaka, and Robert W. Heath, Jr.

Abstract—We propose an adaptive coding approach to achieve linear-quadratic-Gaussian (LQG) control with near-minimum bitrate prefix-free feedback. Our approach combines a recent analysis of a quantizer design for minimum rate LQG control with work on universal lossless source coding for sources on countable alphabets. It was recently demonstrated that the aforementioned quantizer’s outputs are an asymptotically stationary, ergodic process. To enable LQG control with provably near-minimum bitrate, the quantizer outputs must be encoded into binary codewords efficiently. This is possible given knowledge of the quantizations’ probability distributions, or of their limiting distribution. Obtaining such knowledge is challenging; the distributions do not readily admit closed form descriptions. This motivates the application of universal source coding. Our main theoretical contribution in this work is a proof that (after an invertible transformation), the quantizer outputs are random variables that fall within an exponential or power-law envelope class (depending on the plant dimension). Using ideas from universal coding on envelope classes, we develop a practical, zero-delay, fixed precision source code for the quantizer outputs. We evaluate the performance of this approach numerically, and demonstrate competitive results with respect to fundamental tradeoffs between bitrate and LQG control performance.

I. INTRODUCTION

In this work, we propose an algorithmic framework to achieve any feasible discrete-time LQG control performance with near minimum bitrate variable-length feedback. Our motivation is feedback control where sensors and actuators are not co-located, and measurements must be fed back from a sensor to a controller over digital communications. In such settings, the feedback bitrate is an effective surrogate for the true *cost* of communication. For many communication architectures, a link’s bitrate corresponds directly to the physical layer resources it uses.

While there are several results in the literature that pertain to LQG control with minimum (variable-length) feedback bitrate [1]–[4], there are barriers to applying these results in practice. This work aims to bridge this gap. We propose practical algorithms that combine the approach to minimum rate LQG control from [4] with approaches to universal coding on countably infinite alphabets from [5]–[8]. A numerically precise implementation of our algorithm follows from a classical approach to arithmetic coding with fixed-precision arithmetic via [9] as well as Elias’s “penultimate”

(omega) encoding for integers [10]. We evaluate the proposed algorithms numerically on a model for an inverted pendulum system [11]. The results demonstrate that our approach achieves good performance with respect to the fundamental directed information bitrate lower bound (cf. [12]).

A. Literature Review

For the single-input linear quadratic regulator problem it was shown in [13] that quantizers with logarithmic level spacing achieve stabilization with a minimal “quantization density”. It is also shown how to achieve stability with a finite number of quantizer levels. The LQG problem we treat in this work is quite different; we consider stochastic plants with unbounded, Gaussian disturbances (which rule out static, finite quantizers, cf. [14]) and aim to attain a fixed quadratic system performance. Furthermore, we consider a “bitrate” notion of communication cost; e.g. quantizations are encoded for transmission over a binary channel.

Data compression architectures for LQG control with near-minimum bitrate variable-length feedback were proposed in [1]–[4]. These approaches propose to quantize sensor measurements with prescribed lattice quantizer designs, and assume that the quantizers’ outputs are efficiently losslessly encoded into (variable-length) binary packets. If this can be accomplished, the architectures in [1]–[4] provably attain bitrates close to established directed information (DI) lower bounds (cf. [3], [12]). To our knowledge, such results has not been obtained with logarithmic quantizers. Accomplishing the efficient lossless encoding in [1]–[4] is quite difficult in practice. In [1]–[4] the quantizer outputs are random variables with countable support, and form a non-stationary random process. In [1]–[3] the lossless encoding is assumed to be adapted precisely to the output distribution. While [4] demonstrated that it suffices to use a time-invariant encoding adapted to the limiting distribution of the quantizer’s outputs; it remains a challenging problem to estimate, and design a code for, this limiting distribution. In contrast to some recent work on LQG control with fixed-rate feedback (cf. e.g. [15]), there are few results on practical algorithms that have been shown to exhibit competitive performance with respect to established fundamental tradeoffs between bitrate and control performance (cf. [3], [12]) in the LQG setting.

Despite recent positive results on LQG control with fixed-length feedback, the less restrictive setting of variable-length feedback offers some advantages. It is well established that a linear plant driven by unbounded process noise cannot be stabilized in the mean square sense with feedback that undergoes time-invariant, memoryless, fixed-length quantization [14]. While several works on control with *adaptive* fixed-

The authors gratefully acknowledge that this work was supported in part by the US Air Force Research Laboratory under grant FA8651-22-1-0013.

T. Cuvelier and T. Tanaka are with the Oden Institute for Computational Engineering and Sciences, The University of Texas at Austin, TX, 78712 USA e-mails: tcuvelier@utexas.edu, ttanaka@utexas.edu.

R. Heath is with the Department of Electrical and Computer Engineering, North Carolina State University, Raleigh, NC, 27606 USA e-mail: rwheath2@ncsu.edu.

length quantization guarantee stability, e.g. [16], [17], they are not guaranteed to operate near the fundamental rate-cost tradeoff for LQG problems. In [15], fixed-length quantizers that optimize LQG control performance were designed via the Lloyd-Max algorithm. While numerical experiments exhibited competitive performance was achieved with respect to fundamental tradeoffs, the approach in [15] requires the Lloyd-Max quantizer to be “re-designed” at every timestep. Stability is not guaranteed. In contrast, it is relatively easy to guarantee a fixed LQG performance with variable-length feedback; rare events that would effectively saturate fixed-length quantizers are encoded into long bitstrings. In this work, we use variable-length feedback, which allows us to give long-term guarantees on control performance.

In this work, we combine the quantizer design and analysis from [4] with relatively recent work on universal lossless source coding on countable alphabets. Broadly speaking, “universal coding” refers to techniques for the design of lossless codes for sources with unknown distributions. The quantizer outputs from [4] are an asymptotically stationary random process with countable support. Their limiting distribution evades closed form description. While there is a wealth of literature devoted to universal coding on alphabets with finite support, there are comparatively fewer results for alphabets with countably infinite support [5]–[8]. In particular, [5] introduced the notion of universal coding on *envelope classes*. An envelope class is a set of probability mass functions on \mathbb{N}_+ (interpreted as potential source distributions) that are upper bounded by a monotonically decreasing, ℓ^1 , *envelope functions*. Relevant envelope functions include power-laws and exponentials. Universal block coding techniques for independent and identically distributed (IID) stationary sources with marginals that fall into envelope classes are introduced in [5]–[8]. The coding algorithms in [5]–[8] take a censoring approach; symbols that are below a (time-varying) cutoff are encoded using arithmetic coding, and large-valued symbol realizations that exceed the cutoff are treated differently. If the symbol value exceeds the cutoff, an escape symbol is encoded arithmetically and the observed symbol value (or the difference between the symbol value and the cutoff) is encoded via one of Elias’s universal prefix codes for integers [10]. The probability model used for arithmetic coding is an essentially an estimator of the source probability mass function. The particular model used in [5]–[8] was investigated in [18]. It is memoryless and empirically adapted to reflect the source.

While the algorithms in [5]–[8] permit sequential encoding and decoding, they are fundamentally block codes. Arithmetic coding sometimes requires that bits from n encoded symbols be received before the first symbol can be decoded. Given a block of n symbols to encode, the algorithms in [5] and [6] can be shown to achieve a redundancy (e.g. a codeword length in excess of the entropy of the source) that has $o(n)$. If, instead of using arithmetic coding, the symbols below the cutoff are encoded with a Shannon-Fano-Elias (SFE) code ([4, Section IV.A.1]) using the empirical source model, the encoding schemes in [5]–[8] can be made to incur

zero encoding and decoding delay at the expense of at most 1 bit per encoded symbol. We discuss [5]–[8] in the context of our present work in the next subsection.

B. Our contributions

This work makes one theoretical and one practical contribution. On the theoretical side, we prove that, after a transformation, the limiting distribution of the quantizer outputs in [4]’s source coding architecture fall into an exponential or power-law envelope class (depending on the plant dimension). The approaches in [5]–[8] apply to *stationary* sources in envelope classes. Given the asymptotic stationarity and the Kullback–Leibler (KL) sense convergence of the quantizer’s outputs to their limiting distribution in [4], we conjecture that using a “zero-delay version” of an algorithm from [5] or [6] to encode the quantizer’s outputs will achieve the same bitrates predicted by [4, Theorem IV.3 (ii)]. The proof of this conjecture is left to future work.

The coding schemes in [5]–[8] are not strictly practical, as they require arithmetic-style prefix coding to be performed on larger and larger alphabets. This requires the (arithmetic) precision with which calculations are performed at the encoder and decoder to likewise expand [9]. We devise a simplified, zero-delay version of the censoring codes using fixed-precision arithmetic. We demonstrate competitive performance with respect to the lower bounds from [12] and the upper bounds in [4].

C. Notation

We use lower case serif letters for vectors v , upper case serif letters for matrices V , and sans serif lower case letters v for scalars. Random variables are in boldface \mathbf{v} . For sequences $\{x_0, x_1, \dots\}$ we let x^n denote $\{x_0, \dots, x_n\}$. If $x \in \mathbb{R}^m$, let $[x]_i$ denote the element of x in entry i (we use zero-based indexing). If \mathbf{x} is a discrete random variable, we let $\mathbb{P}_{\mathbf{x}}(x) = \mathbb{P}_{\mathbf{x}}[\mathbf{x} = x]$. The entropy of the discrete random variable \mathbf{q} , in bits, is denoted $H(\mathbf{q})$. The Kullback–Leibler (KL) divergence between $\mathbb{P}_{\mathbf{a}}$ and $\mathbb{P}_{\mathbf{b}}$, in bits, is denoted $D_{\text{KL}}(\mathbf{a}||\mathbf{b})$. We write $\mathbf{x} \stackrel{D}{=} \mathbf{y}$ if \mathbf{x} and \mathbf{y} are identically distributed. Denote the max singular value of the matrix X via $\|X\|_2$ and the spectral radius (largest magnitude eigenvalue) by $\rho_{\text{max}}(X)$. We use standard notation for vector norms. The set of finite length binary strings is denoted $\{0, 1\}^*$, and the length of $a \in \{0, 1\}^*$, in bits, is $\ell(a)$.

II. SYSTEM MODEL AND PROBLEM FORMULATION

We consider the system model depicted in Figure 1, which depicts a discrete-time, Gauss-Markov plant that is controlled using feedback that is encoded into variable-length binary strings (or packets). Denote the state vector $\mathbf{x}_t \in \mathbb{R}^m$, the control input $\mathbf{u}_t \in \mathbb{R}^u$, and let $\mathbf{w}_t \sim \mathcal{N}(0_m, W)$ denote processes noise assumed to be IID over time. We assume $W \succ 0_{m \times m}$. We assume the initial state has $\mathbf{x}_0 \sim \mathcal{N}(0, X_0)$, for $X_0 \succeq 0$. For some $A \in \mathbb{R}^{m \times m}$ and $B \in \mathbb{R}^{m \times u}$, the plant dynamics for $t \geq 0$ are $\mathbf{x}_{t+1} = A\mathbf{x}_t + B\mathbf{u}_t + \mathbf{w}_t$. To ensure finite control cost is attainable, we assume (A, B) is stabilizable. The plant is fully observable to a sensor/encoder




Fig. 1. The codeword \mathbf{a}_t can be generated randomly conditioned on the information known at the encoder at time t . Likewise, after receiving \mathbf{a}_t , decoder can randomly generate \mathbf{u}_t given \mathbf{a}_t and its previous knowledge. The dither δ_t is generated independently of all past system variables. Under the dynamics, \mathbf{x}_0^t is a deterministic function of \mathbf{x}_0 , \mathbf{a}_0^{t-1} , \mathbf{u}_0^{t-1} , and \mathbf{w}_0^{t-1} . For $t \geq 0$ we assume the factorizations $\mathbb{P}_{\mathbf{a}_{t+1}, \mathbf{u}_{t+1} | \mathbf{a}_0^t, \delta_0^t, \mathbf{u}_0^t, \mathbf{w}_0^t, \mathbf{x}_0} = \mathbb{P}_{E, t+1} \mathbb{P}_{C, t+1}$ and $\mathbb{P}_{\mathbf{a}_{t+1}, \delta_{t+1}, \mathbf{u}_{t+1}, \mathbf{w}_{t+1} | \mathbf{a}_0^t, \delta_0^t, \mathbf{u}_0^t, \mathbf{w}_0^t, \mathbf{x}_0} = \mathbb{P}_{E, t+1} \mathbb{P}_{C, t+1} \mathbb{P}_{\delta} \mathbb{P}_{\mathbf{w}}$, where $\mathbb{P}_{\mathbf{w}}$ and \mathbb{P}_{δ} are the marginal distributions of the process noise and dither. We assume that initially $\mathbb{P}_{\mathbf{a}_0, \delta_0, \mathbf{u}_0, \mathbf{w}_0 | \mathbf{x}_0} = \mathbb{P}_{E, 0} \mathbb{P}_{C, 0} \mathbb{P}_{\delta} \mathbb{P}_{\mathbf{w}}$. These factorizations encode assumed conditional independence assumptions between system variables. The lower bounds in [12] follow as a consequence.

block. At each discrete timestep t , the sensor/encoder makes measurements of the plant, and encodes its measurements into the codeword $\mathbf{a}_t \in \{0, 1\}^*$ via an arbitrary causal encoding policy. The encoder then transmits its codeword \mathbf{a}_t over the feedback channel to a combined decoder/controller. The decoder/controller uses the packets it receives to design the control input, \mathbf{u}_t , again via an arbitrary causal policy. We assume a reliable feedback channel; the decoder receives the packet \mathbf{a}_t without errors and without delay. We assume that the sensor/encoder and decoder/controller have access to common *dither sequence*. The dither is a sequence of exogenous random vectors $\delta_t \in \mathbb{R}^m$ that are revealed causally to both encoder and decoder. We assume that $\{\delta_t\}_i \sim \text{Uniform}([-\frac{\Delta}{2}, \frac{\Delta}{2}])$ IID over i and t . The dither can be used to select codewords/control inputs at the respective blocks. In real-world systems, *shared randomness* of this nature can be approximately accomplished via synchronizing pseudorandom number generators at the encoder and decoder. The encoder/sensor policy in Fig. 1 is a sequence of causally conditioned Borel measurable kernels denoted $\mathbb{P}_{E, t}[\mathbf{a}_0^\infty | \delta_0^\infty, \mathbf{x}_0^\infty] = \left\{ \mathbb{P}_{E, t} = \mathbb{P}_{\mathbf{a}_t | \mathbf{a}_0^{t-1}, \delta_0^t, \mathbf{x}_0^t} \right\}$. The corresponding decoder/controller policy is $\mathbb{P}_{C, t}[\mathbf{u}_0^\infty | \mathbf{a}_0^\infty, \delta_0^\infty] = \left\{ \mathbb{P}_{C, t} = \mathbb{P}_{\mathbf{u}_t | \mathbf{a}_0^t, \delta_0^t, \mathbf{u}_0^{t-1}} \right\}$. Conditional independence assumptions between system variables are imposed via factorizations of the system's transition kernels (cf. Fig. 1).

We require that the codewords produced by the encoder be *prefix-free* in the following sense: for all t and any realizations $(\mathbf{a}_0^{t-1} = a_0^{t-1}, \delta_0^t = \delta_0^t, \mathbf{u}_0^{t-1} = u_0^{t-1})$, for all distinct $a_1, a_2 \in \{0, 1\}^*$ with $\mathbb{P}_{\mathbf{a}_t | \mathbf{a}_0^{t-1}, \delta_0^t, \mathbf{u}_0^{t-1}}[\mathbf{a}_t = a_1 | \mathbf{a}_0^{t-1} = a_0^{t-1}, \delta_0^t = \delta_0^t, \mathbf{u}_0^{t-1} = u_0^{t-1}] > 0$ and $\mathbb{P}_{\mathbf{a}_t | \mathbf{a}_0^{t-1}, \delta_0^t, \mathbf{u}_0^{t-1}}[\mathbf{a}_t = a_2 | \mathbf{a}_0^{t-1} = a_0^{t-1}, \delta_0^t = \delta_0^t, \mathbf{u}_0^{t-1} = u_0^{t-1}] > 0$, a_1 is not a prefix of a_2 . This assumption ensures that the decoder can instantaneously decode the codeword at time t . However, it permits the encoder to use different prefix-free codebooks for different realizations of $(\mathbf{a}_0^{t-1}, \delta_0^t, \mathbf{u}_0^{t-1})$ (cf. [4, Prefix Constraint 1] [12, Assumption 2]).

We are interested in the tradeoff between LQG control performance and communication cost, quantified by the time-average expected codeword length. Mathematically, this is

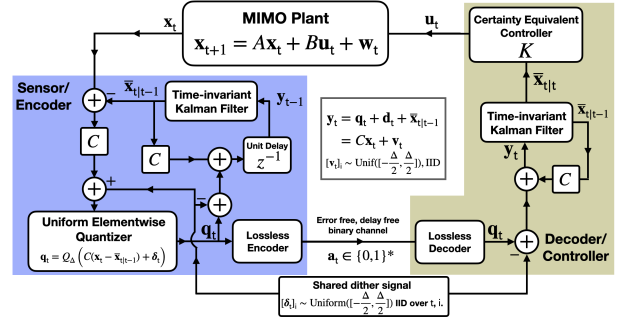


Fig. 2. The quantizer and controller design from [4].

described by the optimization over policies $\mathbb{P}_E, \mathbb{P}_C$ that conform to the prefix constraint:

$$\mathcal{L}(\gamma) = \begin{cases} \inf_{\mathbb{P}_E, \mathbb{P}_C} \limsup_{T \rightarrow \infty} \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}[\ell(\mathbf{a}_t)] \\ \text{s.t. } \limsup_{T \rightarrow \infty} \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}[\|\mathbf{x}_{t+1}\|_Q^2 + \|\mathbf{u}_t\|_\Phi^2] \leq \gamma, \end{cases}$$

where $Q \succeq 0$, $\Phi \succ 0$, and γ is the maximum tolerable LQG cost. Let S be a stabilizing solution to the discrete algebraic Riccati equation (DARE) $A^T S A - S - A^T S B (B^T S B + \Phi)^{-1} B^T S A + Q = 0$, $K = -(B^T S B + \Phi)^{-1} B^T S A$, and $\Theta = K^T (B^T S B + \Phi) K$. Consider the optimization

$$\mathcal{R}(\gamma) = \begin{cases} \inf_{\substack{P, \Pi \in \mathbb{R}^m \times \mathbb{R}^m \\ P, \Pi \succeq 0}} \frac{1}{2} (-\log_2 \det \Pi + \log_2 \det W) \\ \text{s.t. } \text{Tr}(\Theta P) + \text{Tr}(W S) \leq \gamma, \\ P \preceq A P A^T + W, \\ \begin{bmatrix} P - \Pi & P A^T \\ A P & A P A^T + W \end{bmatrix} \succeq 0 \end{cases} \quad (1)$$

$\mathcal{R}(\gamma)$ is the minimum DI that must be incurred by any policy that achieves an LQG cost less than or equal to γ . This DI is a fundamental lower bound on the bitrate $\mathcal{L}(\gamma)$, i.e. $\mathcal{R}(\gamma) \leq \mathcal{L}(\gamma)$ [12, Theorem 1] [19]. The minimizing P from (1), denoted \hat{P} , can be used to construct encoder and decoder policies that nearly achieve the lower bound $\mathcal{R}(\gamma)$ [4]. This approach, depicted in Fig. 2, is presently summarized.

Given \hat{P} , define $\hat{P}_+ = A \hat{P} A^T + W$. Assume that $\Delta = 1$, and that C is chosen such that $\frac{1}{12}(\hat{P}^{-1} - \hat{P}_+^{-1}) = C^T C$. The encoder and decoder operate synchronized time-invariant Kalman filters (KF) that are updated with measurements constructed from the encoded codewords. Let $\bar{\mathbf{x}}_{t|t-1}$ denote these filters' a priori estimate, and $\bar{\mathbf{x}}_{t|t}$ their posterior estimate. Assume initially that $\bar{\mathbf{x}}_{1|0} = 0$. By construction, both the encoder and decoder will have computed the same a priori estimate $\bar{\mathbf{x}}_{t|t-1}$ by the beginning of timestep t . Define the Kalman filter's a priori error via $\mathbf{e}_t = \mathbf{x}_t - \bar{\mathbf{x}}_{t|t-1}$. Let $Q_\Delta : \mathbb{R}^m \rightarrow \mathbb{Z}^m$ be a function that rounds its input elementwise to the nearest integer, e.g. $[Q_\Delta(x)]_i = k$ if $[x]_i \in [k - \frac{1}{2}, k + \frac{1}{2})$. At every time t the encoder produces a *dithered quantization* of the Kalman innovation via $\mathbf{q}_t = Q_\Delta(C\mathbf{e}_t + \delta_t)$. The quantization \mathbf{q}_t is a discrete random variable with countably infinite support on \mathbb{Z}^m . The quantization \mathbf{q}_t is encoded into

the codeword \mathbf{a}_t losslessly. Since the decoder receives \mathbf{a}_t without error, it can reconstruct \mathbf{q}_t exactly. The decoder then computes a centered, reconstructed measurement $\mathbf{y}_t = \mathbf{q}_t - \delta_t + C\bar{\mathbf{x}}_{t|t-1}$. It is shown in [4, Eqn. (19)] that $\mathbf{y}_t = C\mathbf{x}_t + \mathbf{v}_t$, where $[\mathbf{v}_t]_i \sim \text{Uniform}([-1/2, 1/2])$ IID over i and $\mathbf{v}_t \perp \mathbf{x}_0^t$. Since the encoder knows the dither signal, it can also compute \mathbf{y}_t . Both the encoder and decoder update their KFs using the time-invariant gain $J = \hat{P}_+ C^T (C\hat{P}_+ C^T + I_m)^{-1}$, computing $\bar{\mathbf{x}}_{t|t} = \bar{\mathbf{x}}_{t|t-1} + J(\mathbf{y}_t - C\bar{\mathbf{x}}_{t|t-1})$. The decoder then applies certainty equivalent control, selecting $\mathbf{u}_t = K\bar{\mathbf{x}}_{t|t}$. The encoder can likewise compute \mathbf{u}_t , and both the encoder and decoder KFs compute the predict update via $\bar{\mathbf{x}}_{t+1|t} = A\bar{\mathbf{x}}_{t|t} + B\mathbf{u}_t$. The next proposition highlights some relevant results from [4] that pertain to this architecture.

Proposition 2.1: Let $L = AJ$ and $R = A - LC$. Let $\{\nu_t\}$ denote an IID sequence of random variables uniformly distributed on $[-1/2, 1/2]^m$, let $\{\omega_t\}$ be IID with $\omega_t \sim \mathcal{N}(0_m, W)$, and let $\chi \sim \mathcal{N}(0_m, X_0)$. Let $\{\omega_t\}$, $\{\nu_t\}$, and χ be mutually independent. We have that R is (discrete-time) globally asymptotically stable, e.g. $\rho_{\max}(R) < 1$, that

$$\mathbf{e}_t \stackrel{D}{=} R^t \chi + \sum_{i=0}^{t-1} R^i (\omega_i - L\nu_i), \quad (2)$$

and that there exists a random variable $\mathbf{e} \sim \mathbb{P}_e$ such that (\mathbf{e}_t, δ_t) converge in total variation (and thus weakly) to (\mathbf{e}, δ) where $\delta \perp \mathbf{e}$ and $\delta \sim \text{Uniform}([-1/2, 1/2]^m)$. Furthermore $\mathbf{e} \in \mathbb{R}^m$ has a strictly positive probability density function and the process (\mathbf{e}_t, δ_t) is ergodic in the sense that that if $f : \mathbb{R}^m \times [-1/2, 1/2]^m \rightarrow \mathbb{R}$ has $\mathbb{E}[|f(\mathbf{e}, \delta)|] < \infty$ then

$$\lim_{T \rightarrow \infty} \frac{1}{T} \sum_{i=0}^{T-1} f(\mathbf{e}_i, \delta_i) = \mathbb{E}[f(\mathbf{e}, \delta)], \text{ almost surely.} \quad (3)$$

Defining $\mathbf{q} = Q_\Delta(C\mathbf{e} + \delta)$, we have that the \mathbf{q}_t converge in total variation (and thus weakly) to \mathbf{q} and furthermore that $\lim_{t \rightarrow \infty} D_{\text{KL}}(\mathbf{q}_t | \mathbf{q}) = 0$. Finally, we have that both $\limsup_{t \rightarrow \infty} H(\mathbf{q}_t) \leq R(\gamma) + b$ and $H(\mathbf{q}) \leq R(\gamma) + b$ where $b = m(1 + \frac{1}{2} \log_2(\frac{2\pi e}{12}))$. So long as the \mathbf{q}_t are losslessly conveyed to the decoder, LQG control cost satisfies

$$\limsup_{T \rightarrow \infty} \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}[\|\mathbf{x}_{t+1}\|_Q^2 + \|\mathbf{u}_t\|_\Phi^2] \leq \gamma. \quad (4)$$

Proof: That $\rho_{\max}(R) < 1$ is addressed in [4, Section IV.C], (2) is addressed in the second paragraph of the proof of [4, Lemma 4.8]. The existence and ergodic properties of the limiting measures $\mathbb{P}_{e,\delta}$ and \mathbb{P}_q are main results of [4]. These results are discussed in [4, Theorem IV.3 and Section IV.C, in particular Lemma IV.5]. The KL-sense convergence of \mathbf{q}_t to \mathbf{q} is [4, Lemma IV.8]. The entropy bounds are via [4, Lemma IV.7] and (4) is [4, Theorem IV.3 (iii)]. ■

In the encoder/decoder in Fig. 2, at every timestep \mathbf{q}_t is encoded into the prefix-free codeword $\mathbf{a}_t \in \{0, 1\}^*$. The encoding is lossless and zero delay. Upon receiving \mathbf{a}_t , decoder recovers \mathbf{q}_t . The remainder of this work pertains to this lossless encoding. Let $\mathcal{C}_t : \mathbb{Z}^m \rightarrow \{0, 1\}^*$ be the encoding function used at time t , e.g. assume $\mathbf{a}_t = \mathcal{C}_t(\mathbf{q}_t)$. If the probability mass function $\mathbb{P}_{\mathbf{q}_t}$ is known at every t , \mathcal{C}_t can be chosen as a SFE code adapted to the distribution

$\mathbb{P}_{\mathbf{q}_t}$ and achieve a codeword length $\mathbb{E}[\ell(\mathbf{a}_t)] \leq H(\mathbf{q}_t) + 1$ (cf. e.g. [4, Sec. IV.A.1]). In [4], we proved that if the limiting distribution \mathbb{P}_q is known and \mathcal{C}_t is chosen as a fixed (time-invariant) SFE code adapted to the distribution $\mathbb{P}_{\mathbf{q}_t}$, then $\limsup_{t \rightarrow \infty} \mathbb{E}[\ell(\mathbf{a}_t)] \leq \limsup_{t \rightarrow \infty} H(\mathbf{q}_t) + 1$. Thus, if *either* the marginal probability mass functions (PMFs) $\mathbb{P}_{\mathbf{q}_t}$ or the limiting PMF \mathbb{P}_q is known, we can losslessly encode $\{\mathbf{q}_t\}$ so that the prefix constraint is satisfied and

$$\limsup_{T \rightarrow \infty} \frac{1}{T} \sum_{i=0}^{T-1} \mathbb{E}[\ell(\mathbf{a}_i)] \leq \mathcal{R}(\gamma) + b + 1, \quad (5)$$

e.g. the system achieves a time-average codeword length that is at most b bits above the fundamental rate-distortion lower bound in (1). While the bound in (5) can be achieved if either the $\mathbb{P}_{\mathbf{q}_t}$ or \mathbb{P}_q are known, this is difficult to accomplish in practice. It is doubtful that they admit closed form characterizations. This motivates pursuing designs for \mathcal{C}_t via universal coding. In our present context universal coding functions will be constructed via past observations of the source process. To make this explicit, we will denote $\mathcal{C}_t = \mathcal{C}_{t|\mathbf{q}^{t-1}}$ when \mathcal{C}_t is a function of the ‘‘past realizations’’ \mathbf{q}^{t-1} . We discuss universal coding in the next section.

III. TOWARDS UNIVERSAL CODING FOR THE QUANTIZER OUTPUTS

We begin with some definitions. A probability mass function $\mathbb{P}_z : \mathbb{N}_+ \rightarrow [0, 1]$ belongs to the *power-law envelope class* with parameters (α, β) if for $\alpha > 1$, $\beta > 2^\alpha / (\sum_{i=1}^\infty \frac{1}{i^\alpha})$, $\mathbb{P}_z(x) \leq \min(\frac{\beta}{x^\alpha}, 1)$ for $x \in \mathbb{N}_+$ [5]. Likewise, a probability mass function $\mathbb{P}_z : \mathbb{N}_+ \rightarrow [0, 1]$ belongs to the *exponential envelope class* with parameters (α, β) if for $\alpha > 0$, $\beta > e^{2\alpha}$ such that $\mathbb{P}_z(x) \leq \min(\beta e^{-\alpha x}, 1)$, for $x \in \mathbb{N}_+$. Finally, a (scalar) random variable \mathbf{x} is *σ -subgaussian* if $\mathbb{E}[\mathbf{x}] = 0$ and $\mathbb{E}[e^{\lambda \mathbf{x}}] \leq e^{\lambda^2 \sigma^2 / 2}$ for all $\lambda \in \mathbb{R}$. A random vector $\mathbf{x} \in \mathbb{R}^m$ is *σ -subgaussian* if for every $c \in \mathbb{R}^m$ with $\|c\|_2 = 1$, $c^T \mathbf{x}$ is σ -subgaussian.

Let \mathbf{x}_t be a stationary, IID source on \mathbb{N}_+ . Assume $\mathbf{x}_t \sim \mathbb{P}_x$ for all t . Let $\mathbf{x} \sim \mathbb{P}_x$. Assume that \mathbb{P}_x falls into either a power-law or exponential envelope class. A consequence of [5] and [6] is that one can construct a sequence of functions $\mathcal{C}_{t|\mathbf{x}^{t-1}} : \mathbb{N}_+ \rightarrow \{0, 1\}^*$ such that

$$\limsup_{T \rightarrow \infty} \frac{1}{T} \sum_{i=0}^{T-1} \mathbb{E}[\ell(\mathcal{C}_{t|\mathbf{x}^{t-1}}(\mathbf{x}_t))] \leq H(\mathbf{x}) + 1. \quad (6)$$

Further, if $\mathbb{P}_x[\mathbf{x}_t = x_1 | \mathbf{x}^{t-1}]$, $\mathbb{P}_x[\mathbf{x}_t = x_2 | \mathbf{x}^{t-1}] > 0$ then $\mathcal{C}_{t|\mathbf{x}^{t-1}}(x_1)$ is not a prefix of $\mathcal{C}_{t|\mathbf{x}^{t-1}}(x_2)$ and vice-versa. This follows from replacing arithmetic coding with SFE coding (under the same probability model) in the encoding algorithms for power-law ([5]) or exponential enveloped sources ([6]). Each $\mathcal{C}_{t|\mathbf{x}^{t-1}}$ is constructed using the previous observations \mathbf{x}^{t-1} without knowledge of the PMF \mathbb{P}_x .

The envelope classes, and encoding schemes in [5]–[8] apply to sources on \mathbb{N}_+ , however, the quantizer output discussed in Section II has $\mathbf{q}_t \in \mathbb{Z}^m$. We will develop a bijection $g : \mathbb{Z}^m \rightarrow \mathbb{N}_+$ such that the limiting distribution of

$g(\mathbf{q}_t)$ falls into an envelope class. In particular, assume that g is such that if $\|a\|_\infty > \|b\|_\infty$ then $g(a) > g(b)$. In other words, let g be any bijection that respects the partial order induced by the infinity norm on \mathbb{Z}^m . Clearly such a mapping exists. For $i \in \mathbb{N}$, define $\mathcal{B}_i = \{z \in \mathbb{Z}^m : \|z\|_\infty \leq i\}$, which is the set of points in \mathbb{Z}^m that lie within an origin-centered hypercube with edges of length $2i$. Note that the cardinality of \mathcal{B}_i is $|\mathcal{B}_i| = (2i+1)^m$. To define g , choose $g(0_m) = 1$, and then map the points in $\mathcal{B}_1/\mathcal{B}_0$ arbitrarily to 2 to 3^m and so on. We have the following theorem.

Theorem 3.1: Let $g : \mathbb{Z}^m \rightarrow \mathbb{N}_+$ be any bijection such that if $\|a\|_\infty > \|b\|_\infty$ then $g(a) > g(b)$. Let \mathbf{q} be as described in Proposition 2.1, and define $\bar{\mathbf{q}} = g(\mathbf{q})$. If $m = 1$ or $m = 2$, $\bar{\mathbf{q}}$ falls into the exponential envelope class. If $m > 2$, have that $\bar{\mathbf{q}}$ is a member of a power-law envelope class.

Proof: Let $\{\omega_t\}$, $\{\nu_t\}$, and χ be as in Prop. 2.1, and let $\delta_t \perp\!\!\!\perp (\chi, \{\omega_t\}, \{\nu_t\})$. Define $\mathbf{z}_t = C\mathbf{e}_t + \delta_t$. Via (2),

$$\mathbf{z}_t \stackrel{D}{=} C \left(R^t \chi + \sum_{i=0}^{t-1} R^i (\omega_i - L\nu_i) \right) + \delta_t. \quad (7)$$

Take $c \in \mathbb{R}^m$ with $\|c\|_2 = 1$. Let $\kappa_1 = \lim_{t \rightarrow \infty} \sum_{i=0}^{t-1} \|R^i\|_2$, $\kappa_2 = \sup_t \|R^t\|_2^2$, and $\kappa_3 = \lim_{t \rightarrow \infty} \sum_{i=0}^{t-1} \|R^i\|_2^2$. Note that since $\rho_{\max}(R) < 1$, $\kappa_1, \kappa_2, \kappa_3$ are finite via Gelfand's theorem (cf. e.g. Proposition A.4 in [4]). Define $\mathbf{z}_t = -C \sum_{i=0}^{t-1} R^i L\nu_i + \delta_t$ and let $\bar{\mathbf{z}}_t = C(R^t \chi + \sum_{i=0}^{t-1} R^i \omega_i)$ so that $\mathbf{z}_t \stackrel{D}{=} \underline{\mathbf{z}}_t + \bar{\mathbf{z}}_t$ with $\mathbb{E}[\bar{\mathbf{z}}_t] = \mathbb{E}[\underline{\mathbf{z}}_t] = 0$. For all t , $c^T \underline{\mathbf{z}}_t$ has bounded support, i.e.

$$|c^T \underline{\mathbf{z}}_t| \leq (\kappa_1 \|C\|_2 \|L\|_2 + 1) \frac{\sqrt{m}}{2} \quad (8)$$

which follows via the triangle inequality, Cauchy-Schwartz, and the submultiplicativity of matrix norms, and since $\|\nu_t\|_2, \|\delta_t\|_2 \leq \frac{\sqrt{m}}{2}$. Let $\sigma_1 = (\kappa_1 \|C\|_2 \|L\|_2 + 1) \frac{\sqrt{m}}{2}$. Given (8) we have that for every t and c with $\|c\|_2 = 1$, $c^T \underline{\mathbf{z}}_t$ is σ_1 -subgaussian [20, Ex. 5.6 (b)].

Define $\Omega_t = C \left(R^t X_0 (R^t)^T + \sum_{i=0}^{t-1} R^i W (R^i)^T \right) C^T$. We have $c^T \bar{\mathbf{z}}_t \sim \mathcal{N}(0, c^T \Omega_t c)$. If we denote $\sigma_2^2 = \|C\|_2^2 (\|X_0\|_2 \kappa_2 + \|W\|_2 \kappa_3)$, for all t we have $c^T \Omega_t c \leq \sigma_2^2$, which follows analogously by the triangle inequality, Cauchy-Schwartz, the submultiplicativity of matrix norms. Since $c^T \bar{\mathbf{z}}_t$ is a zero-mean Gaussian with a variance upper bounded by σ_2^2 we have that $c^T \bar{\mathbf{z}}_t$ is σ_2 -subgaussian. Note that σ_2 does not depend on c , and that this holds for any c with $\|c\|_2 = 1$. Since $c^T \mathbf{z}_t \stackrel{D}{=} c^T \underline{\mathbf{z}}_t + c^T \bar{\mathbf{z}}_t$, via [20, Lemma 5.4 (b)] $c^T \mathbf{z}_t$ is $\sigma = \sqrt{\sigma_1^2 + \sigma_2^2}$ -subgaussian for all t and c with $\|c\|_2 = 1$. Thus the \mathbf{z}_t is σ -subgaussian for all t .

Define $\mathbf{q}_t = Q_\Delta(\mathbf{z}_t)$ and $\bar{\mathbf{q}}_t = g(\mathbf{q}_t)$. Note that for $r \in \mathbb{N}_0$ by definition the set \mathcal{B}_r contains $|\mathcal{B}_r| = (2r+1)^m$ points. By definition of the bijection g , for $r \in \mathbb{N}_+$, $\|\mathbf{q}_t\|_\infty \geq r$ if and only if $\bar{\mathbf{q}}_t \geq (2(r-1)+1)^m + 1$. Take $z \in \mathbb{N}_+$. We have

$$\begin{aligned} \mathbb{P}[\bar{\mathbf{q}}_t = z] &\leq \mathbb{P}[\bar{\mathbf{q}}_t \geq z] & (9) \\ &\leq \mathbb{P} \left[\bar{\mathbf{q}}_t \geq \left(2 \left(\lfloor \frac{\sqrt[3]{z-1}+1}{2} \rfloor - 1 \right) + 1 \right)^m + 1 \right] \\ &\leq \mathbb{P} \left[\|\mathbf{q}_t\|_\infty \geq \lfloor \frac{\sqrt[3]{z-1}+1}{2} \rfloor \right], & (10) \end{aligned}$$

where the $z = 1$ case in (10) holds trivially and for $z > 1$ (10) follows from our observation above (9). Note that for $r \in \mathbb{N}_0$, $\|\mathbf{q}_t\|_\infty \geq r$ if and only if $\|C\mathbf{e}_t + \delta_t\|_\infty \geq r - \frac{1}{2}$. Thus, further relaxing (10)

$$\begin{aligned} \mathbb{P}[\bar{\mathbf{q}}_t = z] &\leq \mathbb{P} \left[\|C\mathbf{e}_t + \delta_t\|_\infty \geq \lfloor \frac{\sqrt[3]{z-1}+1}{2} \rfloor - \frac{1}{2} \right] \\ &\leq \mathbb{P} \left[\|C\mathbf{e}_t + \delta_t\|_\infty \geq \left(\frac{\sqrt[3]{z-1}}{2} - 1 \right) \right] & (11) \\ &\leq 2m e^{-\frac{1}{2\sigma^2} \left(\frac{1}{2} \sqrt[3]{z-1} - 1 \right)^2} & (12) \end{aligned}$$

where (12) follows from a maximal inequality for subgaussian random vectors (cf. [21, Theorem 2.2.1]). Let $\zeta = 1/(8\sigma^2)$. We have $z^2 \leq e^{\zeta (\sqrt[3]{z-1}-2)^2}$ for z sufficiently large, and thus, $\mathbb{P}[\bar{\mathbf{q}}_t = z] \leq \frac{2m}{z^2}$ for z sufficiently large. This demonstrates that irrespective of m , the quantizer output, when wrapped by the function g , falls into a power-law envelope class with parameter $\alpha = 2$ and a β that depends on ζ and m . If $m = 1$, once we have immediately that $\bar{\mathbf{q}}_t$ is in the exponential class with $\alpha = \zeta$ and a β that depends on ζ . If $m = 2$, $\bar{\mathbf{q}}_t$ is in the exponential class with $\alpha = \frac{\zeta}{2}$ and a β that depends on ζ . Since the bound in (12) holds for all t , it holds for the limiting distribution $\bar{\mathbf{q}} = g(\mathbf{q})$. ■

Note that $H(\mathbf{q}_t) = H(\bar{\mathbf{q}}_t)$ and $H(\mathbf{q}) = H(\bar{\mathbf{q}})$ since g is a bijection. By Proposition 2.1 $D_{\text{KL}}(\bar{\mathbf{q}}_t \| \bar{\mathbf{q}}) \rightarrow 0$. A corollary to this latter fact and Theorem 3.1 is that $\lim_{t \rightarrow \infty} H(\mathbf{q}_t) = H(\mathbf{q})$ [22, Theorem 21]. Notably, the results in [5] and [6] hold only for *stationary sources* on \mathbb{N}_+ , meanwhile $\bar{\mathbf{q}}_t$ is only asymptotically stationary. However, given Proposition 2.1 and Theorem 3.1 we conjecture the following.

Conjecture 3.2: Let g be a bijection as described in Theorem 3.1. Assume that the lossless encoder in Fig. 2 first computes $\bar{\mathbf{q}}_t = g(\mathbf{q}_t)$, and then encodes $\bar{\mathbf{q}}_t$ via $\mathbf{a}_t = \mathcal{C}_{t|\bar{\mathbf{q}}^{t-1}}(\bar{\mathbf{q}}_t)$, where $\mathcal{C}_{t|\bar{\mathbf{q}}^{t-1}}$ is the encoding function constructed via either [6] (in the case that $m \leq 2$) or [5] with the zero-delay modification (replacing arithmetic coding with SFE coding, cf. (6)). The encoding is prefix-free in the sense of Section II (indeed, particular prefix-free encoding used at time t depends only on $\bar{\mathbf{q}}^{t-1}$), and the decoder can reconstruct \mathbf{q}_t exactly, thus ensuring that the constraint on LQG control performance is satisfied. We conjecture that the time-average expected codeword lengths will satisfy $\lim_{T \rightarrow \infty} \frac{1}{T} \sum_{i=0}^{T-1} \mathbb{E}[\ell(\mathbf{a}_t)] = H(\mathbf{q}) + 1$.

Mirroring the approaches in [5], [6], Theorem 3.1 provides the first step required to affirmatively establish Conjecture 3.2. While proving (or disproving) Conjecture 3.2 is a topic for future work; we suspect that it holds given the ergodicity of $\{\mathbf{q}_t\}$ and since $\mathbf{q}_t \rightarrow \mathbf{q}$ in the KL-sense. A proof of Conjecture 3.2 would require analyzing the lengths of both the Shannon-Fano-Elias style encoding (cf. [5, Lemma 2]) and the encoding of "overflows" (cf. [5, Lemma 3]). In the next section, we depart from theory and propose a practical algorithm, based on [5], [6], for the lossless encoding of the quantizer outputs. The proposed algorithm demonstrates competitive performance in Section V.

IV. A PRACTICAL ADAPTIVE ALGORITHM

A key element of the coding schemes in [5]–[8] is cutoff censoring. Generally speaking, if a source produces a symbol below a cutoff at a given time, the symbol is encoded via adaptive arithmetic coding. If the source produces a symbol that exceeds the cutoff, an escape symbol is encoded arithmetically, and the overflowing symbol itself was encoded using the Elias omega code. In [5]–[8], the cutoffs grow with time, and (in our case) will diverge as large numbers of symbols are encoded. Cutoffs that grow with time are impractical as they require that the arithmetic precision used to implement arithmetic, or, in the zero-delay case, SFE, encoders and decoders similarly expands over time [9]. Since our interest is in long term, infinite horizon communication cost, we propose instead to fix the cutoffs a priori and account for fixed arithmetic precision. We describe our proposed algorithm in the remainder of this section.

We first describe our notion of “censoring” before describing the encoding. Let p be the precision, in bits, in which (unsigned integer) arithmetic operations are to be performed at the encoder and decoder. We first transform \mathbf{q}_t into a source on \mathbb{N}_+^m by computing the vector \mathbf{s}_t with elements

$$[\mathbf{s}_t]_i = \begin{cases} 2[\mathbf{q}_t]_i, & [\mathbf{q}_t]_i > 0 \\ -2[\mathbf{q}_t]_i + 1, & [\mathbf{q}_t]_i \leq 0 \end{cases}. \quad (13)$$

Let $k \in \mathbb{N}_+^m$ be a vector of *cutoffs*. The k are fixed a priori; and are hyperparameters of our encoding algorithm. We require that $n = \prod_{j=1}^m ([k]_j + 1)$ has $n < 2^{\frac{p}{2}}$. Define the truncation operator $\text{trunk}_k : \mathbb{N}_+^m \rightarrow \mathbb{N}_0^m$ via

$$[\text{trunk}_k(s)]_i = \begin{cases} [s]_i, & [s]_i \leq [k]_i \\ 0, & \text{otherwise} \end{cases}. \quad (14)$$

Define the post-truncation symbol tuple $\bar{\mathbf{s}}_t = \text{trunk}_k(\mathbf{s}_t)$, which is a source on an alphabet of cardinality n . Denote the sequence of symbols that were truncated $\hat{\mathbf{s}}_t$, so that $[\hat{\mathbf{s}}_t]_i$ is the i^{th} symbol truncated from \mathbf{s}_t . Note that the dimension of the vector $\hat{\mathbf{s}}_t$ is random. Define the *linear indexing* function $\lambda : \mathbb{N}_0^m \rightarrow \mathbb{N}_0$ via $\lambda(s) = \sum_{i=1}^{m-1} ([s]_i \prod_{j=i+1}^m ([k]_j + 1) + [s]_m)$. We have that λ is a bijection from the range of \mathbf{s}_t to the set $\{0, 1, \dots, n-1\}$. Let $\mathbf{m}_t = \lambda(\bar{\mathbf{s}}_t)$. In our algorithm, we encode \mathbf{m}_t using a SFE code (cf. [23], [4, Section IV.A.1]). Subsequently, we encode each element of $\hat{\mathbf{s}}_t$ with the Elias omega code [10]. The decoder decodes \mathbf{m}_t , recovering $\bar{\mathbf{s}}_t$ via inverting the indexing function λ . It then counts the overflow “0” symbols in $\bar{\mathbf{s}}_t$, and decodes the omega-encoded $\hat{\mathbf{s}}_t$ from the remaining bits. Finally, it reconstructs \mathbf{s}_t and \mathbf{q}_t .

We implement the SFE encoding along the same lines as the fixed-precision implementation of adaptive arithmetic coding in [9]. To encode \mathbf{m}_t via an SFE codec, a probability mass function for \mathbf{m}_t is required. We encode \mathbf{m}_t using an empirical model based on \mathbf{m}^{t-1} . Let $\mathbf{c}_t \in \mathbb{N}_0^n$, and assume $[\mathbf{c}_{t-1}]_i = 1$ for $i \in \{0, \dots, n-1\}$. For all t , let $\mathbf{r}_t = \sum_{i=0}^n [\mathbf{c}_t]_i$. The PMF used for SFE encoding at time $t+1$ is based on

the empirical frequencies

$$[\mathbf{c}_t]_i = \begin{cases} \frac{1}{2}([\mathbf{c}_{t-1}]_i - 1) + 1 + \mathbb{1}_{\mathbf{m}_t=i}, & \text{if } \mathbf{r}_{t-1} = 2^{p/2} - 1 \\ [\mathbf{c}_{t-1}]_i + \mathbb{1}_{\mathbf{m}_t=i}, & \text{otherwise.} \end{cases}$$

We encode \mathbf{m}_t using the PMF $\mathbb{P}_{\mathbf{c}_{t-1}}(i) = [\mathbf{c}_{t-1}]_i$. Both the encoder and decoder begin with the same initial model \mathbf{c}_{-1} , and the encoder and decoder updates their models via after encoding/decoding so that they remain synchronized. The update rule periodically re-scales to ensure that the arithmetic operations required for SFE encoding/decoding can be carried out in p bits of precision. For more details, see [24, Chapter 11]. Since the model used for SFE encoding at time t depends only on \mathbf{m}^{t-1} , the SFE coding is prefix-free in the sense of Section II (given \mathbf{m}^{t-1}). The Elias omega coding is likewise instantaneously decodable, and thus the jointly encoding \mathbf{q}_t via SFE encoding $\bar{\mathbf{s}}_t$ and $\hat{\mathbf{s}}_t$ satisfies our desired prefix constraint. We now characterize this algorithm’s performance numerically.

V. NUMERICAL RESULTS

We simulated the proposed approach with linearized model for an inverted pendulum system [11]. The system is described in detail in Appendix I. The system’s state vector consists of $m = 4$ dimensions including a horizontal position and velocity as well as the pendulum’s azimuthal angle, and angular velocity, with respect to the normal from the cart’s platform. The control input is $u = 1$ dimensional. We discretized the system’s dynamics at a sampling frequency of 100 samples per second and assume that the continuous-time control input is via sample-and-hold at the same frequency. The discrete-time system dynamics are assumed to be $\mathbf{x}_{t+1} = A_\tau \mathbf{x}_t + B_\tau \mathbf{u}_t + \mathbf{w}_t$ where $A_\tau \in \mathbb{R}^{4 \times 4}$, $B_\tau \in \mathbb{R}^{4 \times 1}$, and \mathbf{w}_t is IID process noise such that $\mathbf{w}_t \perp \mathbf{u}_0^{t-1}, \mathbf{x}_0^t, \mathbf{w}_0^{t-1}$ and $\mathbf{w}_t \sim \mathcal{N}(0, W_\tau)$. We assume a diagonal W_τ such that $W_\tau = .005I_4$, and an uncertain initial configuration such that $\mathbf{x}_0 \sim \mathcal{N}(0, .05I_4)$. We will assume that the LQG cost weights are given by $Q = I_{4 \times 4}$ and $\Phi = 1$.

All simulations were performed using MATLAB R2022B [25]. We implemented the quantizer design from Section II and encoded the quantizations using the compression algorithm of Section IV. To obtain the solution to the rate-distortion optimization in (1), and to obtain the measurement matrix C , we used the YALMIP toolbox [26] with the MOSEK solver [27]. We used the “sorted” version of SFE coding in our implementation of the lossless compression algorithm; this tends to reduce redundancy (cf. [4, Section IV.A.1]). The SFE coding was implemented using unsigned 64-bit arithmetic, and was based on the arithmetic coding implementation in [9]. After some initial tuning, we settled on a cutoff vector of $k = [8191, 3, 3, 3]^T$. More details can be found in [24, Chapter 11].

Figure 3 empirically investigates the bitrate/LQG cost tradeoff. We plot the lower bound, $\mathcal{R}(\gamma)$ (cf. (1)) and the upper bound from (5) (which is $\approx \mathcal{R}(\gamma) + 1.26m + 1$). We simulated the algorithm from Sec. IV for $T = 400000$ samples, and plotted the average control cost versus the average codeword length over the entire horizon. We found




Fig. 3. The control cost constraint, γ is plotted on the horizontal, while the data rate in bits is plotted along the vertical. The behavior of the rate-distortion lower bound is typical; the horizontal asymptote of the lower bound corresponds to the minimum feasible LQG performance with fully observable state feedback (at this sample rate). The lower bound on data rate rises sharply near this asymptotic and becomes quite modest as a higher cost can be tolerated.

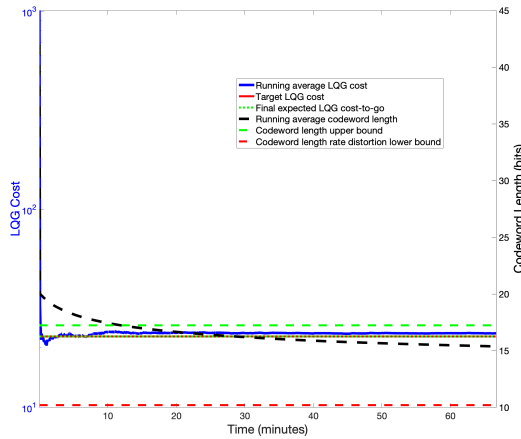


Fig. 4. The running average control and communication costs on the same plot for a particular fixed target control cost. We depict upper and lower bounds on codeword length, the target control cost, as well expected cost to go after $T = 400000$ samples. The expected cost to go is seen to converge to the target control cost.

the control cost slow to settle. The target control costs for most, but not all, of the points displayed are below the computed empirical average cost. Figure 4 illustrates the convergence of the running average bitrate and control cost.

VI. CONCLUSIONS

One immediate takeaway from Figure 3 is that the algorithm as proposed seems to outperform the upper bound from [4]. This suggests that the bounds from [4] can be tightened. If one uses prefix-free feedback at each sample instant, the true date rate (in bits per second) tends to infinity as the sampling rate increases. Likewise, a low sampling rate results in bad continuous-time control performance, and may result in long codewords. A more precise characterization of the




Fig. 5. The inverted pendulum, or “cart pole” system consists of a motorized wheeled cart that can move in one dimension along the z axis. The “inverted pendulum” is affixed to the top of the cart, and consists of a slender armature anchored to a fulcrum on the cart. The control input can accelerate the cart along its axis of motion in an effort to stabilize the pendulum about its unstable equilibrium at $\theta = 0$.

Variable	Descriptions	Value
μ_{cart}	mass of cart	.5 kilograms
μ_{pend}	mass of pendulum	.2 kilograms
κ	coefficient of friction for cart	.1 newton/(meters sec)
ψ	mass moment of inertia for pendulum	.006 kilogram meters ²
ϵ	length of pendulum to center of mass	.3 meters
\mathbf{g}	gravitational acceleration	9.8 meters/(sec) ²

TABLE I

PARAMETERS OF THE INVERTED PENDULUM SYSTEM IN FIG. 5 [11].

upper bounds from [4] would be useful for optimizing the sampling rate. In interpreting these numerical results, one must take care to not interpret them as a “real-world” experiment. They are performed using pseudorandom numbers and finite precision arithmetic using a standard consumer laptop. Such simulations doubtlessly suffer from numerical inaccuracies that may be significant in some applications.

Instead of encoding the linear index of the truncated symbol $\text{trunk}_k(\mathbf{s}_t)$ with a (sorted) SFE code, one could encode $\text{trunk}_k(\mathbf{s}_t)$ using arithmetic coding over the dimensions of the vector; viewing the elements of the vector as a Markov source. For example, the encoder and decoder could store one model for the first component $[\text{trunk}_k(\mathbf{s}_t)]_0$ of the source, then $[k]_0 + 1$ models for $[\text{trunk}_k(\mathbf{s}_t)]_1$, one for each potential realization of $[\text{trunk}_k(\mathbf{s}_t)]_0$, and so on. This approach allows a more accurate characterization of the source probability mass function at the expense of greater spatial complexity. It also reduces the frequency of model rescalings.

APPENDIX I

THE INVERTED PENDULUM’S SYSTEM DYNAMICS

We perform our experiments using a linearized model for the inverted pendulum system from [11]. The inverted pendulum system is depicted and described in Fig. 5. The systems state is considered in $m = 4$ dimensions, with \mathbf{z} the lateral position of the cart along its axis of motion (in meters), $\dot{\mathbf{z}} = \frac{d}{dt}\mathbf{z}$ the associated velocity (in meters/second), θ the angle of the pole from the vertical (in radians), and $\dot{\theta} = \frac{d}{dt}\theta$ the associated angular velocity (in radians/sec). The control input \mathbf{u} is one dimensional. For $t \in \mathbb{R}$, $t \geq 0$, denote

the continuous time state vector

$$\mathbf{x}(t) = [\mathbf{z}(t), \dot{\mathbf{z}}(t), \boldsymbol{\theta}(t), \dot{\boldsymbol{\theta}}(t)]^\top. \quad (15)$$

The true dynamics of the system are nonlinear, however, after linearizing, the continuous-time dynamics are assumed to be

$$d\mathbf{x}(t) = A_{\text{ct}}\mathbf{x}(t) + B_{\text{ct}}\mathbf{u}(t) + W_{\text{ct}}^{\frac{1}{2}}d\mathbf{w}(t), \quad (16)$$

where the system matrix A_{ct} and feedback matrix B_{ct} are functions of the system parameters given in Table I [11]. Let $\rho = \psi(\mu_{\text{pend}} + \mu_{\text{cart}}) + \mu_{\text{pend}}\mu_{\text{cart}}\epsilon^2$. Explicitly, we have

$$A_{\text{ct}} = \begin{bmatrix} 0 & 1 & 0 & 0 \\ 0 & -\frac{(\psi + \mu_{\text{pend}}\epsilon^2)\kappa}{\rho} & \frac{(\mu_{\text{pend}}^2\epsilon^2\mathbf{g})}{\rho} & 0 \\ 0 & 0 & 0 & 1 \\ 0 & -\frac{(\mu_{\text{pend}}\epsilon\kappa)}{\rho} & \frac{\mu_{\text{pend}}\mathbf{g}\epsilon(\mu_{\text{cart}} + \mu_{\text{pend}})}{\rho} & 0 \end{bmatrix}, \quad (17)$$

$$B_{\text{ct}} = \left[0, \frac{(\psi + \mu_{\text{pend}}\epsilon^2)}{\rho}, 0, \frac{\mu_{\text{pend}}\epsilon}{\rho} \right]^\top,$$

and $d\mathbf{w}(t)$ is standard Brownian motion that accounts for modeling error and unmodeled dynamics (wind, etc). We now construct a discrete-time version of the model in (16), assuming a sampling period τ . We assume a sample-and-hold feedback policy where $u(t) = u_{\lfloor \frac{t}{\tau} \rfloor}$. Let

$$A_\tau = e^{A_{\text{ct}}\tau}, B_\tau = \int_0^\tau e^{A_{\text{ct}}s} B_{\text{ct}} ds. \quad (18a)$$

and

$$W_\tau = \left(\int_0^\tau e^{A_{\text{ct}}s} W_{\text{ct}} e^{A_{\text{ct}}^\top s} ds \right)^{\frac{1}{2}}, \quad (18b)$$

where for a matrix M argument, e^M refers to the matrix exponential of M . For $t \in \mathbb{N}_0$, we let $\mathbf{x}(t\tau) = \mathbf{x}_t$ and $\mathbf{u}(t\tau) = \mathbf{x}_t$.

REFERENCES

- [1] E. I. Silva, M. S. Derpich, and J. Ostergaard, "A framework for control system design subject to average data-rate constraints," *IEEE Trans. Automat. Contr.*, vol. 56, no. 8, pp. 1886–1899, 2011.
- [2] T. Tanaka, K. H. Johansson, T. Oechtering, H. Sandberg, and M. Skoglund, "Rate of prefix-free codes in LQG control systems," in *Proc. IEEE ISIT*, 2016, pp. 2399–2403.
- [3] V. Kostina and B. Hassibi, "Rate-cost tradeoffs in control," *IEEE Trans. on Automat. Contr.*, vol. 64, no. 11, pp. 4525–4540, 2019.
- [4] T. Cuvelier, T. Tanaka, and R. W. Heath, "Time-invariant prefix coding for LQG control," *IEEE J. Sel. Areas Inf. Theory*, pp. 1–1, 2022.
- [5] S. Boucheron, A. Garivier, and E. Gassiat, "Coding on countably infinite alphabets," *IEEE Trans. Inf. Theory*, vol. 55, no. 1, pp. 358–373, 2009.
- [6] D. Bontemps, "Universal coding on infinite alphabets: Exponentially decreasing envelopes," *IEEE Trans. Inf. Theory*, vol. 57, no. 3, pp. 1466–1478, 2011.
- [7] D. Bontemps, S. Boucheron, and E. Gassiat, "About adaptive coding on countable alphabets," *IEEE Trans. Inf. Theory*, vol. 60, no. 2, pp. 808–821, 2014.
- [8] S. Boucheron, E. Gassiat, and M. I. Ohannessian, "About adaptive coding on countable alphabets: Max-stable envelope classes," *IEEE Trans. Inf. Theory*, vol. 61, no. 9, pp. 4948–4967, 2015.
- [9] I. H. Witten, R. M. Neal, and J. G. Cleary, "Arithmetic coding for data compression," *Commun. ACM*, vol. 30, no. 6, p. 520–540, Jun 1987. [Online]. Available: <https://doi.org/10.1145/214762.214771>

- [10] P. Elias, "Universal codeword sets and representations of the integers," *IEEE Trans. Inf. Theory*, vol. 21, no. 2, pp. 194–203, 1975.
- [11] D. Tilbury, B. Messner, J. Luntz, R. Hill, J. Taylor, and S. Das. (2019 Accessed Feb. 2023) Control tutorials for Matlab and Simulink: Inverted pendulum. [Online]. Available: <https://ctms.engin.umich.edu/CTMS/index.php?example=InvertedPendulum§ion=SystemModeling>
- [12] T. C. Cuvelier, T. Tanaka, and R. W. Heath, "A lower-bound for variable-length source coding in Linear-Quadratic-Gaussian Control with shared randomness," *IEEE Control Systems Letters*, vol. 6, pp. 2918–2923, 2022. [Online]. Available: <https://arxiv.org/abs/2203.12467>
- [13] N. Elia and S. Mitter, "Stabilization of linear systems with limited information," *IEEE Transactions on Automatic Control*, vol. 46, no. 9, pp. 1384–1400, 2001.
- [14] G. N. Nair and R. J. Evans, "Stabilizability of stochastic linear systems with finite feedback data rates," *SIAM J. Contr. Optim.*, vol. 43, no. 2, pp. 413–436, 2004. [Online]. Available: <https://doi.org/10.1137/S0363012902402116>
- [15] A. Khina, Y. Nakahira, Y. Su, and B. Hassibi, "Algorithms for optimal control with fixed-rate feedback," in *Proc. IEEE CDC*, 2017, pp. 6015–6020.
- [16] S. Yüksel, "Stochastic stabilization of noisy linear systems with fixed-rate limited feedback," *IEEE Trans. Automat. Contr.*, vol. 55, no. 12, pp. 2847–2853, 2010.
- [17] O. Sabag, V. Kostina, and B. Hassibi, "Stabilizing dynamical systems with fixed-rate feedback using constrained quantizers," in *Proc. IEEE ISIT*, 2020, pp. 2855–2860.
- [18] R. Krichevsky and V. Trofimov, "The performance of universal encoding," *IEEE Trans. Inf. Theory*, vol. 27, no. 2, pp. 199–207, 1981.
- [19] T. Tanaka, P. M. Esfahani, and S. K. Mitter, "LQG control with minimum directed information: Semidefinite programming approach," *IEEE Trans. Automat. Contr.*, vol. 63, no. 1, pp. 37–52, 2018.
- [20] T. Lattimore and C. Szepesvári, *Bandit Algorithms*. Cambridge Univ. Press, 2020.
- [21] E. Pauwel, "Statistics, optimization and algorithms in high dimension," Sep. 2020. [Online]. Available: <https://www.math.univ-toulouse.fr/~epauwels/M2RI/poly.pdf>
- [22] P. Harremoës, *Information Topologies with Applications*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2007, pp. 113–150. [Online]. Available: https://doi.org/10.1007/978-3-540-32777-6_5
- [23] T. M. Cover and J. A. Thomas, *Elements of Information Theory*, 2nd ed. Wiley-Interscience, 2006.
- [24] T. Cuvelier, "Enabling reliable control with communication constraints," Ph.D. dissertation, The University of Texas at Austin, 2023. [Online]. Available: <https://www.traviscuvelier.io/research>
- [25] MATLAB, *version R2022b*. Natick, Massachusetts: The MathWorks Inc., 2022.
- [26] J. Löfberg, "Yalmip : A toolbox for modeling and optimization in matlab," in *Prof. IEEE CACSD*, Taipei, Taiwan, 2004.
- [27] M. ApS, *The MOSEK optimization toolbox for MATLAB manual. Version 10.0.38*, 2023. [Online]. Available: <https://www.mosek.com/>