2023 62nd IEEE Conference on Decision and Control (CDC)
December 13-15, 2023. Marina Bay Sands, Singapore

Steering a linear system at the minimum information rate: quantifying redundancy via covariance assignment theory

Eric Wendel, John Baillieul, and Joseph Hollmann

Abstract— **We compute fundamental performance limitations in the data rate constrained control of continuous-time linear stochastic control systems using information theoretic tools and principles. Specifically, we find the minimum achievable mean square error for steering a linear system to sequences of uncertain steering objectives under a rate constraint as the solution of a convex optimization problem. We propose the redundancy of a control system as a measure of the relative inefficiency of information transmission through a linear control system vs. an ideal communications channel.**

I. Introduction

In almost all practical applications of interest, a control or state estimation algorithm is implemented on an embedded system and, as a result, subject to data rate constraints on both its inputs and outputs. Maximum clock rates and bit depths limit the amount of data, in bits/sec, that an algorithm may employ towards a specific task. These digital constraints are also universal, in the sense that the same maximum data rate applies to a controller irrespective of its design paradigm, e.g. state space, machine learning, or neuromimetic. We may ask the following question: *to what extent is the performance of a control algorithm determined by its data rate constraint?*

We answer this question in the context of a prototypical *steering problem*, in which a linear control system must repeatedly steer itself to a sequence of desired points in state space. The performance of the control algorithm is a mean square error measure. In addition to data rate constraints, stochastic effects and uncertainty are also seen as intrinsic to digital implementations and the performance of a controller is therefore limited by the effects of thermal noise, clock jitter, and other errors. This steering problem is intended to model a wide range of typical applications in which the control algorithm does not know the next steering point in the sequence, and is limited by uncertainty and process noise.

Our approach to quantifying the fundamental performance limitations of data rate constrained steering in the presence of noise and uncertainty is based on principles of information theory, specifically those of rate distortion theory. Our main result is the formulation of a convex optimization problem whose solution yields the minimum achievable mean square steering error by any controller for a continuous-time stochastic linear system in the presence of noise, uncertainty, and rate constraints. Our results inform the design of sampling and quantization strategies in digital control and contribute to a growing body of interdisciplinary work at the intersection of control and information theory.

Related literature. Inspired by information-theoretic principles, information-based approaches to state estimation and control often involve maximization of an appropriate information measure, such as the mutual information between the input and output of a measurements channel [12], or the Kullback-Leibler divergence between an uninformative prior and a state distribution [14]. In [8], [9], [15] the authors solve optimization problems whose objectives are the entropy between states, controls, and measurements. Entropy-based approaches are often justified by the observation that entropy serves as an indirect proxy for the amount of uncertainty in a random variable.

Our methodology is similar in some respects to these information-based approaches, except that our primary performance measure is a mean square steering error. The operational significance of our information measure, the mutual information between the steering objective and the system state, is established by fundamental source coding theorems.

Our main result, Theorem III.2, applies both covariance assignment theory and rate distortion theory in order to find the optimal distribution of system state minimizing the mean square steering error, in the presence of a data rate constraint. Our work is thus related to prior work on linear quadratic Gaussian density steering [11], [5]. The objective of these and other works on density steering is to steer the control system's marginal statistics to a desired marginal density with the aid of information-based measures such as the Kullback-Leibler divergence or Wasserstein metric.

In communications and information theory, as in this work, the objective is to design deterministic encoders and decoders that achieve an acceptable level of average performance. Analysis of a large class of such deterministic mappings is enabled by constructing statistical descriptions of their behavior in the presence of a source of uncertainty. Consequently, although the main contributions of this work can be interpreted as results on rate-constrained density steering, the analysis is relevant to the design of controllers and estimators for steering a control system as accurately as possible to specific points in state space.

Contributions and organization. Our first contribution is a new representation of the so-called *rate distortion function* of a multivariate Gaussian source of information as

E. Wendel is with the Division of Systems Engineering, Boston University, Boston MA 02215 USA, and also a Draper Scholar with The Charles Stark Draper Laboratory, Inc., Cambridge, MA 02139 USA edbw@bu.edu.

J. Baillieul is with the Division of Systems Engineering, Boston University, Boston MA 02215 USA johnb@bu.edu.

J. Hollmann is with The Charles Stark Draper Laboratory, Inc., Cambridge, MA 02139 USA jhollmann@draper.com.

a convex optimization problem in Theorem II.2. We show its equivalence to the classical reverse waterfilling algorithm [2], [6]. Our second main result, Theorem II.3, establishes the operational significance of the rate distortion function to the steering problem, viewed as an instance of a remote or noisy source coding problem in information theory. These two main results are presented in Section II along with a review of rate distortion theory.

We apply these theorems in Section III, where we formally define the steering problem and provide our final main result, regarding the minimum achievable mean square steering error (MMSE) in rate-constrained control. Our analysis investigates the sensitivity of this MMSE to sampling rate and system stability. We then briefly define and discuss the information-based notion of control system redundancy. Inspired by the coding redundancy relevant to source coding problems, the redundancy we propose measures the difference between the behavior of a control system and that of an ideal communications channel, in the context of the steering problem. Section IV concludes with remarks on future work.

II. The Information Rate of a Memoryless Gaussian Source

Let x_d be a continuous vector-valued random variable taking values on some subset $\mathcal{V} \subseteq \mathbb{R}^n$, and $p(x_d)$ its density. If all samples from this density are iid then we say $x_d \sim p(x_d)$ is a (stationary) *memoryless source*. A *reproduction* of x_d is a random variable $\hat{x}_d := g \circ f(x_d)$ where the *compressor* $f : \mathcal{V} \to \mathcal{U}$ maps a realization of the random variable x_d onto its representation in some set $\mathcal{U}$ and the *decompressor* $g : \mathcal{U} \to \mathcal{V}$ maps the representation back. If the set $\mathcal{U} = \{\mathsf{u}_k\}_{k=1}^K \subset \mathbb{R}^n$ is finite, $K < \infty$, then we call the elements of $\mathcal{U}$ *codevectors* or *symbols* and the f, g together with the sets $\mathcal{V}$, $\mathcal{U}$ a *source code* of *size* K. If the maps f and g are suitably extended to operate on a sequence of L iid samples drawn from $p(x_d)$ then we say the source code has *blocklength* L. This defines the *source code rate*

$$\bar{R} := \frac{1}{L}\log_2(K).$$

The error in the reproduction of a source is called its *distortion*. In this work we are only concerned with mean square error distortion. A source code design $g \circ f$ is said to be *admissible* with distortion D if its mean square error satisfies

$$\bar{D} := E[\|x_d - g \circ f(x_d)\|_2^2] \leq D. \tag{1}$$

Given the marginal density $p(x_d)$, recall that the *mutual information* $I(x_d; \hat{x}_d) = H(x_d) - H(x_d|\hat{x}_d)$ between the source random variable x_d and its reproduction $\hat{x}_d$ is a function only of the conditional density function $p(x_d|\hat{x}_d)$. This density may be viewed as a statistical description of a given deterministic source code design $g \circ f$. The *rate distortion function* is defined as the solution of the following variational problem over the set of all statistical characterizations, or *backwards test channels*, of admissible source codes:

$$R(D) := \min_{p(x_d|\hat{x}_d)} \; I(x_d; \hat{x}_d) \quad \text{s.t.} \quad E[\|x - \hat{x}\|_2^2] \leq D \tag{2}$$

The Converse Source Coding Theorem [2, Theorem 3.2.2] states that the code rate $\bar{R}$ of any admissible source code with any blocklength $L > 0$ and distortion $\bar{D} < D$ is such that $\bar{R} \geq R(D)$. An important consequence of this fundamental theorem, the Information Transmission Theorem establishes the operational significance of the rate distortion function as the *necessary and sufficient code rate* of an admissible source code. The following is Theorems 3.3.1-2 of [2] restated for a memoryless source $p(x_d)$.

Theorem II.1 (Information Transmission Theorem [2])**.** *The memoryless source $p(x_d)$ can be reproduced with maximum admissible distortion $D + \epsilon$ for any $\epsilon > 0$ at the output of a channel of capacity C if and only if $R(D) + \epsilon < C$.*

As noted in [2], the operational significance of $R(D)$ does not depend on arbitrarily long blocklengths. Theorem II.1 also applies to both finite and variable length codes.

It is well-known that if the memoryless source is a zero-mean Gaussian with covariance Σ_d, *i.e.* $p(x_d) = \mathcal{N}(0, \Sigma_d)$, then variational problem (2) reduces to a minimization over the set of conditional covariance matrices of the (backwards) test channel. Computing $R(D)$ is then possible via a simple procedure known colloquially as *reverse waterfilling* (cf. [6, Theorem 10.3.3], [2, equation (4.5.21)]):

$$\begin{aligned} R(D) = R_L(D) &:= \frac{1}{2}\sum_{i=1}^n \log\left(\frac{\sigma_i^2}{D_i(\theta)}\right), \\ D &= \sum_{i=1}^n D_i(\theta), \quad D_i(\theta) = \min\{\theta, \sigma_i^2\} \end{aligned} \tag{3}$$

where σ_i^2 is the i^{th} eigenvalue of Σ_d. Note that if $\sigma_i^2 < \theta$ then $\log(\sigma_i^2/D_i(\theta)) = 0$ and the i^{th} principal component of the source contributes no bits to the minimum admissible code rate $R_L(D)$. Reverse waterfilling transforms the source into n independent components and truncates those with variance below the "water level" θ.

The following observation leads to an equivalent characterization of $R_L(D)$ as the solution of a *maximum determinant* (max-det) problem. Let $e := x_d - \hat{x}_d$, $\Lambda := \mathrm{Cov}(e) = \mathrm{Cov}(x_d|\hat{x}_d)$, and observe that

$$H(x_d|\hat{x}_d) = H(x_d - \hat{x}_d|\hat{x}_d) \leq H(e) \leq H(z) \tag{4}$$

where $z \sim \mathcal{N}(0, \Lambda)$ has the same covariance as the error e. Since $I(x_d; \hat{x}_d) = H(x_d) - H(x_d|\hat{x}_d)$ we obtain

$$\begin{aligned} R_L(D) = \frac{1}{2}\log\det(\Sigma_d) - \max_{\Lambda \succeq 0} \; &\frac{1}{2}\log\det(\Lambda) \\ \text{s.t.} \;\; &\Lambda \preceq \Sigma_d, \; \mathrm{tr}(\Lambda) \leq D \end{aligned} \tag{5}$$

Equation (5) expresses the *Shannon lower bound* for this source, cf. [2, eqn. (4.3.11)], as a max-det problem over the set of all backwards channel conditional covariances Λ.

Our first result represents $R(D)$ as a max-det problem in the parameters defining the *forwards channel* $p(\hat{x}_d|x_d)$. The proof may be found in the appendix.

Theorem II.2. *The rate distortion function of the n-dimensional multivariate Gaussian memoryless source*

$p(x_d) = \mathcal{N}(\mu_d, \Sigma_d)$ *is the solution of the max-det problem*

$$\begin{aligned} R(D) = \min_{\beta,\Lambda,\Sigma} \quad & \frac{1}{2}\log\det(\Sigma_d) - \frac{1}{2}\log\det(\Lambda) \\ \textit{s.t.} \quad & \begin{bmatrix} \Sigma_d - \Lambda & \beta^T \\ \beta & \Sigma \end{bmatrix} \succeq 0, \quad \Lambda, \Sigma \succeq 0 \\ & \mathrm{tr}(\Sigma - 2\beta + \Sigma_d) \leq D, \ \beta \in \mathbb{R}^{n\times n}. \end{aligned} \tag{6}$$

The optimal β, Σ *define the optimal forwards test channel*

$$p(\hat{x}_d|x_d) = \mathcal{N}(\mu_d + \beta\Sigma_d^{-1}(x_d - \mu_d), \ \Sigma - \beta\Sigma_d^{-1}\beta^T). \tag{7}$$

It is not surprising that representation (6) of the standard rate distortion function does not often, or perhaps at all, appear in the literature given (3) and (5). In [16] the authors compute the *sequential* rate distortion function relevant to applications involving exclusively causal source coding. A special case of their max-det problem reduces to (5).

Expressing the standard rate distortion function as a max-det problem in the forwards channel parameters is a necessary first step in our analysis of the steering problem in the next section. First, note that so far we have assumed that we may implement any source code $g \circ f$ with desired statistics, a design freedom that is rarely encountered in practice [18]. The noisy or *remote* source coding problem considers the additional distortion introduced by constraints on the design of the source code, usually in the form of additional fixed channels separating the source from its destination. The steering problem is a special case in which the forwards test channel $p(\hat{x}_d|x_d)$ is constrained to yield output statistics $p(\hat{x}_d)$ of a particular form. Consistent with the approach of [2, Theorem 3.5.1], we introduce below a modified rate distortion function and establish, as our second main result, its operational significance as the necessary and sufficient code rate.

For the Gaussian memoryless source $p(x_d) = \mathcal{N}(\mu_d, \Sigma_d)$ it suffices to consider only Gaussian test channels $p(\hat{x}_d|x_d)$. Equation (7) identifies every such channel by parameters $(\beta, \Sigma) \in \mathbb{R}^{n\times n} \times \mathsf{S}_+^n$. Let $H : \mathbb{R}^{n\times n} \times \mathsf{S}_+^n \to \mathbb{R}^{n\times n} \times \mathbb{R}^{n\times n}$ be a linear function in β, Σ and $\mathcal{A} := H^{-1}(0)$ denote its zero level-set. Let $p(\hat{x}_d|x_d) \in \mathcal{A}$ denote, by an abuse of notation, a forwards test channel with $(\beta, \Sigma) \in \mathcal{A}$. Finally, let $R_{\mathcal{A}}(D)$ denote the constrained rate distortion function obtained by solving (6) with the additional linear constraint $H(\beta, \Sigma) = (0, 0)$. The proof of the following can be found in the appendix.

Theorem II.3. *The source* $p(x_d)$ *can be reproduced with fidelity* $D + \epsilon$ *for any* $\epsilon > 0$ *at the output of a channel of maximum capacity* C *per source symbol by a test channel* $p(\hat{x}_d|x_d) \in \mathcal{A}$ *if and only if* $R(D) \leq R_{\mathcal{A}}(D) + \epsilon < C$.

III. THE STEERING PROBLEM

This section addresses the following question: *what is the minimum achievable mean square error in steering a linear control system under a data rate constraint?* Consider the controllable n-dimensional continuous-time stochastic control system

$$dx(t) = \big(A\,x(t) + B\,u(t)\big)dt + dw(t). \tag{8}$$

where $dw(t)$ is Brownian motion with $E[dw(t)dw(t)^T] = N\,dt$. The controller is digital, meaning that at any time t, the m-dimensional control signal $u(t)$ is represented by m (possibly uncompressed) binary numbers. Digital controls are data rate constrained. For example, suppose the processor clock, bus, or communications protocol operates at some maximum rate $f = 1/\Delta t$ Hz. If each control component is also represented with maximum bit depth r, then the controls are subject to a data rate constraint of fmr bits/sec. In order to understand how these digital constraints limit the performance of control system (8) we study its *sampled data representation*

$$\begin{aligned} x_{k+1} &= Fx_k + Gu_k + \Delta w_k, \\ F = e^{A\Delta t}, &\qquad G = \int_0^{\Delta t} e^{A\tau} d\tau\, B, \\ \Delta w_k \sim \mathcal{N}(0, W), &\quad W = \int_0^{\Delta t} e^{A\tau} N e^{A^T\tau}\, d\tau \end{aligned} \tag{9}$$

where $x_{k+1} := x(t_{k+1})$ denotes state at time t_{k+1}.

We are interested in applications where x_k is (fully) observed every M timesteps at total of L times. That is, for every $k = jM$ we observe x_{jM}, where $j \in [L]$ and $[L] := \{1, \ldots, L\}$ denotes the 1-based index set of size L. At each t_{jM} a sequence of M open-loop controls are designed based on the information available at that time.

The steering problem can be seen in this context as a data rate constrained communications problem. Our objective is to find a deterministic controller that minimizes the mean square steering error

$$\bar{D} := \frac{1}{L}\sum_{j\in[L]} \left\| x_d^{(j)} - x_{jM} \right\|_2^2 \tag{10}$$

where $\{x_d^{(j)}\}_{j=1}^L$ is a set of L steering objectives. Assume that the initial condition is uncertain, with $x_0 \sim \mathcal{N}(\mu_0, \Sigma_0)$, that the steering objectives are iid $x_d^{(j)} \sim \mathcal{N}(\mu_d, \Sigma_d)$, and that the $x_0, x_d^{(j)}, \Delta w_k$ are independent for all $j \in [L]$ and for all k. The controller must encode each steering objective into a sequence of M control vectors, each belonging to an *admissible control set* $\mathcal{U} \subset \mathbb{R}^m$. Since the controls are digital, $\mathcal{U}$ is a finite and discrete set of cardinality K, defining the controller's code rate

$$\bar{R} = \frac{1}{L}\log_2(KM)$$

in bits per source symbol. It follows from Theorem II.1 that $\bar{R} \geq R(D)$. Given that steering objectives are iid, if it is known that the steering problem is to be solved repeatedly for large L then a plausible steering strategy is to control the marginal statistics of the x_{jM} to the same stationary distribution. Our main result, Theorem III.2, computes the optimal stationary distribution achieving this minimum mean square steering error under a data rate constraint. We state our results for the case $M = 1$.

First, we recall the following result from covariance assignment theory [13], [19]. Let $\Pi_G^\perp = I - GG^+$ denote projection onto the orthogonal complement of the range of the sampled-data control matrix G of (9).

Theorem III.1 ([19, Theorem 2.1])**.** *There exists a constant feedback gain control,* $u_j = K(x_d^{(j)} - x_j)$*, rendering* $\Sigma \in \mathsf{S}_+^n$ *a fixed point of the discrete-time Lyapunov equation for system (9),*

$$\begin{aligned}\Sigma &= (F - GK)\Sigma(F - GK)^T + W + GK\Sigma_d K^T G^T \\ &\quad + GK\,\mathrm{Cov}(x_d, x_j)(F - GK)^T \\ &\quad + (F - GK)\,\mathrm{Cov}(x_j, x_d)K^T G^T\end{aligned}$$

if and only if

$$\Pi_G^\perp \left(F\Sigma F^T + W - \Sigma\right)\Pi_G^\perp = 0. \tag{11}$$

Condition (11) defines the set of *assignable covariances* as the zero level-set of a linear function on S_+^n. Its members satisfy the Lyapunov equation modulo components in the range space of the control matrix G.

Theorem III.2. *The minimum achievable mean square error of any controller steering system (9) to an assignable covariance under a rate constraint* R *is given by the distortion rate function*

$$\begin{aligned} D_{\Delta t}(R) := \min_{\mu,\beta,\Lambda,\Sigma}\ & \mathrm{tr}(\Sigma - 2\beta + \Sigma_d) \\ \text{s.t.}\quad & \begin{bmatrix}\Sigma_d - \Lambda & \beta^T \\ \beta & \Sigma\end{bmatrix} \succeq 0, \quad \Lambda, \Sigma \succeq 0 \\ & \frac{1}{2}\log\det(\Sigma_d) - \frac{1}{2}\log\det(\Lambda) \leq R, \\ & \Pi_G^\perp \beta = 0, \quad \Sigma \textit{ satisfies } (11), \\ & \Pi_G^\perp (F\mu - \mu_d) = 0, \\ & \mu \in \mathbb{R}^n,\ \beta \in \mathbb{R}^{n\times n}. \end{aligned} \tag{12}$$

Conversely, the necessary and sufficient code rate of any controller that achieves a given a distortion D *is given by the inverse of (12), the rate distortion function* $R_{\Delta t}(D)$*.*

Proof. Let $x \sim p(x) = \mathcal{N}(\mu, \Sigma)$ denote the marginal distribution of the system state. With the controls defined in Theorem III.1, its conditional mean satisfies

$$E[x|x_d] = (F - GK)\mu + GKx_d = \beta\Sigma_d^{-1}x_d + \alpha$$

and so $\alpha = (F - GK)\mu$, $\beta\Sigma_d^{-1} = GK$. This yields a linear constraint $\Pi_G^\perp \beta = 0$.

The mean square error is minimized if $\alpha = (F - GK)\mu = (I - \beta\Sigma_d^{-1})\mu_d$, as in the proof of Theorem II.2. This yields

$$\Pi_G^\perp \alpha = \Pi_G^\perp F\mu = \Pi_G^\perp (I - \beta\Sigma_d^{-1}\alpha)\mu_d = \Pi_G^\perp \mu_d. \tag{13}$$

Let D be some admissible mean square steering error, and define $R_{\Delta t}(D)$ by applying the above linear constraints to (6). If $R_{\Delta t}(D)$ is feasible then it has a unique minimum. It follows from Theorem II.3 that this minimum defines the necessary and sufficient code rate for steering the system's statistics to $p(x)$. If $R_{\Delta t}(D)$ is infeasible then no controller exists that can achieve the desired mean square error D for system (9).

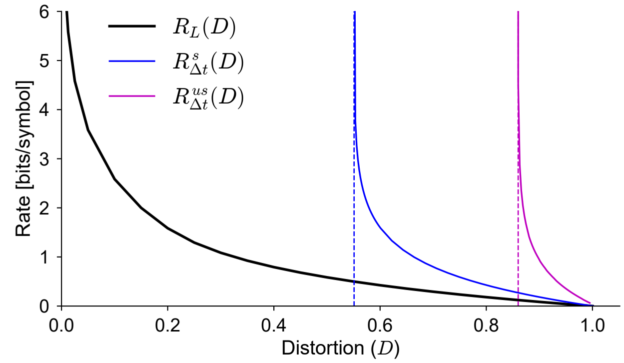

Fig. 1. The Shannon lower bound $R_L(D)$ (black) is compared to the rate distortion curves of a stable ($A_s = \left[\begin{smallmatrix}0 & 1\\ -1 & -2\end{smallmatrix}\right]$, blue) and unstable ($A_{us} = \left[\begin{smallmatrix}2 & 1\\ -1 & 0\end{smallmatrix}\right]$, magenta) linear system, both sampled at 100 Hz and driven by Brownian motion with intensity N = 1e-3 $I\,dt$. The memoryless source of steering objectives is zero-mean with covariance $\Sigma_d = \left[\begin{smallmatrix}0.1 & \\ & 0.9\end{smallmatrix}\right]$. The minimum achievable steering MSE for these systems is $D_s^* = 0.55$ (stable, dashed blue) and $D_{us}^* = 0.86$ (unstable, dashed magenta).

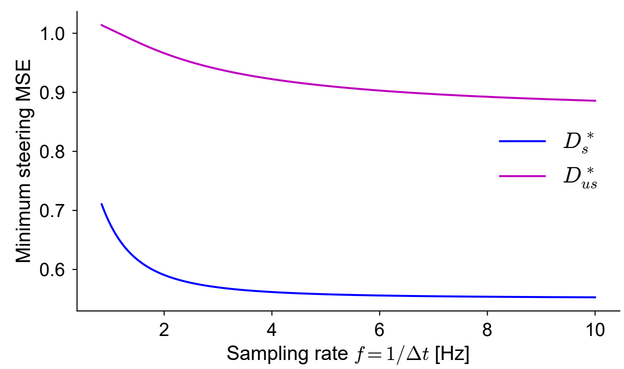

Fig. 2. Increasing the system sampling rate above about 5 Hz does not significantly decrease the minimum achievable steering MSE of the stable (blue) and unstable (magenta) systems of Figure 1.

Construct the inverse of $R_{\Delta t}(D)$, the distortion rate function $D_{\Delta t}(R)$, by swapping the roles of the trace and mutual information terms as cost and constraint. It follows again from the proof of Theorem II.3 that $D_{\Delta t}(R)$ describes the minimum achievable mean square steering error using a controller of code rate R. $\square$

Distortion, sampling rate, and data rate. Both $R_{\Delta t}(D)$ and $D_{\Delta t}(R)$ are straightforward to solve for various sampled data representations of both stable and unstable continuous-time linear time invariant systems. We use [1], [7] to solve all max-det problems. The process noise covariance (9) is computed using a 4th-order adaptive Runge-Kutta method.

Figure 1 compares the rate distortion curves for two different linear systems against the Shannon lower bound. While an ideal communications channel can approach zero distortion at rates exceeding 6 bits/symbol, there is a strict minimum achievable distortion (MMSE) of $D_{\Delta t}^* = 0.55$ and $D_{\Delta t}^* = 0.86$ for the stable and unstable linear systems,

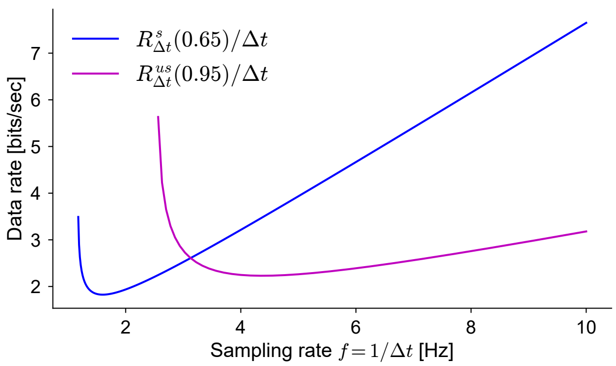

Fig. 3. The necessary and sufficient data rates for steering the stable (blue) and unstable (magenta) systems of Figure 1 with admissible MSE of 0.65 and 0.95, respectively.

respsectively. This performance is achieved by sampling both systems at a rate of $f = 1/\Delta t = 100$ Hz.

Can this sampling rate be lowered without significantly impacting the steering MMSE? Figure 2 shows that the fundamental steering performance is relatively insensitive to changes in sampling rate above 5 Hz. An order of magnitude increase in sampling rate, from 10 to 100 Hz, is evidently required to decrease the MMSE by about 0.1 units.

If the source emits steering objectives at the same system sampling rate, f symbols/sec, then the necessary and sufficient *data* rate for control is $R_{\Delta t}(D) \times f$ bits/sec. Figure 3 shows how this data rate varies with f. Both curves indicate that sampling the stable and unstable systems at rates away from about 1.6 Hz and 4.3 Hz requires increased data transmission in order to maintain the same MMSEs of 0.65 and 0.95, respectively.

These results confirm expectations regarding the effects of linear system stability and temporal sampling on fundamental steering performance: a stable system has a lower necessary and sufficient code rate and MMSE than a comparable unstable system. We have not, however, fully characterized the relationship between rate, distortion, and system stability. This relationship is likely determined by the geometry of the source and the channel, the former defined by Σ_d, and the latter by the system matrix A, the control matrix B, and the Brownian motion intensity N. Consider, for example, the chosen source covariance $\Sigma_d = \begin{bmatrix} 0.1 & \\ & 0.9 \end{bmatrix}$. If the admissible MMSE is large then any small transient motions of the control system in the x-axis direction, due to either process noise or instability, may be relatively unimportant compared to the large variance in steering objectives along the y-axis. This intuitive argument follows from the reverse waterfilling principle discussed in Section II, under which an ideal source code allocates bits first to encoding (steering) components of the source with sufficiently large variance. Controllability is important in this context as well. If the Brownian motion process tends to diffuse in certain directions that are within the range space of the sampled system control matrix, then increases of intensity in those directions may not warrant increases in the necessary and sufficient steering code rate.

On the redundancy of a control system. We conclude this section with remarks on performance measures for non-ideal communications channels. Let $q(x)$ denote the distribution of a discrete memoryless source. In [6] the cost incurred in compressing this source with the wrong source code is called its *redundancy*. Specifically, if the code were designed for some other source with distribution $p(x)$ then its coding redundancy is defined by $\mathcal{R}(q,p) = KL(q;p) := \sum_x p(x) \log(p(x)/q(x))$, the Kullback-Leibler divergence between q and p.

Inspired by this, we propose a measure of the cost in bits of using a linear control system to communicate the memoryless Gaussian source $\mathcal{N}(\mu_d, \Sigma_d)$ of steering objectives. Let $\hat{x}_d \sim q(\hat{x}_d) = \mathcal{N}(\hat{\mu}_d, \hat{\Sigma}_d)$ the marginal density obtained by reverse waterfilling with an ideal test channel. That is, $\hat{\Sigma}_d$ minimizes (6), and $\hat{\mu}_d = \mu_d$.

Definition III.1. The *redundancy* in steering linear control system (9) is the Kullback-Leibler divergence between $q(\hat{x}_d)$ and the marginal $p(x) = \mathcal{N}(\mu, \Sigma)$ where μ, Σ solve (12): $\mathcal{R}(q,p) := KL(q;p)$.

From Figure 1 this information-based notion of steering redundancy is only well-defined for distortions greater than or equal to the control system's MMSE. The *distortion redundancy*, or difference between $D(R)$, as the inverse of (6), and $D_{\Delta t}(R)$ of (12), would be an alternative measure of redundancy directly relevant to system performance. Future work will explore the utility of these and other redundancy measures in contexts where the distribution of steering objectives is not exactly known.

IV. Conclusion

We computed the minimum achievable mean square steering error for data rate constrained steering of a continuous-time stochastic linear control system using fundamental principles of information theory. Our analysis suggests that this MMSE and its corresponding necessary and sufficient code rate can be understood in terms of the geometry of the source, Σ_d, and the geometry of the linear system's sampled data representation. Our results connect with existing work on linear Gaussian density steering, covariance assignment theory, and adds to a growing body of literature on applications of rate distortion theory in control.

Our approach was intrinsically open-loop, due to our continued interest [17] in the design of digital controllers capable of operating on a spectrum between open- and closed-loop control, as proposed in [4]. Future work will analyze the delay-distortion trade-off for this and related state estimation and control problems.

Appendix

The proof of Theorem II.2 is a straightforward application of the following Schur complement lemma, cf. [3].

Lemma. *The Schur complement $\Psi - \beta^T\Sigma^{-1}\beta$, with $\Psi, \Sigma \in \mathsf{S}^n_+$, $\beta \in \mathbb{R}^{n\times n}$, is given by*

$$\Psi - \beta^T\Sigma^{-1}\beta = \arg\min_{\Lambda\succeq 0} -\log\det(\Lambda) \quad \textit{s.t.} \quad \begin{bmatrix}\Psi-\Lambda & \beta^T\\ \beta & \Sigma\end{bmatrix} \succeq 0$$

Proof of Theorem II.2. Let $\beta \in \mathbb{R}^{n\times n}$ and $\Sigma \succeq 0$ be $n\times n$ matrices defining our "test channel," the joint $2n\times 2n$ covariance

$$\mathrm{Cov}\begin{pmatrix}\hat{x}_d\\ x_d\end{pmatrix} := \begin{bmatrix}\Sigma & \beta\\ \beta^T & \Sigma_d\end{bmatrix},$$

with Σ_d given. Let $\Lambda := \mathrm{Cov}(x_d|\hat{x}_d) = \Sigma_d - \beta^T\Sigma^{-1}\beta$ and $\Gamma := \mathrm{Cov}(\hat{x}_d|x_d) = \Sigma - \beta\Sigma_d^{-1}\beta^T$.

Applying the Schur complement Lemma, we obtain the mutual information between x_d and $\hat{x}_d$ as the solution of the max-det problem (for given β and Σ):

$$\begin{aligned} I(x_d;\hat{x}_d) &= \frac{1}{2}\log\det(\Sigma_d) - \frac{1}{2}\log\det(\Lambda)\\ &= \min_{\Lambda\succeq 0} \quad \frac{1}{2}\log\det(\Sigma_d) - \frac{1}{2}\log\det(\Lambda) \qquad (14)\\ &\quad\ \text{s.t.} \quad \begin{bmatrix}\Sigma_d-\Lambda & \beta^T\\ \beta & \Sigma\end{bmatrix}\succeq 0\end{aligned}$$

We may write $E[\hat{x}_d|x_d] = \beta\Sigma_d^{-1}x_d + \alpha$, for some vector α. Setting $\alpha := (I - \beta\Sigma_d^{-1})\mu_d$ removes the quadratic dependence of the MSE on the means.

$$\begin{aligned}\bar{D} &= E[\|\hat{x}_d - x_d\|_2^2]\\ &= \mathrm{tr}\, E\Big[E\big[\hat{x}_d\hat{x}_d^T\big|x_d\big] - 2E\big[\hat{x}_d\big|x_d\big]x_d^T + x_dx_d^T\Big]\\ &= \mathrm{tr}\,\Big(\Gamma + \beta\Sigma_d^{-1}(\Sigma_d + \mu_d\mu_d^T)\Sigma_d^{-1}\beta^T - 2\beta\Sigma_d^{-1}\alpha^T\\ &\quad + \alpha\alpha^T - 2\beta\Sigma_d^{-1}(\Sigma_d+\mu_d\mu_d^T) - 2\alpha\mu_d^T + \Sigma_d + \mu_d\mu_d^T\Big)\\ &= \mathrm{tr}(\Sigma + \Sigma_d - 2\beta).\end{aligned}$$

Minimizing (14) over β, Σ, Λ subject to the constraint $\bar{D} \le D$ with the given α yields the desired result. $\square$

The proof of Theorem II.3 follows similar arguments in [2].

Proof of Theorem II.3. Consider a source code $\hat{x}_d = g\circ f(x_d)$ with code rate $\bar{R}$, blocklength L, and distortion $\bar{D} \le D$. Suppose this code admits a statistical characterization by a conditional density $p(\hat{x}_d|x_d) \in \mathcal{A}$. By definition, this source code reproduces a source sequence $x_d^L := \{x_d^{(j)}\}_{j=1}^L$ as $\hat{x}_d^L := \{\hat{x}_d^{(j)}\}_{j=1}^L$, such that the average distortion D_j in reproducing $x_d^{(j)}$ satisfies $\frac{1}{L}\sum_{j\in[L]} D_j \le D$.

We want to show that $\bar{R} \ge R_{\mathcal{A}}(D) \ge R(D)$. The latter inequality follows from the fact that $R_{\mathcal{A}}(D)$ is a linearly constrained convex optimization over a closed subset of $\mathbb{R}^{n\times n}\times \mathsf{S}^n_+$. To see the first inequality, note that $\frac{1}{L}\sum_{j=1}^L I(x_d^{(j)};\hat{x}_d^{(j)}) \le \bar{R}$ (cf. [2, Theorem 3.2.2]) and that, by construction $R(D_j) \le R_{\mathcal{A}}(D_j) \le I(x_d^{(j)};\hat{x}_d^{(j)})$. Summing over $j\in[L]$ and noting that both R and $R_{\mathcal{A}}$ are convex in their arguments yields the desired result.

Finally, let $\tilde{C}$ be the capacity of a channel, in nats per channel use, and C its capacity in nats per source symbol. We have $M\tilde{C} = LC$, corresponding to transmission of x_d^L in M channel uses. Let u^M denote the sequence of M channel inputs obtained from x_d^L, and y^M the channel output sequence from which $\hat{x}_d^L$ is decompressed. It follows that $I(u^M;y^M) \le M\tilde{C}$. The Data Processing Theorem [10, Theorem 4.3.3] yields $I(x_d^L;\hat{x}_d^L) \le I(u^M;y^M) \le LC$. Applying again the above convexity argument yields

$$R(D) \le R_{\mathcal{A}}(D) \le I(x_d;\hat{x}_d) \le I(u;y) \le C. \qquad \square$$

References

[1] MOSEK ApS. *MOSEK Fusion API for Python 10.0.40*. 2022.

[2] Toby Berger. *Rate-Distortion Theory*. Information and System Sciences. John Wiley & Sons, Inc., Hoboken, NJ, USA, April 2003.

[3] Stephen Boyd, Laurent El Ghaoui, Eric Feron, and Venkataramanan Balakrishnan. *Linear Matrix Inequalities in System and Control Theory*. Society for Industrial and Applied Mathematics, January 1994.

[4] Roger W. Brockett. Minimum attention control. In *Proceedings of the 36th IEEE Conference on Decision and Control*, volume 3, pages 2628–2632, San Diego, CA, USA, 1997. IEEE.

[5] Yongxin Chen, Tryphon T. Georgiou, and Michele Pavon. Optimal Steering of a Linear Stochastic System to a Final Probability Distribution, Part I. *IEEE Transactions on Automatic Control*, 61(5):1158–1169, May 2016.

[6] Thomas M Cover and Joy A Thomas. *Elements of Information Theory*. John Wiley & Sons, Ltd, 2005.

[7] Steven Diamond and Stephen Boyd. CVXPY: A Python-embedded modeling language for convex optimization. *Journal of Machine Learning Research*, 17(83):1–5, 2016.

[8] Huanyu Ding and David A. Castanon. Optimal multi-vehicle adaptive search with entropy objectives. In *2015 53rd Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pages 780–785, Monticello, IL, September 2015. IEEE.

[9] Huanyu Ding and David A. Castañón. Optimal Solutions for Adaptive Search Problems with Entropy Objectives, August 2015.

[10] Robert Gallager. *Information Theory and Reliable Communication*. John Wiley & Sons, Inc., 1968.

[11] Abhishek Halder and Eric D.B. Wendel. Finite horizon linear quadratic Gaussian density regulator with Wasserstein terminal cost. In *2016 American Control Conference (ACC)*, pages 7249–7254, Boston, MA, USA, July 2016. IEEE.

[12] Alfred O. Hero, David A. Castañón, Douglas Cochran, and Keith Kastella, editors. *Foundations and Applications of Sensor Management*. Springer US, Boston, MA, 2008.

[13] C. Hsieh and R.E. Skelton. All covariance controllers for linear discrete-time systems. *IEEE Transactions on Automatic Control*, 35(8):908–915, Aug./1990.

[14] K. Kastella. Discrimination gain to optimize detection and classification. *IEEE Transactions on Systems, Man, and Cybernetics - Part A: Systems and Humans*, 27(1):112–116, January 1997.

[15] Timothy L. Molloy and Girish N. Nair. Entropy-Regularized Partially Observed Markov Decision Processes. *IEEE Transactions on Automatic Control*, pages 1–8, 2023.

[16] Takashi Tanaka, Kwang-Ki K. Kim, Pablo A. Parrilo, and Sanjoy K. Mitter. Semidefinite Programming Approach to Gaussian Sequential Rate-Distortion Trade-Offs. *IEEE Transactions on Automatic Control*, 62(4):1896–1910, April 2017.

[17] Eric Wendel, John Baillieul, and Joseph Hollmann. On the complexity of linear systems: An approach via rate distortion theory and emulating systems. In *2023 American Control Conference (ACC)*, pages 2800–2805, 2023.

[18] J. Wolf and J. Ziv. Transmission of noisy information to a noisy receiver with minimum distortion. *IEEE Transactions on Information Theory*, 16(4):406–411, July 1970.

[19] J.-H. Xu and R.E. Skelton. An improved covariance assignment theory for discrete systems. *IEEE Transactions on Automatic Control*, 37(10):1588–1591, Oct./1992.