Hard Limits on Robust Control over Delayed and Quantized Communication Channels with Applications to Sensorimotor Control

Yorie Nakahira(依恵), Nikolai Matni(日光来), John C. Doyle(道耀)

***Abstract*— The modern view of the nervous system as layering distributed computation and communication for the purpose of sensorimotor control and homeostasis has much experimental evidence but little theoretical foundation, leaving unresolved the connection between diverse components and complex behavior. As a simple starting point, we address a fundamental tradeoff when robust control is done using communication with both delay and quantization error, which are both extremely heterogeneous and highly constrained in human and animal nervous systems. This yields surprisingly simple and tight analytic bounds with clear interpretations and insights regarding hard tradeoffs, optimal coding and control strategies, and their relationship with well known physiology and behavior. These results are similar to reasoning routinely used informally by experimentalists to explain their findings, but very different from those based on information theory and statistical physics (which have dominated theoretical neuroscience). The simple analytic results and their proofs extend to more general models at the expense of less insight and nontrivial (but still scalable) computation. They are also relevant, though less dramatically, to certain cyber-physical systems.**

I. Introduction

Figures 1 and 2 minimally sketch some of the extreme heterogeneity that is present in the speed and resolution of the human nervous system at both the system and component level. While neuroscience has accumulated seemingly endless details underlying these cartoons, there is no theoretical foundation that rigorously connects the system and component scale tradeoffs, a situation we aim to change here.

Figure 1 shows three important control loops for visual tracking in the presence of object and head motion that illustrates the extreme heterogeneity in system level behavior (timescales span $10^6\times$). This is a popular experimental system because many elements are easily manipulated. To demo the effects of these loops, place your hand in front of your face, close enough to clearly see the fine details of the lines in your palm. Then move your hand horizontally back and forth, increasing the frequency until the lines blur, which typically occurs at a few Hertz. In contrast, hold your hand still and shake your head (in a "no" pattern) at increasing frequencies until blurring occurs. This typically occurs at a much higher frequency than when the hand is moved, and most people are limited by the speed of head motion before there is significant blurring. Hand motion is tracked by slow (delay $>$ 200ms) vision which uses much of the cortex, whereas head motion is tracked by a completely separate fast vestibulo-ocular reflex (VOR, delay $\approx$ 10ms). Both systems

Yorie Nakahira, Nikolai Matni, and John C. Doyle are with the Control and Dynamical Systems Department of California Institute of Technology. Email: {ynakahir,doyle,nmatni}@caltech.edu

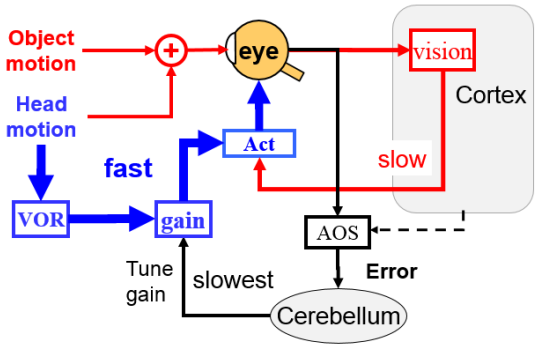

Fig. 1: Control of visual tracking via the actuation ("Act") of eye muscles to compensate for object and head motion. Object tracking uses slow (cortical) visual feedback (delay $>$ 200ms) whereas head uses the separate, fast vestibulo-ocular reflex (VOR, delay $\approx$ 10ms). The gains of these two loops must match, so neural mechanisms involving the cerebellum and the accessory optical system (AOS) tune the VOR gain (minutes to days).

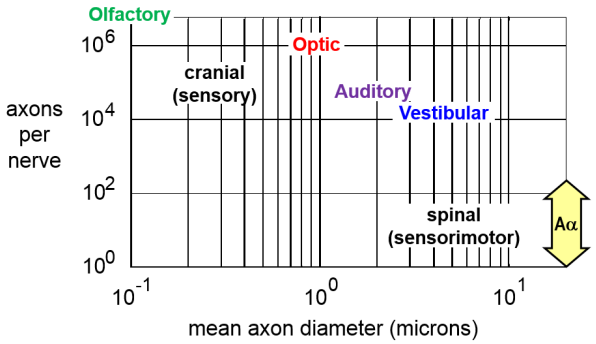

Fig. 2: Axons per nerve ($\propto$ resolution) versus mean axon diameter ($\propto$ speed) for four key cranial nerves, and the largest $(A\alpha)$ sensorimotor axon which occurs in spinal and peripheral nerves (in copies from 1 to hundreds).

are needed for athletes (or hunters) as objects can be at any distance and move at high speeds, but the greatest angular rates that must be tracked by eye movements arise from the motion of the athlete's own head and body, for which vision is too slow. Since vision and VOR gains must match, the slow tuning mechanism shown involving the cerebellum is necessary (but can be disrupted by e.g. prism glasses or ethanol).

While this demo is trivial and artificial, what it illustrates regarding the source of errors at the extremes of performance (when the hand begins to blur) is both fundamental and essential. Sensorimotor control performance is dominated by tradeoffs at the hardware level in the nervous system between temporal and spatial resolution. The hand (or other object with fine detail) at a sufficient distance, or in the visual periphery, will also blur due to resolution limits even when there is no motion. So either delays in tracking motion or limits of visual resolution, or both, can cause blurring, and it

is the tradeoff between speed and resolution (temporal versus spatial [1]) that is the most fundamental to understanding sensorimotor control (and Figure 1). Of course, blurring per se is of consequence only when it effects the ability to take timely and accurate action involving muscles (and tools), so the speed/resolution tradeoff extends throughout sensorimotor control, and not just perception.

Speed versus resolution tradeoffs also occur at the neuron level [2]. Figure 2 shows axon size versus number for four of the human's twelve cranial nerves, plus a single large axon of the types that appear in 31 spinal nerves. These nerves are bundles of fibers (axons) that carry action potentials throughout the body, and their composition is extremely variable. Axons are extreme in size and number, with peripheral nerves having single large ($> 20\mu$) axons, while the olfactory nerve has 6 million small (mean $\approx .1\mu$) axons. So axons vary $> 200\times$ in diameter, and $> 40,000\times$ in area (which is more relevant to the metabolic cost to build and maintain the nerves). The result is an extremely broad tradeoff in action potential propagation speed ($\propto$ axon diameter) and nerve resolution ($\propto$ number of axons), as well as spike rate (not shown in Figure 2 but also roughly $\propto$ axon diameter).

Our main goal is a theoretical foundation to rigorously connect the extreme tradeoffs in Figures 1 and 2. We start with a minimal but physiologically motivated model and derive simple analytic formulas for the impact of delay and quantization on robust performance, and argue that the resulting tradeoffs between speed and resolution are essential the function and evolution of the human nervous system. We later discuss how the math scalably extends to more complex and biologically plausible models. While a significant departure from the theories that have dominated neuroscience, every piece of our new framework is remarkably well established in neuroscience and robust control theory, if not previously integrated or appreciated. The literature is so extensive that a scholarly review is beyond the scope of this short paper, but will hopefully be the subject of future work. Here we will aim only to give a few pointers to relevant review papers and books, wherein more complete reference lists can be found.

Notation: Let $\mathbb{N}$ be the set of non-negative integers, $\mathbb{C}$ be the set of complex numbers. We use lower case boldface letters to denote sequences, i.e., $\mathbf{x} = (x(0), x(1), \dots)$, and denote the space of all such bounded sequences as l_∞. This is a Banach space when equipped with the norm $\|x\|_\infty \triangleq \sup_{t\in\mathbb{N}} |x(t)|$. We use upper case boldface letters such as $\mathbf{P}$ to denote maps from l_∞ to l_∞. We define the L_1 norm of such a map $\mathbf{P}$ to be $\|\mathbf{P}\|_{L_1} := \sup_{\|\mathbf{w}\|_\infty \le 1} \|\mathbf{P}\mathbf{w}\|_\infty$. Finally, for a sequence $\mathbf{x}$, we use $x(t_a : t_b)$ to denote the truncated sequence from time t_a to t_b, i.e., $x(t_a : t_b) = \{x(t_a), x(t_a + 1), \dots, x(t_b - 1), x(t_b)\}$.

II. Minimal Model and Theory

In this section we derive robust performance limits for a simplified model of sensorimotor control and communication that is delayed and quantized. We later use these results to connect Figures 1 and 2, and the issues that they raise.

A. Simplified Model

Consider an initial minimal model with dynamics

$$\begin{aligned} x(t+1) &= ax(t) + w(t - T_w) + \mathbf{Q}(u(t - T_u)) \\ u(t) &= \mathbf{K}(x(0:t), w(0:t), u(0:t-1)), \end{aligned} \tag{1}$$

where $x(t) \in \mathbb{R}$ is the state, $w(t) \in \mathbb{R}$ is the disturbance, $u(t) \in \mathbb{R}$ is the control action generated by the controller $\mathbf{K}$, and $\mathbf{Q} : \mathbb{R} \to \mathscr{S}_R$, for $\mathscr{S}_R \subset \mathbb{R}$ a finite set of cardinality 2^R, is a quantizer that limits communication between the controller and the actuator to R bits/sampling interval. The form of the control law in system (1) implies that the controller is *Full Information* (FI), as the control signal $u(t)$ is allowed to depend on all current and past states $x(0 : t)$, current and past disturbances $w(0 : t)$ and past control actions $u(0 : t - 1)$.

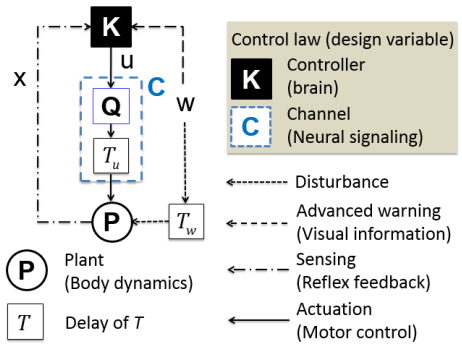

Fig. 3: Feedback system model for sensorimotor control.

A schematic for this model is shown in Fig. 3, where we use $\mathbf{P}$ to denote the plant defined by equation (1). The control signal $\mathbf{u}$ is transmitted to the actuator (colocated with the physical plant $\mathbf{P}$) via the communication channel $\mathbf{C}$, which is defined by the composition of the quantizer $\mathbf{Q}$ with the delay block T_u. This delay block implies that the controller command $u(t)$ takes $T_u (\ge 0)$ sampling intervals to reach and be executed by the actuators, i.e., $u(t)$ only affects the plant $T_u + 1$ sampling intervals later. We initially assume that the channel $\mathbf{C}$ is memoryless and stationary with rate R, allowing us to restrict the quantizer $\mathbf{Q}$ to be memoryless and static as well. Note that because $\mathbf{Q}$ and the delay block commute (as $\mathbf{Q}$ is assumed to be static) the dynamics (1) and Fig. 3 are indeed consistent. We assume $\|\mathbf{w}\|_\infty \le 1$ and $x(0) = 0$. The disturbance is known to the controller with an advanced warning of $T_w (\ge 0)$ sampling intervals, i.e., the controller has access to $w(0 : t)$ even though the disturbance only affects the plant $T_w + 1$ sampling intervals later.

The robust control problem can then be posed as

$$\begin{aligned} \underset{(\mathbf{K},\mathbf{Q})\in\mathcal{Q}_R}{\text{minimize}} \quad & \sup_{\|\mathbf{w}\|_\infty \le 1} \|\mathbf{x}\|_\infty \\ \text{s.t.} \quad & \text{dynamics (1)} \end{aligned} \tag{2}$$

where $\mathcal{Q}_R$ is the space of control laws defined by the pair of mappings $(\mathbf{K}, \mathbf{Q})$, with $\mathbf{Q}$ constrained to be a static memoryless quantizer of rate R, i.e., $\mathbf{Q} : \mathbb{R} \to \mathscr{S}_R$. This

cost function is standard in L_1 robust control [3], except that a communication channel $\mathbf{C}$, composed of a quantizer $\mathbf{Q}$ and a delay T_u, is inserted into the feedback loop. Perhaps surprisingly, this problem formulation still allows for a simple and intuitive analytic solution.

Remark 1: Without quantization or delay, the control law $u(t) = -ax(t) - w(t)$ ensures that $x(t+1) = 0$. Thus any errors in the state is a direct consequences of quantization and/or delay, or to saturation of the control signal $\mathbf{u}$.

B. Fundamental Limits due to Delay and Quantization

In this subsection, we provide an exact solution to the robust control problem (2): in particular, we show that the worst-case state-deviation can be expressed as a function of the plant pole a, the channel rate R, and the *net delay* of the system $T \triangleq T_u - T_w$. The achievable performance takes a different form depending on the net delay regime that the system is operating under. When the net delay T is positive ($T > 0$), this corresponds to a system in which the control action $u(t)$ can only affect the plant T sampling intervals *after* the disturbance $w(t)$ affects the state. Conversely, when the net delay T is non-positive ($T \leq 0$), this corresponds to a system in which there is *advanced warning* of the disturbance, allowing the controller to act in advance. These two qualitatively different cases are treated separately.

Theorem 1: Suppose that $2^R > |a|$. Then the minimal state-deviation achievable in robust control problem (2) is

$$\begin{cases} \sum_{i=1}^{T} |a^{i-1}| + |a^T| \Big(2^R - |a|\Big)^{-1} & \text{if } T > 0 \\ \Big(2^R - |a|\Big)^{-1} & \text{if } T \leq 0. \end{cases} \tag{3}$$

Conversely, if $|a| \geq 2^R$, then the system cannot be stabilized, and the optimal value to optimization problem (2) is infinite.

Remark 2: In the plant dynamics (1), and consequently in Theorem 1, we impose *a priori* that the channel be defined by a static memoryless quantizer $\mathbf{Q} : \mathbb{R} \to \mathscr{S}_R$. In §VI, we state and prove a more general result, which shows that a static and memoryless quantizer is indeed optimal within a broad class of encoder/decoder pairs operating subject to a data-rate limit of R bits per sampling interval. Theorem 1 then follows as a special case of this more general result.

The performance limits (3) are remarkably simple and intuitive. The net warning case ($T \leq 0$) has only one term due to quantization, with the stabilizability condition $2^R > |a|$ well-known from existing literature [4]. With no dynamics ($a = 0$) this reduces to a trivial rate distortion theorem with error 2^{-R}.

The net delayed case ($T < 0$) is more interesting, with the first term due to the delay alone, and the second term the additional contribution due to quantization. As expected, both grow rapidly with increased net delay T and unstable $a > 1$, for reasons familiar and intuitive. In order to prove this part of the result, we establish a separation principle and show that the detrimental effects of the net delay T and the errors causes by the quantizer $\mathbf{Q}$ contribute to the worst-case state-deviation in an additive and independent way.

Specifically, we show that the state $x(t)$ can be decomposed into two components, $x_d(t)$ and $x_q(t)$, such that $x(t) = x_d(t) + x_q(t)$, where $x_d(t)$ is the state-deviation due solely to the net delay $T > 0$, and $x_q(t)$ is the state-deviation due solely to the quantizer $\mathbf{Q}$. Our first step is to notice that the net delay $T > 0$ implies that the effects of $w(t - T_u : t - T_w)$ are independent of the signal $u(0 : t - T_u - 1)$, as this information about the disturbance is not yet available to the controller for times $t < t - T_u$. It follows that the effects of this subset of the disturbance signal $\mathbf{w}$ on $x(t)$ is

$$x_d(t) \triangleq \sum_{i=1}^{T} a^{i-1} w(t - i - T_w). \tag{4}$$

Defining $x_q(t) \triangleq x(t) - x_q(t)$, it is straightforward to show that $x_q(t)$ is strictly a function of $w(0 : t - T_u - 1)$ and $u(0 : t - T_u - 1)$. It then follows that optimization problem (2) can be rewritten as

$$\sup_{\|\mathbf{w}\|_\infty \leq 1} \|\mathbf{x}_d\|_\infty + \inf_{(\mathbf{K},\mathbf{Q}) \in \mathcal{Q}_R} \sup_{\|\mathbf{w}\|_\infty \leq 1} \|\mathbf{x}_q\|_\infty$$

subject to the dynamics (1). To see this, it suffices to note that for a fixed t, disjoint subsets of the disturbance signal $\mathbf{w}$ affect the delay induced error $x_d(t)$ and the quantization induced error $x_q(t)$, and that $x_d(t)$ is independent of the chosen control law. Finally, the result follows by showing

$$\sup_{\|\mathbf{w}\|_\infty \leq 1} \|\mathbf{x}_d\|_\infty = \sum_{i=1}^{T} |a^{i-1}|, \text{ and} \tag{5}$$

$$\inf_{(\mathbf{K},\mathbf{Q}) \in \mathcal{Q}_R} \sup_{\|\mathbf{w}\|_\infty \leq 1} \|\mathbf{x}_q\|_\infty = |a^T| \Big(2^R - |a|\Big)^{-1}. \tag{6}$$

Although we do not emphasize this fact in Theorem 1, our proof also provides an exact expression for the optimal controller $\mathbf{K}$, allowing us to explicitly quantify $\sup_{\|\mathbf{w}\| \leq 1} \|\mathbf{u}\|_\infty$. We use this fact to comment on the biological implications of actuator saturation in §III.B.

III. Neuro-specific Details

Our model assumes that the channel delay T_u and the data rate R are given, but the performance limits (3) allow us to explore neuro-specific tradeoffs between T_u and R. Concretely, reconsider Figure 1 but for a task like stalking and chasing prey while hunting with primitive weapons, a task humans are well adapted to, and presumably drove the evolution of much of our physiology and nervous system, particularly those aspects that differ most from our ape ancestors (and current robots). Modern athletic scenarios include chasing and catching a frisbee or ball. Here the "object" is moving but at least initially at a remote distance, so vision is providing sensorimotor control with substantial advanced (but possibly uncertain) warning about object motion, allowing for advanced planning. Beyond the simplified tracking problem that we model by the dynamics (1), several other complex control tasks must occur. For example, the VOR system must control eye tracking despite head motion due to running and other disturbances, a myriad of other reflexes are needed to run and catch, and neuroendocrine control systems

must maintain internal homeostasis [5]. In contrast to our simple model (1), these systems are distributed and work in parallel, leading to further delays in reaction time.

Thus there are two very different and opposite kinds of delays, modeled initially here by T_u and T_w. Internal delays in vision, VOR, and other reflexes are due to sensing, computation, communication, and actuation, all using slow neural hardware and making control difficult. These are diverse and distributed but in our initial model and theorem are lumped into the single delay T_u. Conversely, sensors like vision and hearing can often sense an external disturbance $w(t)$ such that there is a delay of T_w until the disturbance impacts the plant, thus giving the controller advanced warning and making disturbance rejection easier. These disturbances and sources of advanced warning are also diverse, but we initially just consider a single disturbance and delay.

A. Instabilities, saturation, and worst-case scenarios

A universal feature of almost all biological systems is that the physical "plants" being controlled are unstable. In our minimal model this corresponds to $|a| > 1$ and large a aggravates all the tradeoffs, but particularly those for the delayed cases. Concretely, while running or cycling, all variables of interest from upright balance to internal temperature and arterial blood pressure and oxygenation must be tightly controlled by very complex neuroendocrine control systems or they will crash, often fatally.

Further aggravating the impact of instability is that biological actuators saturate, just like their engineering counterparts, so in addition to state and output error due to disturbances, delays, and quantization, we must also pay close attention to $\sup_{\|\mathbf{w}\|_\infty \le 1} \|\mathbf{u}\|_\infty$. For unstable systems ($|a| \ge 1$), the minimum stabilizing control saturation level (cf. Proof of Theorem 2 in §VI) is given by

$$|a^T| + |a^{T+1}| \Big(2^R - |a|\Big)^{-1} \tag{7}$$

Surprisingly, the optimal control action to minimize $\|\mathbf{x}\|_\infty$ in Theorem 1 is $u(t - T_u) = -a^T w(t - T_u) - a x_q(t)$, yielding a worst-case control cost that is exactly the same as the expression (7). Thus there is no tradeoff in this problem between minimizing worst case error versus control. This optimal $\|\mathbf{u}\|_\infty$ increases rapidly with delay and instability at a slightly faster rate than the state error. Since the control and error norms are so similar we focus on error, but keep in mind the importance of control saturation.

The biological importance of actuator saturation is one among many motivations for studying $\|\mathbf{u}\|_\infty$ and our other assumptions. While we can only raise the most basic questions here, the best problem formulation is a rich topic about which we hope to engage both controls and neuroscience researchers. We aim to show that limits on the achievable robust performance of the human nervous system (HNS) are central to understanding its evolution and organization. In considering the various concrete scenarios described above, none of the HNS features matter if the actuators (e.g. muscles) are not strong or fast enough to execute the required actions. In this context $\|\mathbf{u}\|_\infty$ is more relevant than $\|\mathbf{u}\|_2$ or variance, although more internal variables, such muscle fatigue, are essential in defining more realistic models.

All of our motivating scenarios suggest $\|\mathbf{x}\|_\infty$ and worst case disturbances are also appropriate, whether predator or prey, athlete, runner, or biker, or threading needles. Surprisingly, the resulting proofs are simpler than their H_∞ and Bode Integral counterparts (emphasized in [6]), and the inclusion of channels seems more straightforward (compare with [7]). Most obviously different in this l_∞/L_1 theory is the lack of frequency domain or stochastic interpretations, but we expect this to be more feature than bug, though stochastic disturbances and crash probabilities might make sense in some situations. The overwhelmingly most attractive feature of our model and theory is that we can easily add essential features of neural physiological tradeoffs and plug these directly into tradeoffs of data rate R against net delay T in (3) and (7), allowing us to make connections between the component and system levels of sensorimotor control.

B. Speed vs. Accuracy Tradeoff

We now add a tradeoff between temporal and spatial resolution in neural signaling to our model via the net delay T and data rate R that we believe is the first important constraint in explaining the large scale features discussed above. The nervous system communicates between components and the body with a variety of nerves, which are bundles of axons. Axons are the wiring by which spiking neurons communicate long range using action potentials, and it is possible to derive some rough tradeoffs from well-known physiology. Figure 2 shows some of the tremendous diversity of axon numbers and sizes among the cranial and peripheral nerves. We argue that much of this arises due to hard constraints on speed versus accuracy.

We suppose that our channel $\mathbf{C}$ (cf. Fig. 3) is a single nerve with uniform signaling delay T_s, and assume that the total delay T_u is the sum $T_u = T_s + T_c$ with an additional fixed delay T_c due to grey matter computation and other communications. Initially we assume that T_c is fixed and given, and that T_s is variable and depends on the nerve composition, as in Figure 2. We also assume that the physical layout of white and grey matter is fixed and given: hence the lengths of axons are also fixed. We assume that a nerve has fixed cross-sectional area α, and use this as a measure of the resources devoted to the nerve. Assume each nerve is composed of m axons of uniform radius ρ, such that $\alpha \approx \pi m \rho^2$. If we further assume a digital, error-free, discrete-time code operating at some maximal firing rate ϕ, with 1 represented by the presence of an action potential, and 0 by its absence, the total bit rate $R = m\phi$ with delay T_s. We want a direct relationship between R and T_s.

Because of energy considerations [2], both the propagation speed and firing rate of an action potential in a myelinated axon is roughly proportional to its radius ρ, so $\phi \propto \rho$ and the signal transmission delay $T_s \propto 1/\rho$. Thus $R = m\phi \propto m\rho \propto \alpha/\rho \propto \alpha T_s$, which we will write simply as

$$R = \lambda_\alpha T_s \tag{8}$$

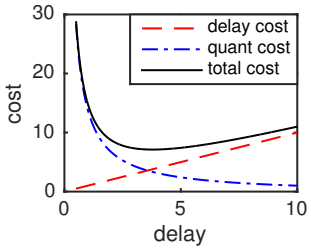

Figure 4: Impact of speed versus accuracy. The cost of delay $\sup_{\|w\|_\infty \le 1} \|\mathbf{x}_d\|_\infty$, the cost of quantization $\sup_{\|w\|_\infty \le 1} \|\mathbf{e}\|_\infty$, and the total cost $\sup_{\|x\|_\infty \le 1} \|\mathbf{x}_d\|_\infty$ is shown with varying delay T_s when $\lambda = .1$, $a = 1$.

where λ_α is a resource measure that scales with the axon area α. (Bit errors would reduce R and coding in the of variation of inter-spike timing would increase R.) This is most plausible for the large nerves and axons on the right half of Figure 2 (i.e. reflexes) [2]. No bit errors is surprisingly plausible for large axons, and worst case bit flips can exactly though expensively be removed with Hamming codes. More seriously, no nerves can persistently fire at their maximal rate so some form of sparse coding would be needed. Both would reduce λ_α. More complex error and rate models would lead to more complicated formulas $R = \lambda_\alpha(T_s)$ but we do not expect these changes to dramatically alter the consequences below. Validating this will be the focus of future work.

IV. Implications

A. Impact of Speed and Accuracy on System Performance

Next we explore the surprisingly rich consequences of the constraint $R = \lambda_\alpha T_s$ on our minimal model of sensorimotor control using Theorem 1. For simplicity, We write λ from now on as the resource dependence is understood.

Corollary 1: If $R = \lambda T_s$ and $T_u \triangleq T_s + T_c$, then the optimal optimal performance specified in Theorem 1 becomes

$$\begin{cases} \sum_{i=1}^{T} |a^{i-1}| + |a^T| \left(2^{\lambda T_s} - |a|\right)^{-1} & \text{if } T > 0 \\ \left(2^{\lambda T_s} - |a|\right)^{-1} & \text{if } T \le 0. \end{cases}$$

Figure 4 shows the system performance when varying delay T_s (and thus channel rate R) for $T_c = T_w = 0$ and a fixed resource level α. Increased delay increases the delay error term $\sup_{\|\mathbf{w}\|_\infty \le 1} \|\mathbf{x}_d\|_\infty$ but reduces the quantization error term $\sup_{\|\mathbf{w}\|_\infty \le 1} \|\mathbf{x}_q\|_\infty$. Consequently, the optimal system level performance is achieved at intermediate levels of delay and channel rate. Because of the exponential dependence there is no analytic formula for the optimum, but the error is convex and the minimum easily found numerically. Next we consider in more detail the consequences of these formulas by varying the additional delays and plotting the resulting optimal errors, bits, and delay.

B. Delayed vs Warned System

Figure 5 shows the optimal delays T_s (and resulting net delay T) and channel rate $R = \lambda T_s$ that achieves the minimum total error when varying $T_w \ge 0$ and $T_c \ge 0$ separately in the two special cases (i) $T = T_u - T_w \le 0$ (warned) and (ii) $T = T_s + T_c > 0$ (delayed). What results are clearly two distinct regimes with distinct physiology. When the computation delay T_c is greater than 0, the system has

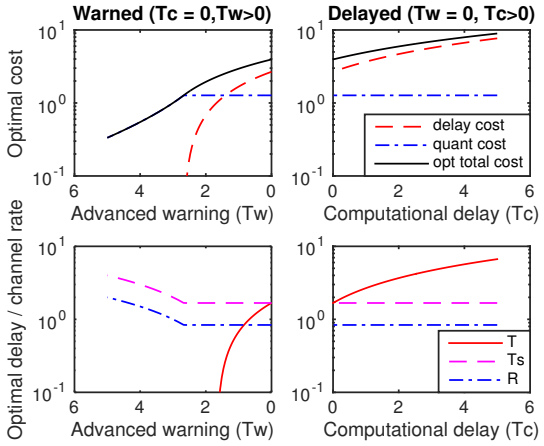

Fig. 5: Delayed versus warned system.

a net delay T and the delay cost increasingly dominates the total cost, leading to both the data rate R and signaling delay T_s becoming constant (with a large and constant radius ρ), independent of T_c. This corresponds to the reflexes on the right half of Figure 2 with nerves having a relatively few large axons. The total error, due mostly to delay, can be much larger than the disturbance. Concretely, in running or cycling on rough terrain or through heavy traffic, a relatively small but well placed perturbation to the foot or wheel can be amplified into a crash, even a fatal one – this effect gets worse at high speeds when the delay is relatively larger. Our nervous system invests in large nerves, axons, and muscles to avoid such crashes, consistent with the theory.

With increasing advanced warning $T_w > 0$ the net delay T becomes non-positive, and in this case the errors due to quantization increasingly dominate the total cost. Further, this total cost goes to zero as T_w increases, exactly the opposite of the delayed case. Further, as the advanced warning T_w increases, so does the data rate R, and consequently the axon radius ρ decreases (as $\alpha \approx \pi R \rho^2$ is fixed). This corresponds to the left side of Figure 2 with many relatively small axons. In running or cycling we can start with huge errors to remotely located objects, and given enough time drive them to zero. Here we are limited largely by the resolution of our vision or other sensing in accurately locating the object, again consistent with the theory.

Thus we have an extremely simple model that connects the system level requirements of advanced warning and planning (e.g. as enabled by vision) to the low-level control implemented by fast reflexes. We further introduce minimal neuro-motivated constraints on accurate versus rapid signaling, allowing us to connect our theoretical results to the known physiology illustrated in Figure 2. In the sequel we explore further aspects of this model, and introduce consider additional constraints and generalizations.

C. A Minimal Network

One of the most important features of a visual system (in Fig. 1) is its distributed nature, in which sensors, actuators, and computational components are interconnected via sparse

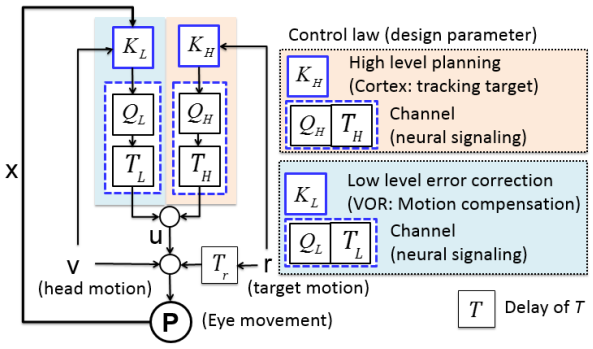

Fig. 6: Modeling Visual processes as a hierarchical control system.

communication. Figure 6 sketches a minimal model of this kind that is composed of two copies of each component in Figure 3. The plant dynamics are given by $x(t+1) = ax(t)+u(t)+w(t)$ except the disturbance is now composed of two terms $w(t) = v(t)+r(t-T_r)$, as is the control action $u(t) = u_L(t-T_L)+u_H(t-T_H)$, each generated by their own sensors, computing, and communication components. Visual trajectory planning is done through the control loop involving $\mathbf{Q}_H$ which is responsible for tracking, via the control signal $u_H(t)$, a visual target whose change in position is expressed by $\mathbf{r}$. We assume a very simplified view of vision whereby remote (in space) sensing means that $r(t)$ is seen but it takes T_r for the disturbance to arrive, effectively creating an advanced warning of T_r, though the physical details are all causal.

On the other hand, local (reflex) compensation is done through the control loop involving $\mathbf{Q}_L$. Disturbances such as those caused by body and head motion are captured by $\mathbf{v}$, and are sensed directly by the VOR, which computes a control action $u_L(t)$ to compensate. The control commands $(u_H(t), u_L(t))$ from both loops are sent to the plant through different signaling pathways, modeled by channels with rates R_H and R_L and delays T_H and T_L, respectively, after which their gains are summed to produced the final previously described control action $u(t) = u_L(t-T_L) + u_H(t-T_H)$.

Using the tradeoff (8) in both signaling pathways, and bounding $\|\mathbf{v}\|_\infty$ and $\|\mathbf{r}\|_\infty$ from above by 1 and δ, respectively, the optimal performance is then given by

$$\left\{\sum_{i=1}^{T_L}|a^{i-1}| + |a^{T_L}|\left(2^{R_L}-|a|\right)^{-1}\right\} + \delta\left(2^{R_H}-|a|\right)^{-1}.$$

This result follows by noting that the total system can be decomposed into two independent subsystems, corresponding to the $\mathbf{Q}_H$ and $\mathbf{Q}_L$ loops, and thus so can its performance. The first subsystem is a delayed system driven by $\mathbf{v}$ and controlled by $\mathbf{u}_L$, while the second subsystem is a warned system driven by $\mathbf{r}$ and controlled by $\mathbf{u}_H$. From our previous analysis, it is expected that the first system achieves better performance when its nerves are composed of a few large and fast axons, whereas the second system achieves better performance when its nerves are composed of many small and slow axons. This phenomena can be indeed observed in the real visual systems [2]. Specifically, the optic nerve has approximately 1M axons of mean diameter $0.64\mu m$ with CV $0.46\mu m$, while the 20K vestibular axons have mean diameter $2.88\mu m$ with CV 0.41, significantly larger and less numerous and slightly less variable.

D. Additional reading

The related literature in both neuroscience and control theory is vast and there is not space for a scholarly or historical review, but we'll give a few pointers for further reading, emphasizing reviews and books. While the study of tradeoffs has a long and rich history, there is surprisingly little connecting the levels of description in Figs. 1-3. Hardware tradeoffs as in Fig. 3 are beautifully and extensively explored in a new book [2]. Optimal (and more recently robust [8]) control theory have been extensively applied as in Fig. 3, but rarely connect to mechanisms even as in Fig. 1, even though experimentalists often interpret their findings using such diagrams (see [1] for particularly insightful examples). Indeed, the kind of reasoning we used here, including speed vs resolution tradeoffs, is ubiquitous among experimentalists, but is rarely formalized. In retrospect, it seems clear that a major disconnect has been that control and learning theory [8] treats delay, but information theory (e.g. [9]) largely ignores delay but dominates theory in the hardware layers [2], [1]. This paper seems to be the first attempt to impose minimal but plausible hardware constraints on system level robust performance. It is admittedly a small step but hopefully in a useful new direction, though one that will admittedly be strongly rejected by many mainstream theorists for its emphasis on the importance of delay in sensorimotor control. On the control theory side, we can now exploit a vast and exciting but also fragmented collection of results in distributed optimal control. In particular, recent results have identified a broad class of distributed control problems that are *convex* [10], [11] and that admit solutions that are *scalable* to compute and implement [12]. A key feature of these results is that the control problem becomes "easy" if certain architectural requirements (often related to density of actuation, sensing and communication) are met by the controller – tractable approaches to designing such favorable architectures have also been developed [13], [14]. An important remaining challenge is to combine these results with complementary ones from network control [4].

V. General Model and Proof

In this section, we show that even if we allow the controller to be time-varying and the encoder, channel and decoder to have memory, the control architecture assumed in system (1) and Fig. 4 remains optimal for the robust control problem (2). We state and prove this general result in this section, from which Theorem 1 follows immediately as a special case.

Consider a feedback system with dynamics

$$\begin{aligned} x(t+1) &= ax(t) + w(t-T_w) + u(t) \\ s(t) &= \mathbf{E}_t(x(0:t), w(0:t), s(0:t-1)) \\ s(t-T_u) &= \mathbf{C}(s(0:t-T_u)) \\ u(t) &= \mathbf{D}_t(s(0:t-T_u)) \end{aligned} \tag{9}$$

where $x(t) \in \mathbb{R}$ is the state, $w(t) \in \mathbb{R}$ is the disturbance, $s(t) \in \mathscr{S}_R$ is the codeword produced by the encoder $\mathbf{E}_t$, and $u(t) \in \mathbb{R}$ is the control action produced by the decoder $\mathbf{D}_t$. Here $\mathbf{E}_t$ is a possibly time-varying encoder that maps the state, noise and symbol history to a finite alphabet $\mathscr{S}_R$ of cardinality 2^R. The cardinality of the space $\mathscr{S}_R$ to which the encoder $\mathbf{E}_t$ and channel $\mathbf{C}$ map imply that R bits per sampling interval arrive at the decoder $\mathbf{D}_t$ with a delay of T_u. We allow the decoder to utilize all past codewords to decide the current control action, and we allow the overall control strategy to be time varying.

We assume that $x(0) = 0$, that $T_u \geq 0$, and that $T_w \geq 0$. The control action $u(t)$ is thus generated by the encoder $\mathbf{E}_t$ and the decoder $\mathbf{D}_t$. From the definition of $s(t)$ in equation (9), it follows that the controller is still FI. The cardinality of the space $\mathscr{S}_R$ to which the decoder $\mathbf{E}_t$ and channel $\mathbf{C}$ map imply that R bits per sampling interval arrive at the decoder $\mathbf{D}_t$ with a delay of T_u. We allow the decoder to utilize all past codewords to decide the current control action, and we allow the overall control strategy to be time varying.

The robust control problem of interest is then

$$\inf_{(\mathbf{E},\mathbf{D})\in\mathcal{G}_R} \quad \sup_{\|\mathbf{w}\|_\infty \leq 1} \|\mathbf{x}\|_\infty \tag{10}$$

where $\mathcal{G}_R$ is the space of control laws defined by sequences of mappings $\mathbf{E} = \{\mathbf{E}_t\}$ and $\mathbf{D} = \{\mathbf{D}_t\}$ as described after (9). We show below the solution presented in Theorem 1 is in fact optimal for the robust control problem (10).

Theorem 2: A control law of the form $(\mathbf{K}, \mathbf{Q}) \in \mathcal{Q}_R$, is the solution to the robust control problem (10), *i.e.*,

$$\inf_{(\mathbf{E},\mathbf{D})\in\mathcal{G}_R} \sup_{\|\mathbf{w}\|_\infty \leq 1} \|\mathbf{x}\|_\infty = \inf_{(\mathbf{K},\mathbf{Q})\in\mathcal{Q}_R} \sup_{\|\mathbf{w}\|_\infty \leq 1} \|\mathbf{x}\|_\infty. \tag{11}$$

In particular, the encoder $\mathbf{E}$ can be taken to be the composition of a state-feedback controller $\mathbf{K}$ with a static memoryless quantizer $\mathbf{Q}$, and the decoder $\mathbf{D}$ can be taken to satisfy $\mathbf{D}_t(s(0:t-T_u)) = s(0:t-T_u)$.

In the interest of brevity, we only prove the result here for the case of $T > 0$ – a nearly identical (and slightly simpler) argument holds for the case of $T \leq 0$.

We begin with a technical lemma.

Lemma 1: Assume that $T > 0$, and let

$$\begin{array}{rcl} x_q(t) & \triangleq & s(t-T_u-1) - u^*(t-T_u-1) \\ u^*(t-T_u) & \triangleq & -a^T w(t-T_u) - a x_q(t), \end{array} \tag{12}$$

and $x_d(t)$ be as defined in equation (4). Then if $\mathbf{D}_t(s(0:t-T_u)) = s(0:t-T_u)$, it holds that

$$\begin{array}{rcl} x(t+1) & = & x_d(t+1) + x_q(t+1) \\ & = & x_d(t+1) + s(t-T_u) - u^*(t-T_u). \end{array} \tag{13}$$

In particular, $u^*(t)$ is a *causal* function of $w(t)$, $s(t-1)$ and $u^*(t-1)$.

Proof: The claim holds trivially for all $t \leq T_u$, as $s(t) = 0$ and $u^*(t) = 0$ for $t < T_u$. We now proceed by induction, and assume that relation (1) holds at time t. It can be checked by direct calculation that

$$a x_d(t) + w(t-T_w) = x_d(t+1) + a^T w(t-T_u). \tag{14}$$

It then follows that

$$\begin{array}{rcl} x(t+1) & = & a x(t) + w(t-T_w) + s(t-T_u) \\ & = & a x_d(t) + w(t-T_w) + a x_q(t) + s(t-T_u) \\ & = & x_d(t+1) + a^T w(t-T_u) \\ & & \quad + a x_q(t) + s(t-T_u) \\ & = & x_d(t+1) + s(t-T_u) - u^*(t-T_u) \\ & = & x_d(t+1) + x_q(t+1), \end{array}$$

where the first equality follows from the dynamics (9) and the assumed form of the decoder $\mathbf{D}$, the second from the induction hypothesis, the third from relation (14), the fourth from the definition (12) of $u^*(t-T_u)$ and the fifth from the definition (12) of $x_q(t+1)$. ∎

With this technical lemma in hand, we may now make formal the argument made after Theorem 1.

Proof: (Theorem 2) Using Lemma 1, and following the argument after Theorem 1, the optimal performance can be written as:

$$\begin{array}{l} \inf\limits_{(\mathbf{E},\mathbf{D})\in\mathcal{G}_R} \sup\limits_{\|\mathbf{w}\|_\infty \leq 1} \|\mathbf{x}\|_\infty = \\ \quad \inf\limits_{(\mathbf{E},\mathbf{D})\in\mathcal{G}_R} \sup\limits_{\|\mathbf{w}\|_\infty \leq 1} \|\mathbf{x}_d\|_\infty + \inf\limits_{(\mathbf{E},\mathbf{D})\in\mathcal{G}_R} \sup\limits_{\|\mathbf{w}\|_\infty \leq 1} \|\mathbf{x}_q\|_\infty \\ \quad = \sup\limits_{\|\mathbf{w}\|_\infty \leq 1} \|\mathbf{x}_d\|_\infty + \inf\limits_{(\mathbf{E},\mathbf{D})\in\mathcal{G}_R} \sup\limits_{\|\mathbf{w}\|_\infty \leq 1} \|\mathbf{x}_q\|_\infty. \end{array} \tag{15}$$

Recall that for a fixed t, $x_d(t)$ is the component of the state response that is uncontrollable due to the net delay T. Thus, when $T \leq 0$ (advanced warning case), it follows that $x_d(t)$ is identically zero, and when $T > 0$, the terms $x_d(t)$ and $x_q(t)$ are independent – specifically, $x_d(t)$ is a linear function of $w(t-T_u : t-T_w-1)$, and $x_q(t)$ is the error due to quantization in the control action that is attempting to cancel the effects of $w(0 : t-T_u-1)$ on the state. Thus, the supremums of $\sup_{\|\mathbf{w}\|_\infty \leq 1} |x_d(t)|$ and $\sup_{\|\mathbf{w}\|_\infty \leq 1} |x_q(t)|$ are simultaneously attainable by a single disturbance signal $\mathbf{w}$, as disjoint subsets of the disturbance affect each of $x_d(t)$ and $x_q(t)$ separately.

Thus it remains to show that the value of these terms are given by equations (5) and (6). The equality (5) follows immediately from definition (4). The equality (6) is proved by computing upper and low bounds on $\inf_{(\mathbf{E},\mathbf{D})\in\mathcal{G}_R} \sup_{\|\mathbf{w}\|_\infty \leq 1} \|\mathbf{x}_q\|_\infty$.

Achievable performance (upper bound). Let the encoder be defined by the mapping $\mathbf{E}$ that generates codeword $s(t)$ according to the following procedure. Let $u^*(t-T_u) = -a^T w(t-T_u) - a x_q(t)$ denote the nominal control law that would be applied if no data-rate limits where present. Define $s(t) = \mathbf{E}_t(x(0:t), w(0:t), s(0:t-1))$ as follows:

$$s(t-T_u) = \begin{cases} (1-\frac{1}{2^R})\Psi & \text{if } u^*(t-T_u) \in [(1-\frac{2}{2^R})\Psi, \infty] \\ (1-\frac{3}{2^R})\Psi & \text{if } u^*(t-T_u) \in [(1-\frac{4}{2^R})\Psi, (1-\frac{2}{2^R})\Psi) \\ \quad \vdots & \\ (-1+\frac{1}{2^R})\Psi & \text{if } u^*(t-T_u) \in [-\infty, -(1-\frac{2}{2^R})\Psi) \end{cases},$$

where $\Psi \triangleq |a^{T+1}|(2^R - |a|)^{-1} + |a^T|$. It is clear that $s(t)$ can take one of the at most 2^R values in $\mathscr{S}_R \triangleq \{(1-\frac{1}{2^R})\Psi, (1-$

$\frac{3}{2^R})\Psi, \ldots, (-1+\frac{1}{2^R})\Psi\}$. It then follows that the output $s(t)$ from the channel $\mathbf{C}$ is simply given by a suitably delayed form of the above expression, and that the optimal decoder $\mathbf{D}$ is simply the identity map taking $\mathscr{S}_R$ into $\mathbb{R}$, generating the control action $u(t) = \mathbf{D}_t(s(0:t-T_u)) = s(t-T_u)$.

We use mathematical induction to show that $\sup_{\|w\|_\infty \leq 1} \|\mathbf{x}_q\|_\infty \leq |a^T|(2^R - |a|)^{-1}$. Assume that $|x_q(t)| \leq |a^T|(2^R - |a|)^{-1}$ holds for time $x_q(t)$ (this is trivially true for $t = 0$). It follows that $|u^*(t-T_u)| = |-a^T w(t-T) - a x_q(t)| \leq |a^T| + |a^{T+1}|(2^R - |a|)^{-1} = \Psi$. Further, from Lemma 1, we have that $x_q(t+1) = s(t-T_u) - u^*(t-T_u)$; since $u^*(t-T_u) \in [-\Psi, \Psi]$, we then have that

$$|x_q(t+1)| = |s(t-T_u) - u^*(t-T_u)| \leq 2^{-R}\Psi = |a^T|(2^R - |a|)^{-1}$$

holds at time $t+1$, thus proving the claim. In particular, we have that

$$\inf_{(\mathbf{E},\mathbf{D}) \in \mathcal{G}_R} \sup_{\|\mathbf{w}\|_\infty \leq 1} \|\mathbf{x}_q\|_\infty \leq |a^T|\big(2^R - |a|\big)^{-1}.$$

Performance limit (lower bound). Seeking a contradiction, we suppose that there exists a control law $(\mathbf{E}, \mathbf{D})$ that achieves $\sup_{\|\mathbf{w}\|_\infty \leq 1} \|\mathbf{x}_q\|_\infty = \Gamma$ for some Γ strictly smaller than $|a^T|(2^R - |a|)^{-1}$.

Following a similar argument to that of Lemma 1 we have that $x_q(t+1) = u(t) - u^*(t-T_u)$, and that $u^*(t-T_u)$ can only take on values lying in the set

$$\begin{aligned}\mathscr{E} &\triangleq \{q \in \mathbb{R} : q = aq + a^T w, q \in [-\Gamma,\ \Gamma], w \in [-1,\ 1]\} \\ &= \Big[-|a|\Gamma - |a^T|,\ \ |a|\Gamma + |a^T|\Big].\end{aligned}$$

Since the channel can only transmit R bits of information, there are at most 2^R different control actions that can be taken. As $u^*(t-T_u) = -a x_q(t) + -a^T w(t-T_u)$ can take any value in the interval, it follows that

$$\sup_{\|\mathbf{w}\|_\infty \leq 1} |u(t) - u^*(t-T_u)| \ \geq\ 2^{-R}\Big(|a|\Gamma + |a^T|\Big)^{-1}$$

is the best performance that one can achieve. Using $|a^T|(2^R - |a|)^{-1} > \Gamma$ twice we obtain

$$2^{-R}\Big(|a|\Gamma + |a^T|\Big)^{-1} > |a^T|\Big(2^R - |a|\Big)^{-1} > \Gamma,$$

contradicting our assumption that $\sup_{\|\mathbf{w}\|_\infty \leq 1} \|\mathbf{x}_q\|_\infty = \Gamma$. Thus we have shown the desired equality (6). Further the control law described above has a clean interpretation: a nominal control action $u^*(t)$ is computed and then quantized (via the static and memoryless encoder $\mathbf{E}$) before being transmitted across the channel. Note in particular that the resulting optimal controller/communication architecture is compatible with the Quantizer + Delay scheme proposed in Theorem 1, thus proving the equality (11). ■

Remark 3: It is shown in [4] that the optimal performance is bounded above by $\frac{|a^T|-1}{|a|-1} + \frac{2^R|a^T|}{2^R-|a|}$ when $a \neq 1$. The first term on right hand side is exactly equal to $\sup_{\|\mathbf{w}\|_\infty \leq 1} \|\mathbf{x}_d\|_\infty$. The second term is an upper bound of $\inf_{(\mathcal{K},(\mathbf{E},\mathbf{D}) \in \mathcal{G}_R)} \|\mathbf{x}_q\|_\infty = \frac{|a^T|}{2^R-|a|}$ because $2^R \leq 1$. This performance gap is due to the fact that in our setting, the encoder has access to $w(0:t)$.

Remark 4: The upper bound for quadratic distortion for a Gaussian process $x(t+1) = x(t) + \eta(t), \eta(t) \overset{i.i.d.}{\sim} N(0,1)$ by a channel with rate R is shown to be $(2^{2R} - 1)^{-1}$ [15]. Notice that our worse-case cost is $\inf \sup_{\|\mathbf{w}\|_\infty \leq 1} \|\mathbf{x}_q\| = (2^R - 1)^{-1}$ when $a = 1$. This suggests the impact of channel capacity on LQG cost scales in a similar manner.

VI. ACKNOWLEDGMENTS

Terry Sejnowski and Simon Laughlin taught us that spiking can be consistent but expensive. We were inspired to make connections by Mike Gazzaniga, Glenn Vinnicombe, and Nuno Martins. We also appreciate discussions with Sri Sarma, John Allman, Xaq Pitkow, Li Zhaoping, Markus Meister, Daniel Wolpert, Joel Burdick, Seungil You, Yoke Peng Leong, Scott Grafton, and Jorge Goncalves. We apologize to Sanjoy, Munther, Partha, and Girish for taking so long to get this.

REFERENCES

[1] L. Zhaoping and Z. Li, *Understanding vision: theory, models, and data*. Oxford University Press, 2014.

[2] P. Sterling and S. Laughlin, *Principles of Neural Design*. MIT Press, 2015.

[3] M. A. Dahleh and I. J. Diaz-Bobillo, *Control of Uncertain Systems: A Linear Programming Approach*. Upper Saddle River, NJ, USA: Prentice-Hall, Inc., 1995.

[4] G. N. Nair, F. Fagnani, S. Zampieri, and R. J. Evans, "Feedback control under data rate constraints: An overview," *Proceedings of the IEEE*, vol. 95, no. 1, pp. 108–137, 2007.

[5] N. Li, J. Cruz, C. S. Chien *et al.*, "Robust efficiency and actuator saturation explain healthy heart rate control and variability," *Proc. of the Natl. Acad. of Sciences*, vol. 111, no. 33, pp. E3476–E3485, 2014.

[6] F. A. Chandra, G. Buzi, and J. C. Doyle, "Glycolytic oscillations and limits on robust efficiency," *Science*, vol. 333, no. 6039, pp. 187–192, 2011.

[7] N. C. Martins and M. A. Dahleh, "Feedback control in the presence of noisy channels:"Bode-like" fundamental limitations of performance," *Auto. Control, IEEE Trans. on*, vol. 53, no. 7, pp. 1604–1615, 2008.

[8] D. W. Franklin and D. M. Wolpert, "Computational mechanisms of sensorimotor control," *Neuron*, vol. 72, no. 3, pp. 425–442, 2011.

[9] F. Rieke, D. Warland, R. d. R. van Steveninck, and W. Bialek, *Spikes: Exploring the Neural Code*. Cambridge, MA: MIT Press, 1999.

[10] M. Rotkowitz and S. Lall, "A characterization of convex problems in decentralized control," *Automatic Control, IEEE Transactions on*, vol. 51, no. 2, pp. 274–286, 2006.

[11] A. Mahajan, N. Martins, M. Rotkowitz, and S. Yuksel, "Information structures in optimal decentralized control," in *Decision and Control (CDC), 2012 IEEE 51st Annual Conference on*, Dec 2012, pp. 1291–1306.

[12] Y.-S. Wang, N. Matni, and J. Doyle, "Localized LQR optimal control," in *Decision and Control (CDC), 2014 IEEE 53rd Annual Conference on*, Dec 2014, pp. 1661–1668.

[13] N. Matni, "Communication delay co-design in $\mathcal{H}_2$ distributed control using atomic norm minimization," *IEEE Transactions on Control Of Network Systems. Conditionally Accepted.*, 2015.

[14] N. Matni and V. Chandrasekaran, "Regularization for design," *Submitted to the IEEE Transactions on Automatic Control*, 2015.

[15] M. S. Derpich and J. Østergaard, "Improved upper bounds to the causal quadratic rate-distortion function for gaussian stationary sources," *Info. Theory, IEEE Trans. on*, vol. 58, no. 5, pp. 3131–3152, 2012.