LQ vs. ℓ_∞ in controller design for systems with delay and quantization

Yorie Nakahira

Abstract— **The normal operation of many cyberphysical, biological, and neural systems fit naturally with robust control, with key variables like lane positions, voltages, temperatures, blood pressures, etc maintained within tight bounds despite diverse uncertainties. However, two challenges in particularly need further theory that this paper addresses. One is that control is distributed with communication having limits on bandwidth and delay. Another is that normal operations can be disrupted and bounds violated, but it is desirable to make such acute situations rare and recoverable without crashing. We take the simplest model that has both normal and acute modes with bandwidth and delay constraints, and focus on two relatively extreme but familiar starting points: i) average case LQG (or $\mathcal{H}_2$) and ii) worst case ℓ_1 control with ℓ_∞ signal bounds. Both have strengths and weaknesses that we highlight, and this leads naturally to a win-win hybrid scheme that has better performance than either alone, with relatively modest computational costs.**

I. INTRODUCTION

Tradeoffs at both system and component levels are ubiquitous in human sensorimotor system, arguably the ultimate cyberphysical system. At the component level, limitations in nervous resources impose stringent speed/accuracy tradeoffs in nerve signaling [1]. At the system level, limitations in hardware imposes tradeoffs between robust performance and controller complexity [2]. The tradeoffs at the system and component level have been studied in isolation in prior works, but there is no theoretical framework to analyze the interactions between the system and component tradeoff, which is crucial to the understanding of the human sensorimotor system. In this paper, we develop control theory with communication constraints, relevant to neuroscience and cyberphysical systems. This theory provides insights into the interaction between the tradeoffs at component and system level. This paper mainly focuses on the control theory, while its application in neuroscience is discussed in [3] [4].

We initially focus on optimal control for simple systems with communication constraints. The LQ (linear-quadratic) optimal control and ℓ_1 optimal control have been studied extensively for systems with perfect communications [5][6][7]. We focus on LQ and ℓ_∞ because they are extremes, different in both average and worst case, and different in underlying norm (ℓ_2 and ℓ_∞). We'll use ℓ_1 and ℓ_∞ interchangeably depending on whether we're emphasizing system or signal norms.

For systems with communication constraints, the stability conditions have been studied in [8][9], and optimal controller

structures have been studied in [10][11] for LQ control and in [12][4] for ℓ_∞ control. Although the work [4] has demonstrated the system level tradeoff between delay cost and quantization cost in the ℓ_∞ system, it was unknown whether that tradeoff holds universally for other systems (or norms). In this paper, we derive a performance lower bound as a function of delay and bandwidth in a LQ system, from which the system level tradeoff between delay cost and quantization cost can also be observed.

The optimal LQ controller for delayed and quantized system satisfies certainty equivalence [10], and is more demanding than the ℓ_∞ system in both computation and memory requirements. This great increment in controller complexity was hinted in [12]: any static memoryless coder cannot reject a disturbance with infinite support. Although a disturbance that takes large value with small probability is a good abstraction of an acute mode disturbance, considering optimality (minimizing LQ cost) leads to designing a controller with heavy computation and large memory usage, which we will ultimately show is avoidable with a hybrid scheme.

Combining the ℓ_∞ and LQ controllers, we propose a hybrid controller that maintains the advantage of both systems. This hybrid controller requires minimal computation and memory most of the time during normal mode, but provides high disturbance attenuation in a rare acute mode, quickly returning the system to normal. Such controllers – having fine normal resolution, but capable of adapting acutely to a wider range of signals –are also ubiquitous in biological systems [13]. The hybrid controller has the tradeoff between system performance and controller complexity, which are also discuss in this paper.

The contribution of this paper is summarized as follows:

The speed/accuracy tradeoff. We provide analytic formulas characterizing the impact of signaling delay and bandwidth on system performance. Both the LQ and ℓ_∞ systems have speed/accuracy tradeoffs that can be optimized in the design of communication hardware.

The performance/complexity tradeoff. We propose a hybrid controller that operates with low complexity most of the time in normal mode, yet is able to reject a disturbance with unbounded support in the rare acute mode, returning quickly to normal mode. The design space of the proposed controller has performance/complexity tradeoffs, and a sweet spot where both good system performance and low controller complexity can be achieved.

Robustness under mixed disturbances. When a band-limited controller needs to stabilize an unstable system that contains both worst case (bounded) and stochastic distur-

bances, the optimal ℓ_∞ controller cannot guarantee the stability, and the LQ controller suffers from a degraded performance. However, the additional design freedom of the proposed hybrid controller helps achieve a consistently robust performance with a right blend of LQ and ℓ_∞ framework.

The rest of this paper is organized as follows. We describe the control system model in Section II. The LQ and ℓ_∞ systems, their performance, and controller properties are studied in Section III. The hybrid controller is proposed and analyzed in Section IV. The discussion about speed/accuracy tradeoffs, performance/complexity tradeoffs, and robustness under mixed disturbances are given in Section V. The proofs of theorems are provided in the technical report [14], and are omitted due to space constraints.

Notation and preliminaries:: We use lower case letters to denote sequences, i.e., $x = \{x_0, x_1, x_2, \dots\}$, and x_τ^t to denote the truncated sequence from τ to t, i.e., $x_\tau^t = \{x_\tau, x_{\tau+1}, \cdots, x_t\}$. The ℓ_∞ norm of a sequence x is defined as $\|x\|_\infty := \sup_t |x(t)|$.

We denote a uniform quantizer of rate R (i.e., with 2^R levels) over the interval $[-1, 1]$ as $Q : \mathbb{R} \to S_R$, where $|S_R| = 2^R$. The output of the map $Q(x)$ is given by the nearest point to x in S_R, i.e., the nearest of the 2^R uniformly distributed points across the interval $[-1, 1]$.

II. Model Description

In this section, we briefly introduce a dynamical system equipped with a feedback controller with communication constraints in Fig. 1 similarly to [4]. The plant satisfies the following dynamics:

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

where $x_t \in \mathbb{R}$ is the state, $w_t \in \mathbb{R}$ is the disturbance, and $u_t \in \mathbb{R}$ is the control action. Let the initial condition be $x_0 = 0$ and $w_t = 0$ for $t < 0$. The control action u_t is generated by a controller K with delay d and bandwidth R, where R is minimum stabilizing $R > \log_2 |A|$ [8]. The controller $K := \{(E_0, D_0), (E_1, D_1), (E_2, D_2) \cdots\}$ is parameterized by the encoder E_t and decoder D_t at each time t. We assume the encoder can utilize the information of $\mathcal{I}_t \subset \{\{x_\tau\}_{\tau=0,\dots,t}, \{w_\tau\}_{\tau=0,\dots,t-1}\}$, is defined by the mapping $E_t : \mathcal{I}_t \to S$ such that

$$s_t = E_t(\{x_\tau\}_{\tau=0,\dots,t}, \{w_\tau\}_{\tau=0,\dots,t-1}) \in S, \tag{2}$$

where $s_t \in S$ is the codeword, and S has cardinality at most 2^R. The decoder can utilize the information $\mathcal{J}_t \in \{\{s_\tau\}_{\tau=0,\dots,t}\}$ with delay d caused by the channel, is defined by the mapping $D_t : \mathcal{J}_{t-d} \to \mathbb{R}$ such that

$$u_t = D_t(\mathcal{J}_{t-d}). \tag{3}$$

We denote $\mathcal{K}(R, d)$ to be the space of controllers with delay d and bandwidth R that has the form above. In addition, we assume there is a tradeoff between building a feedback loop with smaller delay d and greater bandwidth R, *i.e.*,

$$R = \mathcal{T}(d), \tag{4}$$

where $\mathcal{T}$ is an increasing function.

We assume that the plant dynamics (1) and physical limitations $\mathcal{T}$ are given, whereas K, d, R and $(\mathcal{I}_t, \mathcal{J}_t)$ are design variables. The desired properties for the controller is to achieve a good performance (small state deviation) with small control effort (low actuation, computation, and memory usage). We quantify the performance in $\|x\|$ for some norm, low actuation in $\|u\|$, computation using algorithm processes, and memory usage using the information sharing structure $(\mathcal{I}_t, \mathcal{J}_t)$. With a sa abuse of notation, we call the controller processes K as 'software', and the component properties $(d, R, \mathcal{I}_t, \mathcal{J}_t)$ as 'hardware'. In the following sections, we explore controllers with different software K and hardware $(d, R, \mathcal{I}_t, \mathcal{J}_t)$ and their performance .
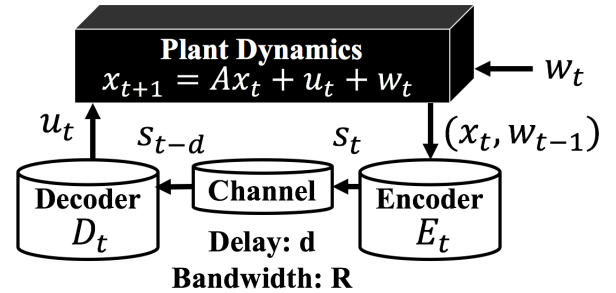

Fig. 1: The system model. A plant is connected to a controller composed of an encoder, channel, and decoder. The channel has delay d and bandwidth R.

Remark 1 (The speed/accuracy tradeoffs): The model described above can be used to model the human sensorimotor and potentially cyberphysical systems. In the sensorimotor system, the body corresponds to the plant, the central nervous system to the controller having delay and limited bandwidth. The speed/accuracy tradeoff (4) for spiking neurons is given by $R \propto d$ in [4], [1] if we assume the space and energy for those neurons are fixed.

III. Two Control Problems

In this section, we study the optimal performance for the LQ/ℓ_∞ system, and the controllers that achieve the optimal performance. We also assume an ideal scenario when the controller has unlimited computational and storage capabilities, and leave the issue of controller complexity to Section IV.

A. The Linear-Quadratic System

In this section, we study the robust control problem for the LQ system with delay and quantization. The disturbance w_t for $t \geq 0$ is assumed to be *i.i.d.* Gaussian noise with zero mean and variance σ^2, i.e. $w_t \overset{i.i.d.}{\sim} \mathcal{N}(0, \sigma^2)$ for $t \geq 0$. Our goal is to solve the robust control problem:

$$\begin{aligned} \underset{K \in \mathcal{K}(R,d)}{\text{minimize}} \quad & \lim_{t \to \infty} \mathbb{E}[x_t' P x_t + u_t' Q u_t] \\ \text{subject to} \quad & \text{plant dynamics (1)} \end{aligned} \tag{5}$$

for some known matrices $P \succeq 0$, $Q \succ 0$. Then, the lower-bound of the LQ cost performance is stated in Theorem 1.

Theorem 1: The optimal performance of the robust control problem (5) is bounded below by

$$\lim_{N \to \infty} \mathbb{E}[x_t' P x_t + u_t' Q u_t] \qquad (6)$$

$$\geq P \sum_{i=0}^{d-1} A^{2i} \sigma^2 + P^\star A^{2d} \sigma^2 + G^\star A^{2d} \frac{\sigma^2}{2^{2R} - A^2}$$

where the matrices $P^\star$, $G^\star$ are the unique solutions of

$$\begin{aligned} P^\star &= A' \left[P^\star + P - P^\star (Q + P^\star)^{-1} P^\star \right] A \\ G^\star &= A' P^\star (Q + P^\star)^{-1} P^\star A. \end{aligned} \qquad (7)$$

Due to space constraints, we omit its proof from this paper and refer to the technical report [14]. The performance limitation in Theorem 1 has a clear interpretation. The first and second terms in the right side of (6), $P \sum_{i=0}^{d-1} A^{2i} \sigma^2 + P^\star A^{2d} \sigma^2$, are due to delayed control action, and the third term $G^\star A^{2d} \frac{\sigma^2}{2^{2R} - A^2}$ is due to limited bandwidth.

The optimal controller for the following finite-horizon optimal control problem is also stated in Lemma 1.

$$\begin{aligned} &\underset{K \in \mathcal{K}(R,d)}{\text{minimize}} \quad \lim_{t \to \infty} \mathbb{E} \left[x_N' P x_N + \sum_{t=0}^{N-1} (x_t' P x_t + u_t' Q u_t) \right] \\ &\text{subject to} \quad \text{plant dynamics (1).} \end{aligned}$$
$$(8)$$

Lemma 1: Let us consider the finite-horizon optimal control problem (8). Then, the optimal controller for (8) has the structure

$$u_t = L_t \, \mathbb{E}[z_t | s^{t-d}], \qquad (9)$$

where the sequence z_t is defined from the recursion

$$z_{t+1} = A z_t + A^d w_{t-d} + u_t, \quad z_0 = 0 \qquad (10)$$

and the matrices L_t, P_t are defined from the recursion

$$\begin{aligned} P_N &:= P \\ P_t &:= A' \left[P_{t+1} + P - P_{t+1}(Q + P_{t+1})^{-1} P_{t+1} \right] A \\ L_t &:= -(Q + P_{t+1})^{-1} P_{t+1} A. \end{aligned}$$

Remark 2 (Certainty equivalence): The definition of certainty equivalence and its extension to quantized systems is given in [10] [11]. The optimal controller structure in Lemma 1 is an extension of certainty equivalence for systems with delay and quantized, and the sequence $\{z_t\}$ plays an important role in Section IV.

Remark 3 (MIMO extension): For a MIMO plant in the state space representation $(A, B, C, 0)$, the optimal controller structure (9) can be re-defined using

$$\begin{aligned} z_{t+1} &= A z_t + A^d w_{t-d} + B u_t \\ P_N &:= P \\ P_t &:= A' \left[P_{t+1} + P - P_{t+1} B (Q + B' P_{t+1} B)^{-1} B' P_{t+1} \right] A \\ L_t &:= -(Q + B' P_{t+1} B)^{-1} B' P_{t+1} A. \end{aligned}$$

The achievable performance. Based on the optimal controller structure from Lemma 1, we propose a controller in Algorithm 1, and denote it as *the LQ controller*. The encoder and decoder in Algorithm 1 essentially performs a Bayes filter for a hidden Markov chain. The operation (12) does the prediction, while the operation (13) does the update. The Lloyd algorithm [15] is used to design an adaptive quantizer based on the prior knowledge of z_t.

Initialize:
1) Set $f(z_d | s^0) = \mathcal{N}(0, \sigma^2)$.
2) Set $z_d = 0, u_0 = 0$.

Encoder: At time t, the encoder performs the following procedures:
1) Update the auxiliary variable

$$z_t = A z_{t-1} + A^d w_{t-d-1} + u_{t-1}. \qquad (11)$$

2) Generate the decoder's prior distribution by

$$f(z_t | s_{t-d-1}) = \int_{-\infty}^{\infty} f(z_t | z_{t-1}) f(z_{t-1} | s_{t-d-1}) dz_t \qquad (12)$$

where the value of $f(z_t | z_{t-1})$ is computed using

$$f(z_t | z_{t-1}) = f(A z_{t-1} + A^d w_{t-d-1} + u_{t-1} | z_{t-1}).$$

3) Run the Lloyd algorithm [16] to obtain a quantization scheme Q_t.
4) Send the codeword $s_{t-d} = Q_t(z_t)$ to the decoder.
5) Generate the decoder's posterior distribution by

$$f(z_t | s_{t-d}) \propto f(s_{t-d} | z_t) f(z_t | s_{t-d-1}) \qquad (13)$$

with appropriate scaling.

Decoder: At time t, the decoder receives the codeword s_{t-d} that was generated d sampling intervals before, and performs the following procedures:
1) Compute the prior information $f(z_t | s_{t-d-1})$ using relation (12).
2) Run the Lloyd algorithm to recover encoder's quantizer Q_t.
3) Use the delayed codeword s_{t-d} to generate the posterior distribution by (13).
4) Calculate the estimate of z_t by

$$\hat{z}_t = \mathbb{E}[z_t | s_{t-d}] = \int_{-\infty}^{\infty} z_t f(z_t | s_{t-d}) dz_t.$$

5) Produce the control action:

$$u_t = -(Q + P^*)^{-1} P^* A \, \hat{z}_t. \qquad (14)$$

Algorithm 1: The LQ controller

The gap between the achievable performance and the fundamental limitation is compared in Fig. 2. It suggests that Algorithm 1 starts to produce near optimal performance when the bandwidth is approximately greater than 5. The derivation of Theorem 1 suggests that the first two terms in (8), $P \sum_{i=0}^{d-1} A^{2i} \sigma^2 + P^\star A^{2d} \sigma^2$, is tight for any delay $d \in \mathbb{N}$ if the controller has the structure (9) (which is the case for Algorithm 1). Then, the performance gap reduces to the difference between the empirical value of $(z_t - \hat{z}_t) G^*(z_t - \hat{z}_t)$ and the lower bound of $\mathbb{E}[(z_t - \hat{z}_t) G^*(z_t - \hat{z}_t)]$.
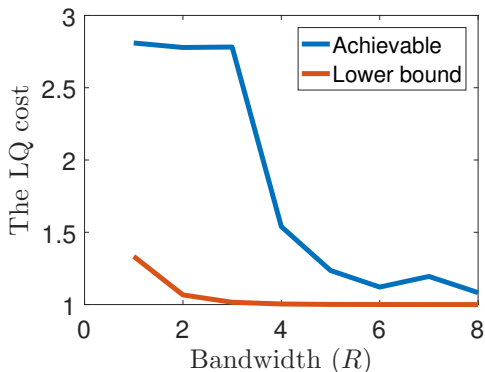
Fig. 2: Tightness of (6) in Theorem 1. The theoretical lower bound in (6) and the achievable performance by Algorithm 1 are compared for $A = 1$, $d = 0$, and $\sigma^2 = 1$.

Encoder:
$$q_t = Q_{\Psi(L)}^{-1}(s_{t-d-1}) - u_{t-1}^*$$
$$z_t = A^d w_{t-d-1} + q_t$$
$$u_t^* = -A z_t$$
$$s_{t-d} = Q_{\Psi(L)}(u_t^*)$$
Decoder:
$$u_t = Q_{\Psi(L)}^{-1}(s_{t-d})$$

Algorithm 2: The ℓ_∞ controller

The LQ controller is demanding in both computation and memory usage. This poses the question whether such a heavy procedure is indeed necessary. Building an adaptive quantizer is necessary for stabilizing an unstable system if the disturbance has an infinite support [17]. This is because, for any fixed quantizer, there is always a non-zero probability to observe a disturbance that exceeds the limit of this quantizer. However, it is not necessary to *always* use an adaptive quantizer nor Algorithm 1. We present a much simpler controller (the solution of optimal ℓ_1/ℓ_∞ control) in Section III-B, and combine the two controllers to reduce controller complexity in Section IV.

B. The ℓ_∞ System

We summarize here the existing robust control theory for ℓ_∞ systems with delay and quantization [12][4]. For disturbance with bounded support $\|w\|_\infty \leq L$ and stabilizing bandwidth $R > \log_2 |A|$, the optimal performance is

$$\max_{\|w\|_\infty \leq L} \|x\|_\infty = \left\{ \sum_{i=0}^{d} |A^i| + \frac{|A^{d+1}|}{(2^R - |A|)^{-1}} \right\} L. \quad (15)$$

Let $\Psi(L) := \left\{ |A^{d+2}|(2^R - |A|)^{-1} + |A^{d+1}| \right\} L$. The optimal performance is achieved by the controller shown in Algorithm 2. We denote this as the ℓ_∞ *controller*.

The advantage of this controller is that it requires little computation and storage: the encoder only needs to store the last codeword and perform minimum computation; the decoder is static and memoryless. In addition, this controller also requires minimum actuation effort when $|A| \geq 1$: the

stabilizing control law that minimizes $\max_{\|w\|_\infty \leq 1} \|u\|_\infty$ is identical to the above control law, which minimizes $\max_{\|w\|_\infty \leq 1} \|x\|_\infty$. However, the low complexity of the ℓ_∞ controller does not come for free. For a disturbance with unbounded support, the fixed quantizer in Algorithm 2 is not stabilizing because there is always a nonzero probability that the quantizer saturates. In next section, we combine the LQ controller with the ℓ_∞ to overcome the limitations of both controllers.

IV. MIXED ROBUST CONTROL PROBLEMS

We have observed in the previous section that the ℓ_∞ and LQ controller have contrasting properties: the ℓ_∞ controller is simple but incapable of rejecting large disturbance, whereas the LQ controller is capable of rejecting large disturbance but heavy in both computation and memory usage. This motivates us to study the middle ground between the two controller design frameworks, aiming at designing a controller that preserves the advantages of both frameworks. In this section, we propose a controller that has low complexity and high disturbance rejection capability.

A. The Hybrid Controller

The controller that mostly runs Algorithm 2 (ℓ_∞ controller) and occasionally Algorithm 1 (LQ controller) has low complexity (in computation and memory) and is capable of rejecting occasionally large disturbances. We use this observation to propose a hybrid controller.

From now on, we assume the LQ cost function has no control cost, *i.e.*, $Q \to 0$ in (5), yielding the optimal LQ controller

$$u_t = -A\hat{z}_t, \quad (16)$$

to replace (9) in Algorithm 1. This simplification allows the sequences z_t in both Algorithms 1 and 2 to take the same value, which lets the ℓ_∞ and LQ controller to be considered in an unified framework.

We allow the proposed controller to have two modes: *normal-mode* that runs the ℓ_∞ controller (Algorithm 2) and *acute-mode* that runs the LQ controller (Algorithm 1). We first explain the switching policies between the ℓ_∞ and LQ controller using a bridging variable z_t and a design parameter L. Observe that the sequences z_t in the ℓ_∞ and LQ controller have identical roles (storing the sum of the quantization error from past control action and the scaled disturbance $A^d w_{t-d-1}$), and thus can serve as a bridging variable to connect the two controllers. We first re-define the sequence q_t by

$$q_{t+1} = Aq_t + u_t + A^{d+1} w_{t-d-1}, \quad (17)$$

where $w_t = 0$ for $t < 0$. The definition (17) does not rely on a particular realization of the controller, and thus q_t are well-defined in both Algorithms 1 and 2. Using q_t, we define the sequence z_t by

$$z_t := A^d w_{t-d-1} + q_t, \quad (18)$$

Initialize: $mode \leftarrow \text{`}normal\text{'}$
$\Psi(L) \leftarrow \{|A^{d+2}|(2^R - |A|)^{-1} + |A^{d+1}|\}L$
for $t \in \mathbb{N}$ **do**
    if *mode = `normal'* **then**
        Perform the ℓ_∞ controller (Algorithm 2)
        if $|z_t| > \Psi(L)/A$ **then**
            $mode \leftarrow \text{`}acute\text{'}$
        end
    else
        Perform the LQ controller (Algorithm 1)
        if $|z_t| \leq \Psi(L)/A$ **then**
            $mode \leftarrow \text{`}normal\text{'}$
        end
    end
end

Algorithm 3: The hybrid controller

with the initial condition $z_t = 0$ for $t \leq d$. z_t in Algorithm 2 can be rewritten as

$$z_{t+1} = Az_t + A^d w_{t-d} + u_t \qquad (19)$$

$$= A^d w_{t-d} + Aq_t + u_t + A^{d+1} w_{t-d-1} \qquad (20)$$

$$= A^d w_{t-d} + q_{t+1}, \qquad (21)$$

The equality (19) holds by definition (11), the equality (20) by definition (18), and the equality (21) by definition (17). Therefore, z_t takes the same value in both Algorithms 1 and 2. The proposed controller sets a threshold on the absolute value of z_t to determine whether the ℓ_∞ or LQ controller should be used.

Let the design parameter $L \in \mathbb{R}$ be the size of the disturbance up to which the controller stays in normal-mode, *i.e.*, $\|w\|_\infty \leq L$ implies normal-mode. Since $\|w_0^{t-d-1}\|_\infty \leq L$ implies $|z_t| \leq \Psi(L)/A$, equivalently $|z_t| > \Psi(L)/A$ implies $|w_\tau| \geq L$ for some $\tau \leq t-d-1$. Thus, the condition

$$|z_t| > \Psi(L)/A \qquad (22)$$

is a sufficient condition for $\|w_0^{t-d-1}\|_\infty > L$. We use this sufficient condition to define the switching policy as follows:

$$mode = \begin{cases} \text{`}normal\text{'} & |z_t| \leq \Psi(L)/A \\ \text{`}acute\text{'} & |z_t| > \Psi(L)/A. \end{cases} \qquad (23)$$

Now we are ready to present the proposed controller in Algorithm 3, which is denoted as the *Hybrid* controller.

The design parameter L impacts the system performance and controller complexity, and there exists a tradeoff between the two. We discuss its design guideline and this performance/complexity tradeoff in subsequent sections.

B. Switching Behavior of the Hybrid Controller

In this section, we analyze the switching behavior of the hybrid controller using the switching time from normal to acute mode and the recovery time from acute to normal mode. We denote the set of times at which the controller switches from normal to acute mode as

$$\mathcal{T}_s = \{t \in \mathbb{N} : |z_t| > \Psi(L)/A \text{ and } |z_{t-1}| \leq \Psi(L)/A\},$$

and the set of time at which the controller switches from acute to normal mode as

$$\mathcal{T}_r = \{t \in \mathbb{N} : |z_t| \leq \Psi(L)/A \text{ and } |z_{t-1}| > \Psi(L)/A\}$$

Let $t_r \in \{0\} \cup \mathcal{T}_r$ be the beginning of a normal mode, then the the switching time T_s can be defined by

$$T_s(t_r) = \min\{t > t_r : |z_t| > \Psi(L)/A\} - t_r. \qquad (24)$$

Let $t_s \in \mathcal{T}_s$ be the beginning of an acute mode, then the recovery time T_r can be similarly defined by

$$T_r(t_s) = \min\{t > t_s : |z_t| \leq \Psi(L)/A\} - t_s. \qquad (25)$$

Long switching time and short recovery time imply the controller stays in normal-mode most of the time, and thus requires less computation and memory. Therefore, the controller complexity can roughly be characterized by the frequency of operating in acute-mode.

Let a random variable w be drawn from the same distribution with the disturbance w_t, *i.e.*, $w, w_t \overset{i.i.d.}{\sim} \mathcal{N}(0, \sigma)$. The following theorem characterizes the relation between the design parameter L and the expected switching time $\mathbb{E}[T_s(t_r)]$.

Theorem 2: Let a map $\hat{T}_s : \mathbb{R} \rightarrow \mathbb{R}_+$ be defined by

$$\hat{T}_s(t_r) = \begin{cases} d + \mathbb{P}(|w| > L)^{-1} & t_r = 0 \\ \mathbb{P}(|w| > L)^{-1} & t_r \in \mathcal{T}_r. \end{cases}$$

The expected switching time $T_s(t_r)$ is lower bounded by

$$\mathbb{E}[T_s(t_r)] \geq \hat{T}_s(t_r), \qquad (26)$$

where the lower bound becomes tight as $R \rightarrow \infty$.
The proof of Theorem 2 uses the concept of *majorization* to approximate the switching time by *geometric distribution*. Due to space constraints, we omit the proof from this paper and refer to the technical report [14]. Theorem 2 suggests that the expected switching time can be approximated by $\mathbb{E}[T_s(t_r)] \approx \hat{T}_s(t_r)$.

Analogously, the expected recovery time $T_r(\cdot)$ can be approximated by

$$\mathbb{E}[T_r(\cdot)] \approx \mathbb{P}(|w| \leq L)^{-1}. \qquad (27)$$

We denote this value as $\hat{T}_r := \mathbb{P}(|w| \leq L)^{-1}$. Recall from (21) that the evolution of z_t follows $z_{t+1} = A^d w_{t-d} + q_{t+1}$ where q_{t+1} is a function of z_t. Assuming the quantizer (defined from the encoder and decoder) is near-optimal, a large z_{t_s} at the beginning of the an acute mode is approximately reduced by rate $|A|2^{-R}$ per unit time, by rate $|A^\tau|2^{-\tau R}$ during τ times. Thus, for sufficiently large $|A|2^{-R}$, the term $A^d w_{t-d}$ in (21) dominates. In this situation, observing a small disturbance, *i.e.*, $|w_{t-d}| \leq L$, is enough to lessen the value of z_t below $\Psi(L)/A$. This explains why the recovery time can be approximated by a geometric distribution with success probability $\mathbb{P}(|w_t| \leq L)$.

Fig. 4 compares the empirical value of the expected switching time $T_s(0)$ with the theoretical approximation $\hat{T}_s(0)$, and the empirical value of the expected recovery time $T_r(\cdot)$ with the theoretical approximation $\hat{T}_r$. It suggests that
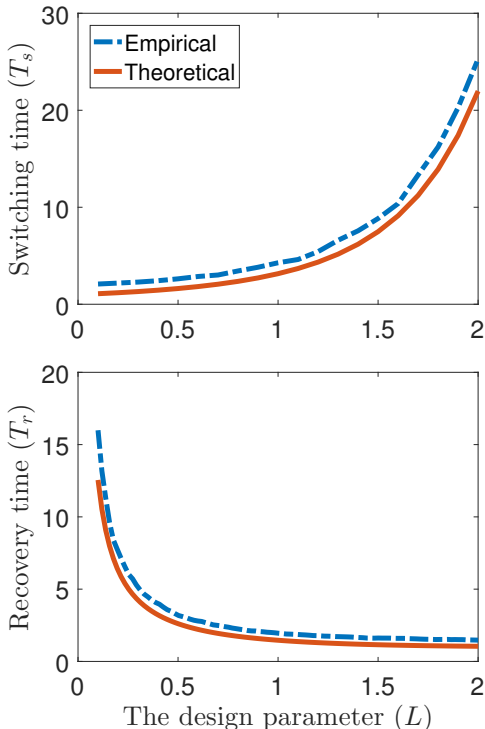
Fig. 3: The switching or recovery time as a function of L. The figure on the top compares the empirical switching time $T_s(t_r)$ and its theoretical lower bound $\hat{T}_s(t_r)$ from Theorem 2. Since the tightness of (26) does not depend on the beginning of a normal mode $t_r \in \{0\} \cup \mathcal{T}_r$, we set $t_r = 0$ without loss of generality. The figure on the bottom compares the empirical recovery time $T_r(\cdot)$ and its approximation $\hat{T}_r(\cdot)$ in (27). We set $A = 1$, $d = 1$ and $R = 6$, and averaged 100 trials when generating the empirical values.

Fig. 4: The accuracy of the theoretical approximation in (26) and (27). The figure on the top shows the convergence speed of $\mathbb{E}[T_s(t_r)] \overset{R \to \infty}{\to} \hat{T}_s(t_r)$ in Theorem 2 (we set $t_r = 0$ without loss of generality). The figure on the bottom shows the error in approximating the expected recovery time (27). We set $A = 1$, $d = 1$. The empirical values of T_s, T_r are first generated by averaging 100 trials for different values of $L \in [0.1, 2]$ and $R \in \{1, 2, \cdots, 9\}$. Then, the approximation errors $|T_s - \hat{T}_s|$, $|T_r - \hat{T}_r|$ are averaged over all L, and their mean values are plotted for different values of R.

the approximation becomes tight from $R \geq 5$ for expected switching time, and from $R \geq 5$ for expected recovery time.

The theoretical approximation suggests that, for sufficiently large bandwidth ($|A|2^{-R} \ll 1$), a greater L implies larger switching time (from $T_s(t_r) \approx \hat{T}_s(t_r) = \mathbb{P}(|w_t| > L)^{-1}$) and smaller recovery time (from $T_r(t_s) \approx \hat{T}_r(t_s) = \mathbb{P}(|w_t| \leq L)^{-1}$). This can also be empirically verified from Fig. 3. Since the switching time is an increasing and the recovery time a decreasing function of L, the complexity of the hybrid controller decreases as L increases. It should also be noted that decrements in controller complexity come with cost of degraded performance because large L also implies a coarser quantizer in Algorithm 2 (and thus larger quantization error). This tradeoff between performance and controller complexity is discussed Section V-B.

V. DISCUSSION

In this section, we address two challenges: A) the optimal system performance under communication constraints, B) system behaviors when normal operations are disrupted. In particular, we discuss in Section V-A how the component level speed/accuracy 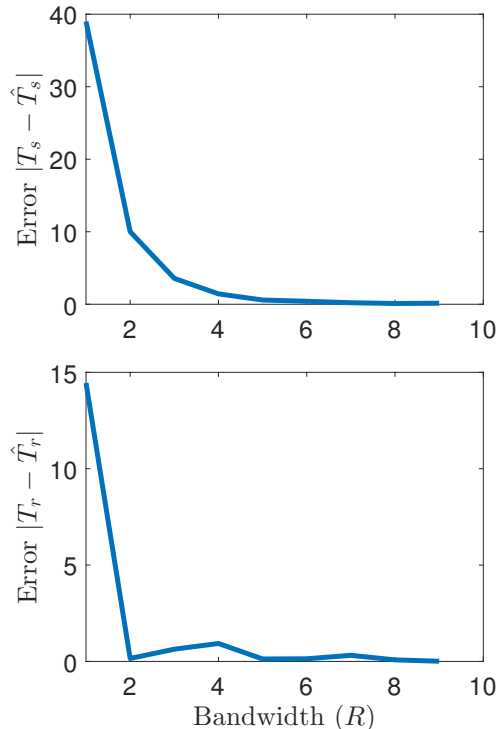tradeoff relates to the system level robust performance, in Section V-B why it is advantageous to use a hybrid controller under computation and memory constraints, and in Section V-C the system robustness under mixtures of worst-case and stochastic disturbances.

A. The Speed vs. Accuracy Tradeoffs

Recall from Remark 1 that the nervous system has a tradeoff between spatial and temporal resolution, $R = \lambda d$. Combining this with the LQ cost (8), we can observe a tradeoff between signaling speed and accuracy on system performance, analogously to the one in the ℓ_∞ system [4]. Larger bandwidth reduces the quantization cost but increases the delay cost. Consequently, the optimal performance is achieved near the minimum stabilizing bandwidth in order to minimize the cost due to delay. Fig. 5 also compares the delay and quantization cost of the LQ system with the ℓ_∞, suggesting that the impact of delay relative to bandwidth is greater in the LQ system than in the ℓ_∞ system.

B. The Performance vs. Complexity Tradeoffs

In this section, we discuss the constraints on controller complexity and its impact on system performance. Assuming
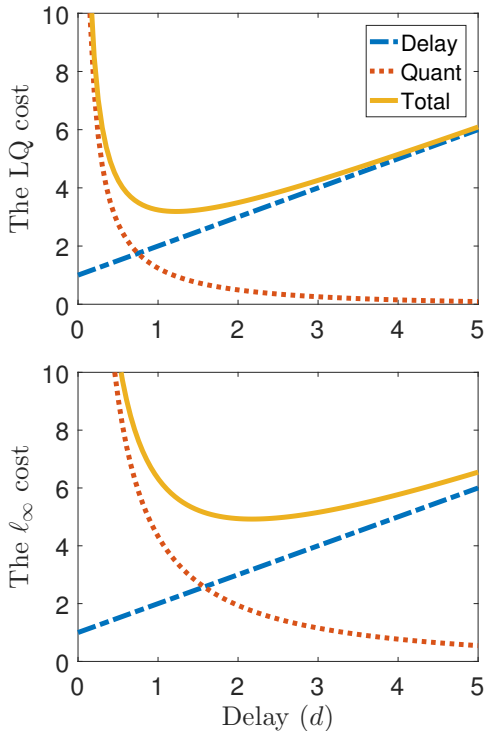
Fig. 5: The system level tradeoffs between delay cost and quantization cost. The figure on the top shows the performance upper bound for the LQ system from Section III-A when $Q \to 0$. Its LQ cost is $d+1+(2^{2R}-A^2)^{-1}$ when $A=1$ and $R = 0.3 \ d$, where the term $d+1$ can be interpreted as the delay cost and $(2^{2R} - A^2)^{-1}$ as the quantization cost. The figure on the bottom shows the optimal performance for the ℓ_∞ system together with its delay and bandwidth cost (see Section III-B and [4] for more details).

Fig. 6: The tradeoffs between switching/recovery time and normal mode ℓ_∞ cost. The figure on the top compares the normal mode performance (28) (in ℓ_∞ cost) and the lower bound of the expected switching time (26). The figure on the bottom compares (28) and the approximation of the expected recovery time $\hat{T}_r$ in (27). We set $A=1$ $d = \{1, 2, 3\}$, where the delays are assumed to satisfy a speed/accuracy tradeoff $R = d$.

the controller is designed to mostly stay in normal mode, we quantify the performance using the normal mode state deviation in ℓ_∞ norm. Let t be the time when the controller is in normal mode. The upper bound of normal mode state derivation is a linear function of L as follows:

$$|x_t| \leq \left(\sum_{i=0}^{d} |A^i| + |A^{d+1}|(2^R - |A|)^{-1} \right) L. \quad (28)$$

This inequality is implied by (23) (the condition for being in normal mode), and is a trivial extension of [4].

Fig. 6 shows switching time (26), recovery time (27) as a function of L. Allowing a small degradation on performance leads to significant increments in switching time as well as significant decrements in recovery time (notice that the axes for switching and recovery time is in log-scale). Recall from (28) that small L implies better normal mode performance, as well as shorter switching time and longer recovery time. This suggests the existence of performance/complexity tradeoffs discussed below.

Fig. 6 compares system performance (28) and the acute mode ratio $T_r/(T_s + T_r)$ (approximated by $\hat{T}_r/(\hat{T}_s + \hat{T}_r)$), where small acute mode ratio implies low complexity in
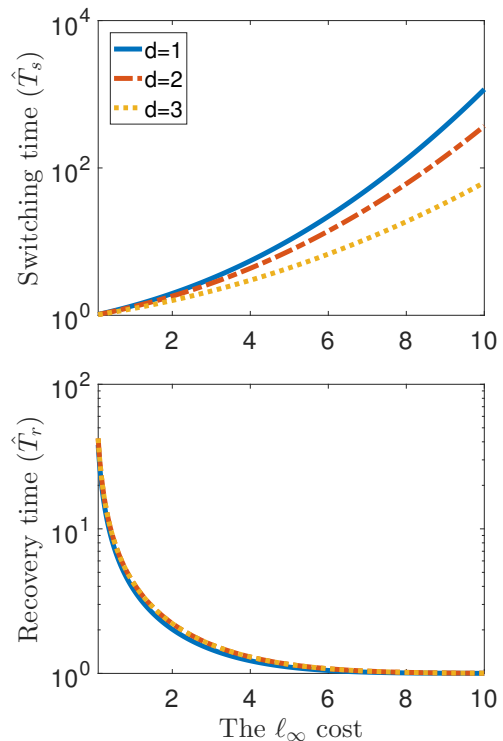
controller. Notice that there exists a sweet spot for the choice of L at $\|x\|_\infty \approx 6$, where the controller operates in normal mode most of the time, i.e., $\hat{T}_r/(\hat{T}_s + \hat{T}_r) \ll 1$.

It also suggests that the systems with different delays have similar performance/complexity tradeoffs, and thus the optimal hardware constrained by the speed/accuracy tradeoff and the optimal software constrained by the performance/complexity tradeoff have little dependence on each other. This simplifies the design process as we can first choose the optimal delay and bandwidth that maximizes the first term in (28), $\sum_{i=0}^{d} |A^i| + |A^{d+1}|(2^R - |A|)^{-1}$, and then design the variable L that so that the controller rarely stay in acute mode.

C. Robustness under Mixed Disturbances

In this section, we discuss the behavior of the proposed controllers (Algorithm 1-3) under mixed disturbances. Consider system (1) with a mixture of worst case and stochastic disturbances, i.e.,

$$w_t = v_t + r_t \quad (29)$$

where $v_t \overset{i.i.d.}{\sim} \mathcal{N}(0, \sigma_v^2)$ and $\|r\|_\infty \leq 1$. We additionally assume that the system is unstable, i.e., $|A| \geq 1$. These
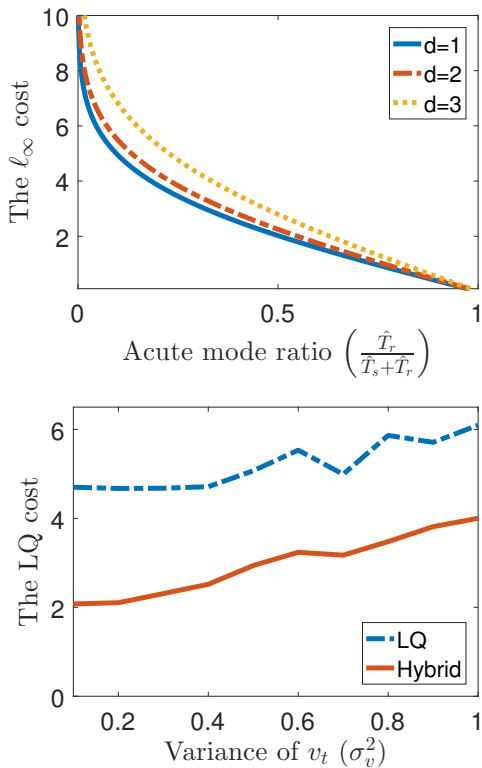
Fig. 7: Behaviors of the hybrid controller. The figure on the top shows the tradeoff between the normal mode performance (in ℓ_∞ cost) and the ratio of normal mode and acute mode $T_s/(T_r + T_s)$. The system parameters are set to be $A = 1$, $d = \{1, 2, 3\}$, and $R = d$. The figure on the bottom shows the the performance (in LQ cost) when the disturbance contains two types of signals, $w_t \overset{i.i.d.}{\sim} \mathcal{N}(0, \sigma^2)$ and $\|v\|_\infty \leq 1$. We set $A = 1$, $d = 1$ and $R = 3$. When running Algorithm 1 in the LQ controller and hybrid controller, we first searched for an optimal variance (the parameter σ^2 in Algorithm 1) to approximate the distribution of $w_t = v_t + r_t$. Then we obtained the LQ cost for 100 trials and averaged.

types of disturbances can be commonly observed in various systems. For example in automated driving, a vehicle experiences bounded disturbances for most of the time, but has a small probability to experience large disturbances (such as an extreme gust or a collision of neighboring vehicles). For a feedback system with perfect communications, the optimal ℓ_∞ controller and LQ controller for the scalar system (1) are identical (assuming no control cost). However, with communication constraints, not only the optimal ℓ_∞ controller and LQ controller are radically different, but also the mixture of disturbances pose great challenges in encoding/decoding strategies (as the system state cannot be defined by either a worst-case or a stochastic framework).

The ℓ_∞ controllers cannot stabilize such system because there is a non-zero probability for the fixed quantizer to saturate. Fig. 7 compares the performance of the LQ controller and the proposed hybrid controller. The LQ controller has

degraded performance when there exists additional disturbance r that cannot be well-defined using probability density function. Despite these difficulties, the proposed hybrid controller consistently produces robust performance for such disturbances. By exploiting the additional dimensions in the controller design space, this right blend of stochastic vs. worst-case, LQ vs. ℓ_∞ enables a robust controller under communication constraints.

VI. ACKNOWLEDGMENTS

REFERENCES

[1] P. Sterling and S. Laughlin, *Principles of neural design.* MIT Press, 2015.
[2] A. J. Nagengast, D. A. Braun, and D. M. Wolpert, "Risk sensitivity in a motor task with speed-accuracy trade-off," *Journal of neurophysiology*, vol. 105, no. 6, pp. 2668–2674, 2011.
[3] Y. Nakahira, Y. P. Leong, N. Matni, and J. C. Doyle, "Neuroscience motivation," in *Conference on Decision and Control*, 2016.
[4] Y. Nakahira, J. C. Doyle, and N. Matni, "Hard limits on robust control over delayed and quantized communications with applications to sensorimotor control," in *54th IEEE Conference on Decision and Control*, 2015, pp. 7522–7529.
[5] K. Zhou, J. C. Doyle, K. Glover *et al.*, *Robust and optimal control.* Prentice hall New Jersey, 1996.
[6] D. P. Bertsekas, *Dynamic programming and optimal control.* Athena Scientific Belmont, 1995.
[7] M. A. Dahleh and I. J. Diaz-Bobillo, *Control of uncertain systems: a linear programming approach.* Prentice-Hall, Inc., 1994.
[8] G. N. Nair and R. J. Evans, "Stabilizability of stochastic linear systems with finite feedback data rates," *SIAM Journal on Control and Optimization*, vol. 43, no. 2, pp. 413–436, 2004.
[9] S. Tatikonda and S. Mitter, "Control under communication constraints," *IEEE Transactions on Automatic Control*, vol. 49, no. 7, pp. 1056–1068, 2004.
[10] Y. Bar-Shalom and E. Tse, "Dual effect, certainty equivalence, and separation in stochastic control," *IEEE Transactions on Automatic Control*, vol. 19, no. 5, pp. 494–500, 1974.
[11] S. Tatikonda, A. Sahai, and S. Mitter, "Stochastic linear control over a communication channel," *IEEE transactions on Automatic Control*, vol. 49, no. 9, pp. 1549–1561, 2004.
[12] 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.
[13] A. J. Nagengast, D. A. Braun, and D. M. Wolpert, "Risk sensitivity in a motor task with speed-accuracy trade-off," *Journal of neurophysiology*, vol. 105, no. 6, pp. 2668–2674, 2011.
[14] Y. Nakahira, "LQ vs. ℓ_∞ in controller design for systems with delay and quantization," Caltech, Tech. Rep., 09 2016. [Online]. Available: http://users.cms.caltech.edu/~ynakahir
[15] S. Lloyd, "Least squares quantization in pcm," *IEEE Transactions on information theory*, vol. 28, no. 2, pp. 129–137, 1982.
[16] B. Sklar, *Digital communications.* Prentice Hall, 2001.
[17] G. N. Nair and R. J. Evans, "Stabilizability of stochastic linear systems with finite feedback data rates," *SIAM Journal on Control and Optimization*, vol. 43, no. 2, pp. 413–436, 2004.

APPENDIX

In this section, we prove two lemmas, followed by the proof of Theorem 1. Lemma 1 utilizes the concept of certainty equivalence [10][11]. Although a naive certainty equivalent control action $u_t = \mathbb{E}[x_t | s^{t-d}]$ is neither stabilizable nor optimal for a delayed system (see Remark 2), the control action can still have a simple structure $u_t = \mathbb{E}[z_t | s^{t-d}]$ for an auxiliary variable z_t.

Proof: (Lemma 1)

We first define the cost-to-go to be:

$$J_k := \mathbb{E}\left[x_N' P x_N + \sum_{t=k}^{N-1} x_t' P x_t + u_t' Q u_t\right]$$

for any $k < N$ and $J_N = \mathbb{E}[x_N' P x_N]$. Due to the delay in control, the action u_t can only depend on the information of w_0^{t-d-1}. Thus, the effect of the disturbance w_{t-d}^{t-1} would always remain in the state x_{t+1}. Let e_t be defined as the sum of those terms in x_t, which can be written as

$$e_t := w_{t-1} + A w_{t-2} + \cdots + A^{d-1} w_{t-d}$$

Consequently, the terms in x_t other than e_t, consisting of w_0^{t-d-1} and past quantization error, can potentially be controlled given a communication channel with infinite bandwidth. Let z_{t-d} be defined as the sum of those terms

$$z_t := x_t - e_t$$

Equivalently, one can also show that this definition is equivalent with

$$z_{t+1} = A z_t + A^d w_{t-d} + B u_t \qquad (30)$$

with the initial condition $z_0 = 0$. The term z_t is generated by w_0^{t-d-1} and the control action in response to it, whereas e_t is generated by w_{t-d}^{t-1}. Thus, z_t and e_t are independent, and we will repeatedly use this property.

Let $\{\bar{s}_t\}$ be the codewords that would be generated at time t if the system (1) has zero control $u_t \equiv 0$, and let

$$\bar{z}_{t+1} = A \bar{z}_t + A^d w_{t-d}, \quad \bar{z}_0 = 0 \qquad (31)$$

be the the variable z_t that would be generated if the system (1) has zero control $u_t \equiv 0$.

We will show by induction that the optimal cost-to-go can be written in the following form:

$$J_t = \mathbb{E}\left[z_t' P_t z_t | s^{t-d}\right] + \alpha_t \qquad (32)$$

for some P_t and that the value of α_t does not depend on the choice of control action, i.e.,

$$\mathbb{E}\left[\alpha_t | s^{t-d}\right] = \mathbb{E}\left[\alpha_t | \bar{s}^{t-d}\right]. \qquad (33)$$

At $t = N$, the cost to go satisfies

$$\begin{aligned}
J_N &:= \mathbb{E}[x_N' P x_N | s^{N-d}] \\
&= \mathbb{E}[(z_N + e_N) P (z_N + e_N) | s^{N-d}] \\
&= \mathbb{E}[z_N' P z_N | s^{N-d}] + \mathbb{E}[e_N' P e_N],
\end{aligned}$$

The last equality holds because z_N are independent from e_N. Since s^{N-d} does not depend on w_{N-d}^{N-1}, we can define $\alpha_N := \mathbb{E}[e_N' P_N e_N]$. Thus, the property (32) holds. Assume now that the property (32) holds for $t = k+1$, the optimal cost to go at time $t = k$ can be derived as follows:

$$\begin{aligned}
J_k &:= \min_{u_k} \mathbb{E}[x_k' P x_k + u_k' Q u_k + J_{k+1} | s^{k-d}] \\
&= \min_{u_k} \mathbb{E}[z_k'(P + A'P_{k+1}A)z_k + u_k'(Q + B'P_{k+1}B)u_k \\
&\quad + u_k' B' P_{k+1} A z_k + z_k' A' P_{k+1} B u_k | s^{k-d}] \\
&\quad + \mathbb{E}[e_k' P e_k] + w_{t-d}(A^d)' P_{k+1} A^d w_{t-d} | s^{k-d}] \\
&\quad + \mathbb{E}[\alpha_{k+1} | s^{k-d}],
\end{aligned}$$

where the last equality is due to the fact α_{k+1}, e_k', and w_{t-d} are independent of the control action u_t. Define the matrix $P_k = A'\left[P_{k+1} + P - P_{k+1}B(Q + B'P_{k+1}B)^{-1}B'P_{k+1}\right]A$. The control action

$$u_k = -(Q + B'P_{k+1}B)^{-1} B' P_{k+1} A \, \mathbb{E}[z_k | s^{k-d}]$$

minimizes the term $\mathbb{E}[z_k'(P + A'P_{k+1}A)z_k + u_k'(Q + B'P_{k+1}B)u_k | s^{k-d}]$. Substituting this control action u_k into J_k, we obtain the optimal cost-to-go

$$J_k = \mathbb{E}\left[z_k' P_k z_k | s^{k-d}\right] + \alpha_k,$$

where the term α_k satisfies the recursion

$$\begin{aligned}
\alpha_k &= \mathbb{E}[(z_k - \mathbb{E}[z_k | s^{k-d}])' G_k (z_k - \mathbb{E}[z_k | s^{k-d}]) | s^{k-d}] \\
&\quad + \mathbb{E}[e_k' P_{k+1} e_k + w_{t-d}(A^d)' P_{k+1} A^d w_{t-d} | s^{k-d}] \\
&\quad + \mathbb{E}[\alpha_{k+1} | s^{k-d}] \\
&= \mathbb{E}[(\bar{z}_k - \mathbb{E}[\bar{z}_k | \bar{s}^{k-d}])' G_k (\bar{z}_k - \mathbb{E}[\bar{z}_k | \bar{s}^{k-d}]) | \bar{s}^{k-d}] \\
&\quad + \mathbb{E}[e_k' P e_k + w_{t-d}(A^d)' P_{k+1} A^d w_{t-d}] \\
&\quad + \mathbb{E}[e_k' P e_k] + \mathbb{E}[\alpha_{k+1} | \bar{s}^{k-d}]
\end{aligned}$$

where $G_k := A'P_{k+1}B(Q + B'P_{k+1}B)^{-1}B'P_{k+1}A$. The last equality follows from the assumption that $\mathbb{E}[\alpha_{k+1} | s^{k-d}] = \mathbb{E}[\alpha_{k+1} | \bar{s}^{k-d}]$. ∎

Lemma 2: Let the sequence z_t be defined in (30), and $\bar{z}_t$ be defined in (31). Recall that the codewords generated by $\{z_t\}$ is $\{s_t\}$, and the codeword generated by $\{\bar{z}_t\}$ is $\{\bar{s}_t\}$. Then, the weighted estimation error covariance given these codewords are identical, i.e.,

$$\begin{aligned}
&\mathbb{E}[(z_t - \mathbb{E}[z_t | s^{t-d}])' G (z_t - \mathbb{E}[z_t | s^{t-d}])] \\
&= \mathbb{E}[(\bar{z}_t - \mathbb{E}[\bar{z}_t | \bar{s}^{t-d}])' G (\bar{z}_t - \mathbb{E}[\bar{z}_t | \bar{s}^{t-d}])].
\end{aligned}$$

Proof: (Lemma 2) Observe that $\mathbb{E}[z_t | s^{t-d}] = \mathbb{E}[\bar{z}_t + \sum_{k=1}^{t} A^{k-1} B u_{t-k} | s^{t-d}] = \mathbb{E}[\bar{z}_t | s^{t-d}] + \sum_{k=1}^{t} A^{k-1} B u_{t-k} = \mathbb{E}[\bar{z}_t | \bar{s}^{t-d}] + \sum_{k=1}^{t} A^{k-1} B u_{t-k}$. Therefore, we can establish the following relation:

$$\begin{aligned}
&z_t - \mathbb{E}[z_t | s^t] \\
&= \bar{z}_t + \sum_{k=1}^{t} A^{k-1} B u_{t-k} - \left(\mathbb{E}[\bar{z}_t | \bar{s}^t] + \sum_{k=1}^{t} A^{k-1} B u_{t-k}\right) \\
&= \bar{z}_t - \mathbb{E}[\bar{z}_t | \bar{s}^{t-d}].
\end{aligned}$$

∎

Remark 4: Since the encoder and decoder has memory, the conditional expectation conditioned on s^t is equivalent with the conditional expectation conditioned on (s^t, u^{t-1}).

Remark 5: Intuitively, the Lemma 2 states that we can negate all the effects of the control action in order to obtain $\bar{z}_t$ because u_0^t is generated from s^{t-d}.

Lemma 3: Consider a scalar Gauss Markov sequence $\{y_t\}$ satisfying

$$y_{t+1} = a y_t + v_t.$$

with the initial condition $y_0 = 0$. Assume $v_t \stackrel{iid}{\sim} \mathcal{N}(0, \sigma^2)$. Assume that we can only communicate $R(> \log_2 |a|)$ bits per sampling interval, i.e., the estimate $\hat{y}_t$ of y_t is generated from the codewords $s^t, s_t \in |S|, |S| = 2^R$, and each codeword s_t is generated from (y^t, s^{t-d}). Then, the following inequality holds:

$$\lim_{t \to \infty} \frac{1}{N} \mathbb{E}\left[\sum_{t=1}^{N} (y_t - \hat{y}_t)^2\right] \geq \frac{\sigma^2}{2^{2R} - a^2}$$

Proof: (Lemma 3) Let D_t be the distortion bound at time t, i.e.,

$$\min_{\hat{y}_t} \mathbb{E}[y_t - \hat{y}_t | \hat{y}^{t-1}] \geq D_t$$

From $y_1 \stackrel{iid}{\sim} \mathcal{N}(0, \sigma^2)$, we obtain $D_1 = 2^{-2R}\sigma^2$. Now we will develop a recursive relation of D_t. Notice that we can rewrite $y_{t+1} = a(y_t - \hat{y}_t) + v_t + a\hat{y}_t$, where $\hat{y}_t$ is known. Since $a(y_t - \hat{y}_t)$

and w_t are uncorrelated, the term $a(y_t - \hat{y}_t) + w_t$ has variance greater than $a^2 D_t + \sigma^2$. Thus, we obtain

$$\mathbb{E}[y_{t+1} - \hat{y}_{t+1} \mid \hat{y}_t] \geq (a^2 D_t + \sigma^2) 2^{-2R},$$

yielding $D_{t+1} = (a^2 D_t + \sigma^2) 2^{-2R}$. Since we have assumed $|a|/2^R < 1$, $\lim_{t\to\infty} D_t = \frac{\sigma^2}{2^{2R} - a^2}$. $\blacksquare$

Proof: (Theorem 1)

From the previous analysis,

$$\lim_{t\to\infty} \mathbb{E}[x_t' P x_t + u_t' Q u_t]$$
$$= \lim_{N\to\infty} \frac{1}{N} \mathbb{E}\left[x_N' P x_N + \sum_{t=0}^{N-1} x_t' P x_t + u_t' Q u_t \right]$$
$$= \lim_{N\to\infty} J_0/N$$
$$= \lim_{N\to\infty} (\mathbb{E}[z_{d+1} P_{d+1} z_{d+1}] + \alpha_0)/N$$

Now we will find the optimal encoder that minimizes α_0. Recall that $G_t = A' P_{t+1} B (Q + B' P_{t+1} B)^{-1} B' P_{t+1} A$. The term α_t satisfies the recursive relation

$$\alpha_t = \mathbb{E}[(\bar{z}_t - \mathbb{E}[\bar{z}_t | \bar{s}^{t-d}])' G_t (z_t - \mathbb{E}[\bar{z}_t | \bar{s}^{t-d}]) \mid \bar{s}^{t-d}]$$
$$+ \mathbb{E}[e_t' P e_t + w_{t-d} (A^d)' P_{k+1} A^d w_{t-d}]$$
$$+ \mathbb{E}[\alpha_{t+1} | \bar{s}^{t-d}].$$

Recall that the matrices $P^\star$ and $G^\star$ are defined in (7). The matrix $P^\star$ from Riccati difference equation (7) has a unique solution since the pair of matrices (A, B) is stabilizable. Combining the above relation with the initial condition $\alpha_N := \mathbb{E}[e_N' P_N e_N]$ yields

$$\lim_{N\to\infty} \frac{1}{N} \alpha_0$$
$$= \lim_{N\to\infty} \frac{1}{N} \sum_{t=0}^{N-1} \mathbb{E}[(\bar{z}_t - \mathbb{E}[\bar{z}_t | \bar{s}^{t-d}])' G_t (z_t - \mathbb{E}[\bar{z}_t | \bar{s}^{t-d}]) \mid \bar{s}^{t-d}]$$
$$+ P(1 + A^2 + A^4 + \cdots + A^{2(d-1)}) \sigma^2 + P^\star A^{2d} \sigma^2$$
$$= \lim_{N\to\infty} \frac{1}{N} \sum_{t=0}^{N-1} \mathbb{E}[(\bar{z}_t - \mathbb{E}[\bar{z}_t | \bar{s}^{t-d}])' G^\star (z_t - \mathbb{E}[\bar{z}_t | \bar{s}^{t-d}]) \mid \bar{s}^{t-d}]$$
$$+ P \sum_{i=0}^{d-1} A^{2i} \sigma^2 + P^\star A^{2d} \sigma^2.$$

From Lemma 3, the first term satisfies

$$\mathbb{E}[(\bar{z}_t - \mathbb{E}[\bar{z}_t | \bar{s}^{t-d}])' G^\star (\bar{z}_t - \mathbb{E}[\bar{z}_t | \bar{s}^{t-d}]) \mid \bar{s}^{t-d}]$$
$$\geq G^\star A^{2d} \frac{\sigma^2}{2^{2R} - A^2}$$

Therefore, we have obtained (8). $\blacksquare$

Proof: (Theorem 2) We first prove the lower bound for $\tau = 0$. Let $\{E_k\}$ be the event that the controller switches at time k, *i.e.*,

$$E_k = \{|z_t| \leq \Psi(L)/A \text{ for all } t < k \text{ and } |z_k| > \Psi(L)/A\},$$

Notice that $\{E_k\}$ a sequence of a mutually exclusive set of events, and that $\mathbb{P}(E_k) = 0$ for $k \leq d$ (since $z_t = 0$ for $t \leq d$ by definition). Let $\{F_k\}$ be the event that the disturbance first exceeds L in amplitude at time k, *i.e.*,

$$F_k = \{|w_t| \leq L \text{ for all } t < k \text{ and } |w_k| > L\}$$

The sequences $\{E_k\}$ are mutually exclusive set of events, and $\lim_{\tau\to\infty} \sum_{i=0}^{\tau} \mathbb{P}(E_i) = 1$. Same conditions also hold for $\{F_k\}$, *i.e.*, $\lim_{\tau\to\infty} \sum_{i=0}^{\tau} \mathbb{P}(F_i) = 1$. From $\cup_{i\geq k} E_i \subset \cup_{i\geq k} F_i$, we obtain

$$\sum_{i=k-d-1}^{\infty} \mathbb{P}(F_i) \leq \sum_{i=k}^{\infty} \mathbb{P}(E_i) \qquad (34)$$

for any $k \in \mathbb{N}$. Using (34), the expected switching time can be bounded below by

$$\mathbb{E}[T_s(\tau)] = \sum_{k=0}^{\infty} k \, \mathbb{P}(E_k)$$
$$= \sum_{k=0}^{\infty} k \, \mathbb{P}(E_k) - \sum_{k=0}^{\infty} k \, \mathbb{P}(E_{k+d}) + \sum_{k=0}^{\infty} k \, \mathbb{P}(E_{k+d})$$
$$= d + \sum_{k=0}^{\infty} k \, \mathbb{P}(E_{k+d})$$
$$= d + \sum_{k=1}^{\infty} \sum_{i=k}^{\infty} \mathbb{P}(E_{i+d})$$
$$\geq d + \sum_{k=1}^{\infty} \sum_{i=k}^{\infty} \mathbb{P}(F_{i-1})$$
$$= d + \sum_{k=1}^{\infty} k \, \mathbb{P}(F_{k-1})$$
$$= d + \sum_{k=1}^{\infty} k \left(1 - \mathbb{P}(|w| > L)\right)^{k-1} \mathbb{P}(|w| > L)$$
$$= d + \mathbb{P}(|w| > L)^{-1},$$

where the last equality can be interpreted as computing the mean of a geometric distribution with failure probability $\mathbb{P}(|w| > L)$.

Next, notice that $|z_t| \leq \Psi(L)/A$ and $|w_{t-d}| \leq L$ implies $|z_{t+1}| \leq \Psi(L)/A$. Thus, we can apply the argument in $\tau = 0$ to obtain the lower bound for $\tau \in \mathcal{T}$:

$$\mathbb{E}[T_s(\tau)] \leq \mathbb{P}(|w| > L)^{-1}.$$

Next, we prove the convergence for $\tau = 0$, *i.e.*, $\mathbb{E}[T_s(0)] \stackrel{R\to\infty}{\to} d + \mathbb{P}(|w| > L)^{-1}$. Since $d + \sum_{k=1}^{\infty} \sum_{i=k}^{\infty} \mathbb{P}(E_{i+d}) \geq d + \sum_{k=1}^{\infty} \sum_{i=k}^{\infty} \mathbb{P}(F_{i-1})$ is the only inequality from the above analysis, it is suffice to show that $\left|\sum_{k=1}^{\infty} \sum_{i=k}^{\infty} \mathbb{P}(E_{i+d}) - \sum_{k=1}^{\infty} \sum_{i=k}^{\infty} \mathbb{P}(F_{i-1})\right| \to 0$. From $\|q\|_\infty \stackrel{R\to\infty}{\to} 0$ [4], $z_t \to A^d w_{t-d-1}$, and thus $\mathbb{P}(F_{t-d-1}) \to \mathbb{P}(E_t)$. This implies that

$$\left| \sum_{i=k-1}^{\infty} \mathbb{P}(F_{i-1}) - \sum_{i=k}^{\infty} \mathbb{P}(E_{i+d}) \right|$$
$$= \left| \left(1 - \sum_{i=0}^{k-2} \mathbb{P}(F_i)\right) - \left(1 - \sum_{i=0}^{k-1} \mathbb{P}(E_{i+d})\right) \right|$$
$$\to 0 \text{ as } R \to \infty.$$

holds for any $k \in \mathbb{N}$. Since both $\sum_{k=1}^{\infty} \sum_{i=k}^{\infty} \mathbb{P}(E_{i+d})$ and $\sum_{k=1}^{\infty} \sum_{i=k}^{\infty} \mathbb{P}(F_{i-1})$ are bounded, for any $\epsilon > 0$ there exits a sufficiently large T such that $\tau > T$ implies

$$\sum_{k=\tau}^{\infty} \sum_{i=k}^{\infty} \mathbb{P}(E_{i+d}) \leq \epsilon/4 \quad \text{and} \quad \sum_{k=\tau}^{\infty} \sum_{i=k}^{\infty} \mathbb{P}(F_{i-1}) \leq \epsilon/4,$$

and sufficiently large $\bar{R}$ such that $R > \bar{R}$ implies

$$\sum_{k=1}^{\tau} \sum_{i=k}^{\infty} \mathbb{P}(E_{i+d}) \leq \epsilon/4 \quad \text{and} \quad \sum_{k=1}^{\tau} \sum_{i=k}^{\infty} \mathbb{P}(F_{i-1}) \leq \epsilon/4,$$

which jointly yields

$$\left| \sum_{k=1}^{\infty} \sum_{i=k}^{\infty} \mathbb{P}(E_{i+d}) - \sum_{k=1}^{\infty} \sum_{i=k}^{\infty} \mathbb{P}(F_{i-1}) \right| \leq \epsilon. \qquad (35)$$

The case for $\tau \in \mathcal{T}$ also follows the same argument and thus is omitted due to space constraint. $\blacksquare$