

The necessity and sufficiency of anytime capacity for stabilization of a linear system over a noisy communication link

Part I: scalar systems

Anant Sahai and Sanjoy Mitter
sahai@eecs.berkeley.edu, mitter@mit.edu

Abstract

We review how Shannon's classical notion of capacity is not enough to characterize a noisy communication channel if the channel is intended to be used as part of a feedback loop to stabilize an unstable scalar linear system. While classical capacity is not enough, another sense of capacity (parametrized by reliability) called "anytime capacity" is shown to be necessary for the stabilization of an unstable process. The required rate is given by the log of the unstable system gain and the required reliability comes from the sense of stability desired. A consequence of this necessity result is a sequential generalization of the Schalkwijk/Kailath scheme for communication over the AWGN channel with feedback.

In cases of sufficiently rich information patterns between the encoder and decoder, adequate anytime capacity is also shown to be sufficient for there to exist a stabilizing controller. These sufficiency results are then generalized to cases with noisy observations, delayed control actions, and without any explicit feedback between the observer and the controller. Both necessary and sufficient conditions are extended to continuous time systems as well. We close with comments discussing a hierarchy of difficulty for communication problems and how these results establish where stabilization problems sit in that hierarchy.

Index Terms

Real-time information theory, reliability functions, error exponents, feedback, anytime decoding, sequential coding, control over noisy channels

A. Sahai is with the Department of Electrical Engineering and Computer Science, U.C. Berkeley. Early portions of this work appeared in his doctoral dissertation and a few other results were presented at the 2004 Conference on Decision and Control.

Department of Electrical Engineering and Computer Science at the Massachusetts Institute of Technology. Support for S.K. Mitter was provided by the Army Research Office under the MURI Grant: Data Fusion in Large Arrays of Microsensors DAAD19-00-1-0466 and the Department of Defense MURI Grant: Complex Adaptive Networks for Cooperative Control Subaward #03-132 and the National Science Foundation Grant CCR-0325774.

The necessity and sufficiency of anytime capacity for stabilization of a linear system over a noisy communication link

Part I: scalar systems

I. INTRODUCTION

For communication theorists, Shannon's classical channel capacity theorems are not just beautiful mathematical results, they are useful in practice as well. They let us summarize a diverse range of channels by a single figure of merit: the capacity. For most non-interactive point-to-point communication applications, the Shannon capacity of a channel provides an upper bound on performance in terms of end-to-end distortion through the distortion-rate function. As far as distortion is concerned, all that matters is the channel capacity and the nature of the source. Given enough tolerance for end-to-end delay, the source can be encoded into bits and those bits can be reliably transported across the noisy channel if the rate is less than the Shannon capacity. As long as the source, distortion, and channel are well-behaved [1], [2], there is asymptotically no loss in separating the problems of source and channel coding. This provides a justification for the layered architecture that lets engineers isolate the problem of reliable communication from that of using the communicated information. Recent advances in coding theory have also made it possible to approach the capacity bounds very closely in practical systems.

In order to extend our understanding of communication to interactive settings, it is essential to have some model for interaction. Schulman and others have studied interaction in the context of distributed computation [3], [4]. The interaction there is between computational agents that have access to some private data and wish to perform a global computation in a distributed way. The computational agents can only communicate with each other through noisy channels. In Schulman's formulation, capacity is not a question of major interest since constant factor slowdowns are considered acceptable.¹ Fundamentally, this is a consequence of being able to design all the system dynamics. The rich field of automatic control provides an interactive context to study capacity requirements since the plant dynamics are given, rather than something that can be designed. In control, we consider interaction between an observer that gets to see the plant and a controller that gets to control it. These two can be connected by a noisy channel.

Shannon himself had suggested looking to control problems for more insight into reliable communication [5].

“... can be pursued further and is related to a duality between past and future² and the notions of control and knowledge. Thus we may have knowledge of the past and cannot control it; we may control the future but have no knowledge of it.”

We are far from the first to attempt to bring together information and control theory. In [7], Ho, Kastner, and Wong drew out a detailed diagram in which they summarized the then known relationships among team theory, signaling, and information theory from the perspective of distributed control. Rather than taking such a broad perspective, we instead ask whether Shannon's classical capacity is the appropriate characterization for communication channels arising in distributed control systems. Our interest is in understanding the fundamental relationship between problems of stabilization and problems of communication.

Tatikonda's recent work on sequential rate distortion theory provides an information-theoretic lower-bound on the achievable performance of a control system over a channel. Because this bound is sometimes infinite, it also implies that there is a fundamental rate of information production, namely the sum of the logs of the unstable eigenvalues of the plant, that is invariantly attached to an unstable linear discrete-time process [8], [9]. This particular notion of rate was justified by showing how to stabilize the system over a noiseless feedback link with capacity greater than the intrinsic rate for the unstable process.³ Nair et al. extended this to cover the case of unbounded disturbances and observation noise under suitable conditions [10], [11]. In addition to noiseless channels, the results were extended for almost-sure stabilization in the context of undisturbed⁴ control systems with bounded initial conditions being stabilized over certain noisy channels [12].

We had previously showed that it is possible to stabilize persistently disturbed controlled Gauss-Markov processes over suitable power-constrained AWGN (Additive White Gaussian Noise) channels [13], [14] where it turns out that Shannon capacity is tight and linear observers and controllers are sufficient to achieve stabilization [15]. In contrast, we showed that the Shannon capacity of the binary erasure channel (BEC) is not sufficient to check stabilizability and introduced the anytime capacity as a candidate figure of merit [16]. Following up on our treatment of the BEC case, Martins et al. have studied more general erasure-type models and have also incorporated bounded model uncertainty in the plant [17]. There is also related work by Elia that uses ideas from robust control to deal with communication uncertainty in a mixed continuous/discrete context, but restricting to linear operations [18], [19]. Basar and his students have also considered such problems and have studied the impact of a noisy channels on both the observations and the controls [20]. The area of control with communications constraints continues to attract attention and the reader is directed to the recent September 2004 issue

¹Furthermore, such constant factor slowdowns appear to be unavoidable when facing the very general class of interactive computational problems.

²The differing roles of the past and future are made clear in [6].

³The sequential rate-distortion bound is generally not attained even at higher rates except in the case of perfectly matched channels.

⁴In seminal work [12], there is no persistent disturbance acting on the unstable plant.

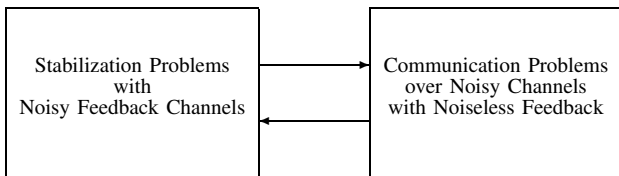


Fig. 1. The “equivalence” between stabilization over noisy feedback channels and reliable communication over noisy channels with feedback is the main result established in this paper.

of IEEE Transactions on Automatic Control and the articles therein for a more comprehensive survey.

Many of the issues that arise in the control context also arise for the conceptually simpler problem of merely estimating an unstable open-loop process⁵, across a noisy channel. For this estimation problem in the limit of large, but finite, end-to-end delays, we have proved a source coding theorem that shows that the distortion-rate bound is achievable. Furthermore, it is possible to characterize the information being produced by an unstable process [23]. It turns out that such processes produce two qualitatively distinct types of information when it comes to transport over a noisy channel. In addition to the classical Shannon-type of information found in traditional rate-distortion settings⁶, there is an essential core of information that captures the unstable nature of the source. While classical Shannon reliability suffices for the classical information, this unstable core requires *anytime* reliability for transport across a noisy channel.⁷ As also discussed in this paper, anytime reliability is a sense of reliable transmission that lies between Shannon’s classical ϵ -sense of reliable transmission and his zero-error reliability [25]. In [23], we also review how the sense of anytime reliability is linked to classical work on sequential tree codes with bounded delay decoding.⁸

The new feature in control systems is their essential interactivity. The information to be communicated is not a message known in advance that is used by some completely separate entity. Rather, it evolves through time and is used to control the very process being encoded. This introduces two interesting issues. First, causality is strictly enforced. The encoder and controller must act in real time and so taking the limit of large delays must be interpreted very carefully. Second, it is unclear what the status of the controlled process is. If the controller succeeds in stabilizing the process, it is no longer unstable. As explored in Section II-D, a purely external non-interactive observer could treat the question of encoding the controlled closed-loop system state using classical tools for the encoding and communication of a stationary ergodic

⁵The unstable open-loop processes discussed here are first-order nonstationary autoregressive processes [21], of which an important special case is the Wiener process considered by Berger [22].

⁶In [23], we show how the classical part of the information determines the shape of the rate-distortion curve, while the unstable core is responsible for a shift of this curve along the rate axis.

⁷How to communicate such unstable processes over noisy channels had been an open problem since Berger had first developed a source-coding theorem for the Wiener process [24]. Berger had conjectured that it was impossible to transport such processes over generic noisy channels with asymptotically finite end-to-end distortion using traditional means.

⁸Reference [26] raised the possibility of such a connection early on.

process. Despite having to observe and encode the exact same closed-loop process, the observer internal to the control system requires a channel as good as that required to communicate the unstable open-loop process. This seemingly paradoxical situation illustrates what can happen when the encoding of information and its use are coupled together by interactivity.

In this paper (Part I), the basic equivalence between feedback stabilization and reliable communication is established. The scalar problem (Figure 2) is formally introduced in Section II where classical capacity concepts are also shown to be inadequate. In Section III, it is shown that adequate feedback anytime capacity is necessary for there to exist an observer/controller pair able to stabilize the unstable system across the noisy channel. This connection is also used to give a sequential anytime version of the Schalkwijk/Kailath scheme for the AWGN channel with noiseless feedback.

Section IV shows the sufficiency of feedback anytime capacity for situations where the observer has noiseless access to the channel outputs. In Section V, these sufficiency results are generalized to the case where the observer only has noisy access to the plant state. Since the necessary and sufficient conditions are tight in many cases, these results show the asymptotic equivalence between the problem of control with “noisy feedback” and the problem of reliable sequential communication with noiseless feedback. In Section VI, these results are further extended to the continuous time setting. Finally, Section VII justifies why the problem of stabilization of an unstable linear control system is “universal” in the same sense that the Shannon formulation of reliable transmission of messages over a noisy channel with (or without) feedback is universal. This is done by introducing a hierarchy of communication problems in which problems at a given level are equivalent to each other in terms of which channels are good enough to solve them. Problems high in the hierarchy are fundamentally more challenging than the ones below them in terms of what they require from the noisy channel.

In Part II, the necessity and sufficiency results are generalized to the case of multivariable control systems on an unstable eigenvalue by eigenvalue basis. The role of anytime capacity is played by a rate region corresponding to a vector of anytime reliabilities. If there is no explicit channel output feedback, the intrinsic delay of the control system’s input-output behavior plays an important role. It shows that two systems with the same unstable eigenvalues can still have potentially different channel requirements. These results establish that in interactive settings, a single “application” can fundamentally require different senses of reliability for its data streams. No single number can adequately summarize the channel and any layered communication architecture should allow applications to adjust reliabilities on bitstreams.

There are many results in this paper. In order not to burden the reader with repetitive details and unnecessarily lengthen this paper, we have adopted a discursive style in some of the proofs. The reader should not have any difficulty in filling in the omitted details.

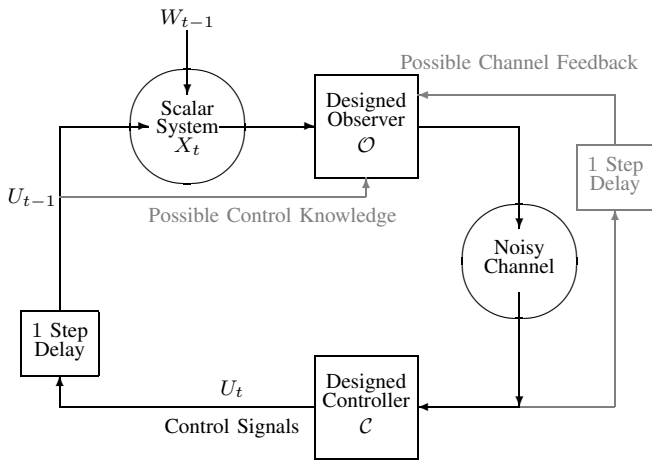


Fig. 2. Control over a noisy communication channel. The unstable scalar system is persistently disturbed by W_t and must be kept stable in closed-loop through the actions of \mathcal{O} , \mathcal{C} .

II. PROBLEM DEFINITION AND BASIC CHALLENGES

Section II-A formally introduces the control problem of stabilizing an unstable scalar linear system driven by both a control signal and a bounded disturbance. In Section II-B, classical notions of capacity are reviewed along with how to stabilize an unstable system with a finite rate noiseless channel. In Section II-C, it is shown by example that the classical concepts are inadequate when it comes to evaluating a noisy channel for control purposes. Shannon's regular capacity is too optimistic and zero-error capacity is too pessimistic. Finally, Section II-D shows that the core issue of interactivity is different than merely requiring the encoders and decoders to be delay-free.

A. The control problem

$$X_{t+1} = \lambda X_t + U_t + W_t, \quad t \geq 0 \quad (1)$$

where $\{X_t\}$ is a \mathbb{R} -valued state process. $\{U_t\}$ is a \mathbb{R} -valued control process and $\{W_t\}$ is a bounded noise/disturbance process s.t. $|W_t| \leq \frac{\Omega}{2}$. This bound is assumed to hold with certainty. For convenience, we also assume a known initial condition $X_0 = 0$.

To make things interesting, consider $\lambda > 1$ so the open-loop system is exponentially unstable. The distributed nature of the problem (shown in Figure 2) comes from having a noisy communication channel in the feedback path. The observer/encoder system \mathcal{O} observes X_t and generates inputs a_t to the channel. It may or may not have access to the control signals U_t or past channel outputs B_{t-1} as well. The decoder/controller⁹ system \mathcal{C} observes channel outputs B_t and generates control signals U_t . Both \mathcal{O} , \mathcal{C} are allowed to have unbounded memory and to be nonlinear in general.

Definition 2.1: A closed-loop dynamic system with state X_t is f -stable if $\mathcal{P}(|X_t| > m) < f(m)$ for all $t \geq 0$.

⁹Because the decoder and controller are both on the same side of the communication channel, they can be lumped together into a single box.

This definition requires the probability of a large state value to be appropriately bounded. A looser sense of stability is given by:

Definition 2.2: A closed-loop dynamic system with state X_t is η -stable if there exists a constant K s.t. $E[|X_t|^\eta] \leq K$ for all $t \geq 0$.

In both definitions, the bound is required to hold for all possible sequences of bounded disturbances $\{W_t\}$ that satisfy the given bound Ω . We do not assume any specific probability model governing the disturbances. Rather than having to specify a specific target for the tail probability f , holding the η -moment within bounds is a way of keeping large deviations rare. The larger η is, the more strongly very large deviations are penalized. The advantage of η -stability is that it allows constant factors to be ignored while making sharp asymptotic statements. Furthermore, Section III-C shows that for generic DMCs, no sense stronger than η -stability is feasible.

The goal in this paper is to find necessary and sufficient conditions on the noisy channel for there to exist an observer \mathcal{O} and controller \mathcal{C} so that the closed loop system shown in Figure 2 is stable in the sense of definitions 2.1 or 2.2. The problem is considered under different information patterns corresponding to different assumptions about what information is available at the observer \mathcal{O} . The controller is always assumed to just have access to the entire past history¹⁰ of channel outputs.

For discrete-time linear systems, the intrinsic rate of information production (in units of bits per time) equals the sum of the logarithms (base 2) of the unstable eigenvalues [9]. In the scalar case studied here, this is just $\log_2 \lambda$. This means that it is generically¹¹ impossible to stabilize the system in any reasonable sense if the feedback channel's Shannon classical capacity $C < \log_2 \lambda$.

B. Classical notions of channels and capacity

Definition 2.3: A discrete time channel is a probabilistic system with an input. At every time step t , it takes an input $a_t \in \mathcal{A}$ and produces an output $b_t \in \mathcal{B}$ with probability¹² $p(B_t|a_1^t, b_1^{t-1})$ where the notation a_1^t is shorthand for the sequence a_1, a_2, \dots, a_t . In general, the current channel output is allowed to depend on all inputs so far as well as on past outputs.

The channel is *memoryless* if conditioned on a_t , B_t is independent of any other random variable in the system that occurs at time t or earlier. All that needs to be specified is $p(B_t|a_t)$.

The maximum rate achievable for a given sense of reliable communication is called the associated capacity. Shannon's

¹⁰In Section III-C.3, it is shown that anything less than that can not work in general.

¹¹There are pathological cases where it is possible to stabilize a system with less rate. These occur when the driving disturbance is particularly structured instead of just being unknown but bounded. An example is when the disturbance only takes on values ± 1 while $\lambda = 4$. Clearly only one bit per unit time is required even though $\log_2 \lambda = 2$.

¹²This is a probability mass function in the case of discrete alphabets \mathcal{B} , but is more generally an appropriate probability measure over the output alphabet \mathcal{B} .

classical reliability requires that after a suitably large end-to-end delay¹³ n that the average probability of error on each bit is below a specified ϵ . *Shannon classical capacity* C can also be calculated in the case of memoryless channels by solving an optimization problem:

$$C = \sup_{\mathcal{P}(A)} I(A; B)$$

where the maximization is over the input probability distribution and $I(A; B)$ represents the mutual information through the channel [1]. This is referred to as a single letter characterization of channel capacity for memoryless channels. Similar formulae exist using limits in cases of channels with memory. There is another sense of reliability and its associated capacity C_0 called *zero-error capacity* which requires the probability of error to be exactly zero with sufficiently large n . It does not have a simple single-letter characterization [25].

Example 2.1: Consider a system (1) with $\Omega = 1$ and $\lambda = \frac{3}{2}$. Suppose that the memoryless communication channel is a noiseless one bit channel. So $\mathcal{A} = \mathcal{B} = \{0, 1\}$ and $p(B_t = 1|a_t = 1) = p(B_t = 0|a_t = 0) = 1$ while $p(B_t = 1|a_t = 0) = p(B_t = 0|a_t = 1) = 0$. This channel has $C_0 = C = 1 > \log_2 \frac{3}{2}$.

Use a memoryless observer

$$\mathcal{O}(x) = \begin{cases} 0 & \text{if } x \leq 0 \\ 1 & \text{if } x > 0 \end{cases}$$

and memoryless controller

$$\mathcal{C}(B) = \begin{cases} +\frac{3}{2} & \text{if } B = 0 \\ -\frac{3}{2} & \text{if } B = 1 \end{cases}$$

Assume that the closed loop system state is within the interval $[-2, +2]$. If it is positive, then it is in the interval $[0, +2]$. At the next time, $\frac{3}{2}X + W$ would be in the interval $[-\frac{1}{2}, \frac{7}{2}]$. The applied control of $-\frac{3}{2}$ shifts the state back to within the interval $[-2, +2]$. The same argument holds by symmetry on the negative side. Since it starts at 0, by induction it will stay within $[-2, +2]$ forever. As a consequence, the second moment will stay less than 4 for all time, and all the other moments will be similarly bounded.

In addition to the Shannon and zero-error senses of reliability, information theory has various reliability functions. Such reliability functions (or error exponents) are traditionally considered an internal matter for channel coding and were viewed as mathematically tractable proxies for the issue of implementation complexity [1]. Reliability functions study how fast the probability of error goes to zero as the relevant system parameter is increased. Thus, the reliability functions for block-codes are given in terms of the block length, reliability functions for convolutional codes in terms of the constraint length[27], and reliability functions for variable-length codes in terms of the expected block length [28]. With the rise of sparse code constructions and iterative decoding, the prominence of error exponents in channel coding has

¹³Traditionally, the community has used block-length for a block code as the fundamental quantity rather than delay. It is easy to see that doing encoding and decoding in blocks of size n corresponds to a delay of between n and $2n$ on the individual bits being communicated.

diminished since the computational burden is not superlinear in the block-length.

For memoryless channels, the presence or absence of feedback does not alter the classical Shannon capacity [1]. More surprisingly, for symmetric DMCs, the fixed block coding reliability functions also do not change with feedback, at least in the high rate regime [29]. From a control perspective, this is the first indication that neither Shannon's capacity nor block-coding reliability functions are the perfect fit for control applications.

C. Counterexample showing classical concepts are inadequate

We use erasure channels to construct a counterexample showing the inadequacy of the Shannon classical capacity in characterizing channels for control. While both erasure and AWGN channels are easy to deal with, it turns out that AWGN channels can not be used for a counterexample since they can be treated in the classical LQG framework [15]. The deeper reason for why AWGN channels do not provide a counterexample is given in Section III-C.4.

1) *Erasure channels:* The packet erasure channel models situations where errors can be reliably detected at the receiver. In the model, sometimes the packet being sent does not make it through with probability δ , but otherwise it makes it through correctly. Explicitly:

Definition 2.4: The L -bit packet erasure channel is a memoryless channel with $\mathcal{A} = \{0, 1\}^L$, $\mathcal{B} = \{0, 1\}^L \cup \{\emptyset\}$ and $p(x|x) = 1 - \delta$ while $p(\emptyset|x) = \delta$.

It is well known that the Shannon capacity of the packet erasure channel is $(1 - \delta)L$ bits per channel use regardless of whether the encoder has feedback or not [1]. Furthermore, because a long string of erasures is always possible, the zero-error capacity C_0 of this channel is 0. There are also variable-length packet erasure channels where the packet-length is something the encoder can choose. See [30] for a discussion of such channels.

To construct a simple counterexample, consider a further abstraction:

Definition 2.5: The *real packet erasure channel* has $\mathcal{A} = \mathcal{B} = \mathbb{R}$ and $p(x|x) = 1 - \delta$ while $p(0|x) = \delta$.

This model has also been explored in the context of Kalman filtering with lossy observations [31], [32]. It has infinite classical capacity since a single real number can carry arbitrarily many bits within its binary expansion, while the zero-error capacity remains 0.

2) *The inadequacy of Shannon capacity:* Consider the problem from example 2.1, except over the real erasure channel instead of the one bit noiseless channel. The goal is for the second moment to be bounded ($\eta = 2$) and recall that $\lambda = \frac{3}{2}$. Let $\delta = \frac{1}{2}$ so that there is a 50% chance of any real number being erased. Assume the bounded disturbance W_t , assume that it is zero-mean and iid with variance σ^2 . By assuming an explicit probability model for the disturbance, the problem is only made easier as compared to the arbitrarily-varying but bounded model introduced earlier.

In this case, the optimal control is obvious — set $a_t = X_t$ as the channel input and use $U_t = -\lambda B_t$ as the control.

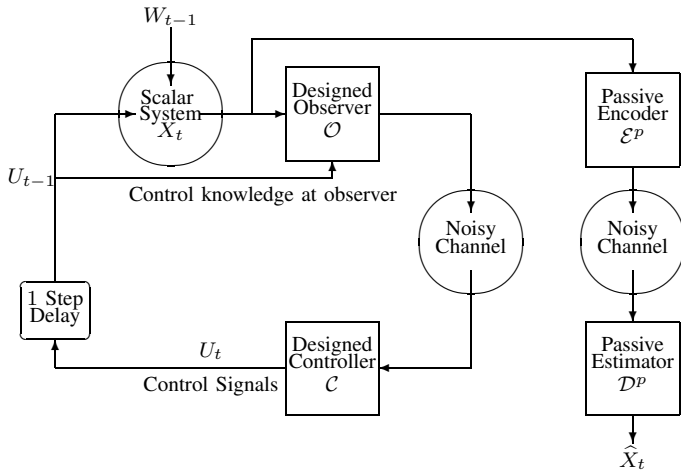


Fig. 3. The control system with an additional passive joint source-channel encoder \mathcal{E}^P watching the closed loop state X_t and communicating it to a passive estimator \mathcal{D}^P . The controller \mathcal{C} implicitly needs a good causal estimate for X_t and the passive estimator \mathcal{D}^P explicitly needs the same thing. Which requires the better channel?

With every successful reception, the system state is reset to the initial condition of zero. For an arbitrary time t , the time since it was last reset is distributed like a geometric- $\frac{1}{2}$ random variable. Thus the second moment is:

$$\begin{aligned}
 E[|X_{t+1}|^2] &> \sum_{i=0}^t \frac{1}{2} \left(\frac{1}{2}\right)^i E\left[\left(\sum_{j=0}^i \left(\frac{3}{2}\right)^j W_{t-j}\right)^2\right] \\
 &= \sum_{i=0}^t \frac{1}{2} \left(\frac{1}{2}\right)^i \sum_{j=0}^i \sum_{k=0}^i \left(\frac{3}{2}\right)^{j+k} E[W_{t-j} W_{t-k}] \\
 &= \sum_{i=0}^t \left(\frac{1}{2}\right)^{i+1} \sum_{j=0}^i \left(\frac{9}{4}\right)^j \sigma^2 \\
 &= \frac{4\sigma^2}{5} \sum_{i=0}^t \left(\left(\frac{9}{8}\right)^{i+1} - \left(\frac{1}{2}\right)^{i+1} \right)
 \end{aligned}$$

This diverges as $t \rightarrow \infty$ since $\frac{9}{8} > 1$.

Notice that the root of the problem is that $(\frac{3}{2})^2(\frac{1}{2}) > 1$. Intuitively, the system is exploding faster than the noisy channel is able to give reliability. This causes the second moment to diverge. In contrast, the first moment $E[|X_t|]$ is bounded for all t since $(\frac{3}{2})(\frac{1}{2}) < 1$.

The adequacy of the channel depends on which moment is required to be bounded. *Thus no single-number characterization like classical capacity can give the figure-of-merit needed to evaluate a channel for control applications.*

D. Non-interactive observation of a closed-loop process

Consider the system shown in Figure 3. In this, there is an additional passive joint source-channel encoder \mathcal{E}^P watching the closed loop state X_t and communicating it to a passive estimator \mathcal{D}^P through a second independent noisy channel. Both the passive and internal observers have access to the same plant state and we can also require the passive encoder and decoder to be causal — no end-to-end delay is permitted.

At first glance, it certainly appears that the communication situations are symmetric. If anything, the internal observer is better off since it also has access to the control signals while the passive observer is denied access to them.

Suppose that the closed-loop process (1) had already been stabilized by the observer and controller system of 2.1, so that the second moment $E[X_t^2] \leq K$ for all t . Suppose that the noisy channel facing the passive encoder is the real $\frac{1}{2}$ -erasure channel of the previous section. It is interesting to consider how well the passive observer does at estimating this process.

The optimal encoding rule is clear, set $a_t = X_t$. It is certainly feasible to use $\hat{X}_t = B_t$ itself as the estimator for the process. This passive observation system clearly achieves $E[(\hat{X}_t - X_t)^2] \leq \frac{K}{2} < K$ since the probability of non-erasure is $\frac{1}{2}$. The causal decoding rule is able to achieve a finite end-to-end squared error distortion over this noisy channel in a causal and memoryless way.

This example makes it clear that the challenge here is arising from interactivity, not simply being forced to be delay-free. The passive external encoder and decoder do not have to face the unstable nature of the source while the internal observer and controller do. An error made while estimating \hat{X}_t by the passive decoder has no consequence for the next state X_{t+1} while a similar error by the controller does.

III. ANYTIME CAPACITY AND ITS NECESSITY

Anytime reliability is introduced and related to classical notions of reliability in [23]. Here, the focus is on the maximum rate achievable for a given sense of reliability rather than the maximum reliability possible at a given rate. The two are of course related since fundamentally there is an underlying region of feasible rate/reliability pairs.

Since the open-loop system state has the potential to grow exponentially, the controller's knowledge of the past must become certain at a fast rate in order to prevent a bad decision made in the past from continuing to corrupt the future. When viewed in the context of reliably communicating bits from an encoder to a decoder, this suggests that the estimates of the bits at the decoder must become increasingly reliable with time. The sense of anytime reliability is made precise in Section III-A. Section III-B then establishes the key result of this paper relating the problem of stabilization to the reliable communication of messages in the anytime sense. Finally, some consequences of this connection are studied in Section III-C. Among these consequences is a sequential generalization of the Schalkwijk/Kailath scheme for communication over an AWGN channel that achieves a doubly-exponential convergence to zero of the probability of bit error universally over all delays simultaneously.

A. Anytime reliability and capacity

The entire message is not assumed to be known ahead of time. Rather, it is made available gradually as time evolves. For simplicity of notation, let M_i be the R bit message that the channel encoder gets at time i . At the channel decoder, no target delay is assumed — *i.e.* the channel decoder does not necessarily know when the message i will be needed by the

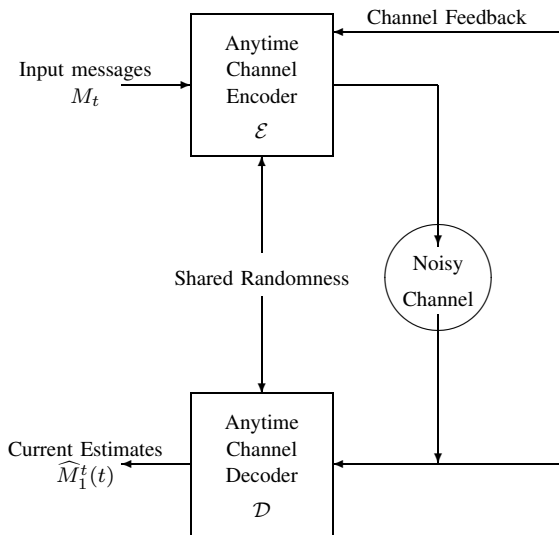


Fig. 4. The problem of communicating messages in an anytime fashion. Both the encoder \mathcal{E} and decoder \mathcal{D} are causal maps and the decoder in principle provides updated estimates for *all* past messages. These estimates must converge to the true message values appropriately rapidly with increasing delay.

application. A past message may even be needed more than once by the application. Consequently, the anytime decoder produces estimates $\widehat{M}_i(t)$ which are the best estimates for message i at time t based on all the channel outputs received so far. If the application is using the past messages with a delay d , the relevant probability of error is $\mathcal{P}(\widehat{M}_1^{t-d}(t) \neq M_1^{t-d})$. This corresponds to an uncorrected error anywhere in the distant past (*ie* on messages M_1, M_2, \dots, M_{t-d}) beyond d channel uses ago.

Definition 3.1: As illustrated in figure 4, a rate R communication system over a noisy channel is an encoder \mathcal{E} and decoder \mathcal{D} pair such that:

- R -bit message M_i enters¹⁴ the encoder at discrete time i
- The encoder produces a channel input at integer times based on all information that it has seen so far. For encoders with access to feedback with delay $1 + \theta$, this also includes the past channel outputs $B_1^{t-1-\theta}$.
- The decoder produces updated channel estimates $\widehat{M}_i(t)$ for all $i \leq t$ based on all channel outputs observed till time t .

A rate R sequential communication system achieves *anytime reliability* α if there exists a constant K such that:

$$\mathcal{P}(\widehat{M}_1^i(t) \neq M_1^i) \leq K2^{-\alpha(t-i)} \quad (2)$$

holds for every i, t . The probability is taken over the channel noise, the R bit messages M_i , and all of the common randomness available in the system.

If (2) holds for every possible realization of the messages M , then the system is said to achieve *uniform anytime reliability* α .

¹⁴In what follows, messages are considered to be composed of bits for simplicity of exposition. The i -th bit arrives at the encoder at time $\frac{i}{R}$ and thus M_i is composed of the bits $S_{\lfloor \frac{i}{R} \rfloor + 1}^{\lfloor \frac{i}{R} \rfloor + R}$.

Communication systems that achieve *anytime reliability* are called *anytime codes* and similarly for *uniform anytime codes*.

We could alternatively have bounded the probability of error by $2^{-\alpha(d - \log_2 K)}$ and interpreted $\log_2 K$ as the minimum delay imposed by the communication system.

Definition 3.2: The α -anytime capacity $C_{\text{any}}(\alpha)$ of a channel is the least upper bound of the rates R (in bits) at which the channel can be used to construct a rate R communication system that achieves uniform anytime reliability α .

Feedback anytime capacity is used to refer to the anytime capacity when the encoder has access to noiseless feedback of the channel outputs with unit delay.

The requirement for exponential decay in the probability of error with delay is reminiscent of the block-coding reliability functions $E(R)$ of a channel given in [1]. There is one crucial difference. With standard error exponents, both the encoder and decoder vary with blocklength or delay n . Here, the encoding is required to be fixed and the decoder in principle has to work at all delays since it must produce updated estimates of the message M_i at all times $t > i$.

This additional requirement is why it is called “anytime” capacity. The decoding process can be queried for a given bit at any time and the answer is required to be increasingly accurate the longer we wait. The anytime reliability α specifies the exponential rate at which the quality of the answers must improve. The anytime sense of reliable transmission lies between that represented by classical zero-error capacity C_0 (probability of error becomes zero at a large but finite delay) and classical capacity C (probability of error becomes something small at a large but finite delay). It is clear that $\forall \alpha, C_0 \leq C_{\text{any}}(\alpha) \leq C$.

By using a random coding argument over infinite tree codes, it is possible to show the existence of anytime codes without using feedback between the encoder and decoder for all rates less than the Shannon capacity. This shows:

$$C_{\text{any}}(E_r(R)) \geq R$$

where $E_r(R)$ is Gallager’s random coding error exponent calculated in base 2 and R is the rate in bits [33], [23]. Since feedback plays an essential role in control, it turns out that we are interested in the anytime capacity with feedback. It is interesting to note that in many cases for which the block-coding error exponents are not increased with feedback, the anytime reliabilities are increased considerably [6].

B. Necessity of anytime capacity

Anytime reliability and capacity are defined in terms of digital messages that must be reliably communicated from point to point. Stability is a notion involving the analog value of the state of a plant in interaction with a controller over a noisy feedback channel. At first glance, these two problems appear to have nothing in common except the noisy channel. Even on that point there is a difference. The observer/encoder \mathcal{O} in the control system may have no explicit access to the noisy output of the channel. It can appear to be using the noisy channel *without feedback*. Despite this, it turns out that

the relevant digital communication problem involves access to the noisy channel *with noiseless channel feedback* coming back to the message encoder.

Theorem 3.3: For a given noisy channel and $\eta > 0$, if there exists an observer \mathcal{O} and controller \mathcal{C} for the unstable scalar system that achieves $E[|X_t|^\eta] < K$ for all sequences of bounded driving noise $|W_t| \leq \frac{\Omega}{2}$, then the channel's feedback anytime capacity $C_{\text{any}}(\eta \log_2 \lambda) \geq \log_2 \lambda$ bits per channel use.

The proof of this spans the next few sections. Assume that there is an observer/controller pair $(\mathcal{O}, \mathcal{C})$ that can η -stabilize an unstable system with a particular λ and are robust to all bounded disturbances of size Ω . The goal is to use the pair to construct a rate $R < \log_2 \lambda$ anytime encoder and decoder for the channel with noiseless feedback, thereby reducing¹⁵ the problem of anytime communication to a problem of stabilization.

The heart of the construction is illustrated in figure 5. The “black-box” observer and controller are wrapped around a simulated plant mimicking (1). Since the $\{U_t\}$ must be generated by the black-box controller \mathcal{C} and the λ is prespecified, the disturbances $\{W_t\}$ must be used to carry the message. So, the encoder must embed the messages $\{M_t\}$ into an appropriate sequence $\{W_t\}$, taking care to stay within the Ω size limit.

While both the observer and controller can be simulated at the encoder thanks to the noiseless channel output feedback, at the decoder only the channel outputs are available. Consequently, these channel outputs are connected to a copy of the black-box controller \mathcal{C} , thereby giving access to the controls $\{U_t\}$ at the decoder. To extract the messages from these control signals, they are first causally preprocessed through a simulated copy of the unstable plant, except with no disturbance input. All past messages are then estimated from the current state of this simulated plant.

The key is to think of the simulated plant state as the sum of the states of two different unstable LTI systems. The first, with state denoted \tilde{X}_t , is driven entirely by the controls and starts in state 0.

$$\tilde{X}_{t+1} = \lambda \tilde{X}_t + U_t \quad (3)$$

\tilde{X} is available at both the decoder and the encoder due to the presence of noiseless feedback.¹⁶ The other, with state denoted \check{X}_t , is driven entirely by a simulated driving noise that is generated from the data stream to be communicated.

$$\check{X}_{t+1} = \lambda \check{X}_t + W_t \quad (4)$$

The sum $X_t = (\tilde{X}_t + \check{X}_t)$ behaves exactly like it was coming from (1) and is fed to the observer which uses it to generate inputs for the noisy channel.

The fact that the original observer/controller pair stabilized the original system implies that $|X_t| = |\check{X} - (-\tilde{X}_t)|$ is small and hence $-\tilde{X}_t$ stays close to \check{X}_t .

¹⁵In traditional rate-distortion theory, this “necessity” direction is shown by going through the mutual information characterizations of both the rate-distortion function and the channel capacity function. In the case of stabilization, mutual information is not discriminating enough and so the reduction of anytime reliable communication to stabilization must be done directly.

¹⁶If the controller is randomized, then the randomness is required to be common and shared between the encoder and decoder.

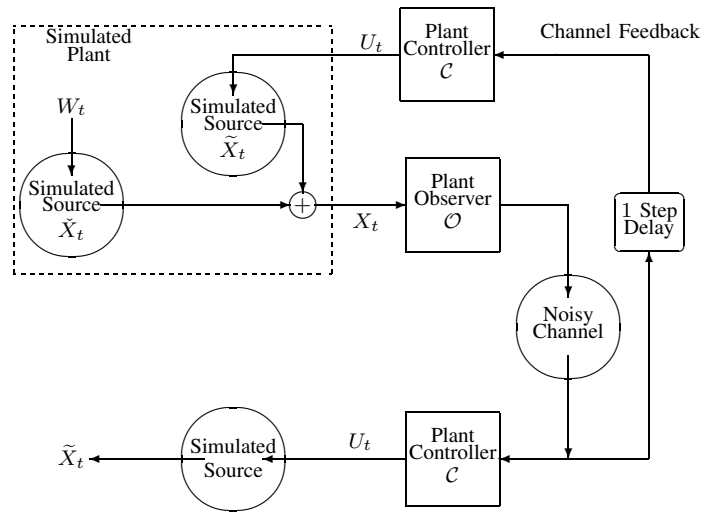


Fig. 5. The construction of a feedback anytime code from a control system. The messages are used to generate the $\{W_t\}$ inputs which are causally combined to generate $\{\check{X}_t\}$ within the encoder. The channel outputs are used to generate control signals at both the encoder and decoder. Since the simulated plant is stable, $-\check{X}$ and \tilde{X} are close to each other. The past message bits are estimated from the \tilde{X} at the decoder.

1) *Encoding data into the state:* As long as the bound Ω is satisfied, the encoder is free to choose any disturbance¹⁷ for the simulated plant. The choice will be determined by the data rate R and the specific messages to be sent. Rather than working with general messages M_i , consider a bitstream S_i with bit i becoming available at time $\frac{i}{R}$. Everything generalizes naturally to non-binary alphabets for the messages, but the notation is cleaner in the binary case with $S_i = \pm 1$.

\check{X}_t is the part of X_t driven only by the $\{W_t\}$.

$$\begin{aligned} \check{X}_t &= \lambda \check{X}_{t-1} + W_{t-1} \\ &= \sum_{i=0}^{t-1} \lambda^i W_{t-1-i} \\ &= \lambda^{t-1} \sum_{j=0}^{t-1} \lambda^{-j} W_j \end{aligned}$$

This looks like the representation of a fractional number in base λ which is then multiplied by λ^{t-1} . This is exploited in the encoding by choosing the bounded disturbance sequence

¹⁷In [23], a similar strategy is followed assuming a specific density for the iid disturbance W_t . In that context, it is important to choose a simulated disturbance sequence that behaves stochastically like W_t . This is accomplished by using common randomness shared between the encoder and decoder to dither the kind of disturbances produced here into ones with the desired density.

so that:¹⁸

$$\tilde{X}_t = \gamma \lambda^t \sum_{k=0}^{\lfloor Rt \rfloor} (2 + \epsilon_1)^{-k} S_k \quad (5)$$

where S_k is the k -th bit¹⁹ of data that the anytime encoder has to send and $\lfloor Rt \rfloor$ is just the total number of bits that are available by time t . γ, ϵ_1 are constants to be specified.

To see that (5) is always possible to achieve by appropriate choice of W , use induction. (5) clearly holds for $t = 0$. Now assume that it holds for time t and consider time $t + 1$:

$$\begin{aligned} \tilde{X}_{t+1} &= \lambda \tilde{X}_t + W_t \\ &= \gamma \lambda^{t+1} \left(\sum_{k=0}^{\lfloor Rt \rfloor} (2 + \epsilon_1)^{-k} S_k \right) + W_t \end{aligned}$$

So setting

$$W_t = \gamma \lambda^{t+1} \sum_{k=\lfloor Rt \rfloor + 1}^{\lfloor R(t+1) \rfloor} (2 + \epsilon_1)^{-k} S_k \quad (6)$$

gives the desired result. Manipulate (6) to get $W_t =$

$$\begin{aligned} &\gamma \lambda^{t+1} (2 + \epsilon_1)^{-\lfloor Rt \rfloor} \sum_{j=1}^{\lfloor R(t+1) \rfloor - \lfloor Rt \rfloor} (2 + \epsilon_1)^{-j} S_{\lfloor Rt \rfloor + j} \\ &= \gamma \lambda \frac{(2 + \epsilon_1)^{Rt - (\lfloor Rt \rfloor)}}{\lambda^{-t(1 - R \frac{\log_2(2 + \epsilon_1)}{\log_2 \lambda})}} \sum_{j=1}^{\lfloor R(t+1) \rfloor - \lfloor Rt \rfloor} (2 + \epsilon_1)^{-j} S_{\lfloor Rt \rfloor + j} \end{aligned}$$

To keep this bounded, choose

$$\epsilon_1 = 2^{\frac{\log_2 \lambda}{R}} - 2 \quad (7)$$

which is strictly positive if $R < \log_2 \lambda$. Applying that substitution gives $|W_t| =$

$$\begin{aligned} &|\gamma \lambda (2 + \epsilon_1)^{Rt - (\lfloor Rt \rfloor)} \sum_{j=1}^{\lfloor R(t+1) \rfloor - \lfloor Rt \rfloor} (2 + \epsilon_1)^{-j} S_{\lfloor Rt \rfloor + j}| \\ &< |\gamma \lambda (2 + \epsilon_1)| \\ &= |\gamma \lambda^{1 + \frac{1}{R}}| \end{aligned}$$

So by choosing

$$\gamma = \frac{\Omega}{2\lambda^{1 + \frac{1}{R}}} \quad (8)$$

the simulated disturbance is guaranteed to stay within the specified bounds.

¹⁸For a rough understanding, ignore the ϵ_1 and suppose that the message were encoded in binary. It is intuitive that any good estimate of the \tilde{X}_t state is going to agree with \tilde{X}_t in all the high order bits. Since the system is unstable, all the encoded bits eventually become high-order bits as time goes on. So no bit error could persist for too long and still keep the estimate close to \tilde{X}_t . The ϵ_1 in the encoding is a technical device to make this reasoning hold uniformly for all bit strings, rather than merely "typical" ones. This is important since we are aiming for exponentially small bounds and so cannot neglect rare events.

¹⁹For the next section, it is convenient to have the disturbances balanced around zero and so we choose to represent the bit S_i as $+1$ or -1 rather than the usual 1 or 0 .

2) *Extracting data bits from the state estimate:*

Lemma 3.1: Given a channel with access to noiseless feedback, for any rate $R < \log_2 \lambda$, it is possible to encode bits into the simulated scalar plant so that the uncontrolled process behaves like (5) by using disturbances given in (6) and the formulas (7) and (8). At the output end of the noisy channel, it is possible to extract estimates $\hat{S}_i(t)$ for the i -th bit sent for which the error event

$$\{\omega | \exists i \leq j, \hat{S}_i(t) \neq S_i(t)\} \subseteq \{\omega | |X_t| \geq \lambda^{t - \frac{j}{R}} \left(\frac{\gamma \epsilon_1}{1 + \epsilon_1} \right)\} \quad (9)$$

and thus:

$$\mathcal{P}(\hat{S}_1^j(t) \neq S_1^j(t)) \leq \mathcal{P}(|X_t| \geq \lambda^{t - \frac{j}{R}} \left(\frac{\gamma \epsilon_1}{1 + \epsilon_1} \right)) \quad (10)$$

Proof: Here ω is used to denote members of the underlying sample space.²⁰

The decoder has $-\tilde{X}_t = \tilde{X}_t - X_t$ which is close to \tilde{X} since X_t is small. To see how to extract bits from $-\tilde{X}_t$, first consider how to recursively extract those bits from \tilde{X}_t .

Starting with the first bit, notice that the set of all possible \tilde{X}_t that have $S_0 = +1$ is separated from the set of all possible \tilde{X}_t that have $S_0 = -1$ by a gap of

$$\begin{aligned} &\gamma \lambda^t \left(\left(1 - \sum_{k=1}^{\lfloor Rt \rfloor} (2 + \epsilon_1)^{-k} \right) - \left(-1 + \sum_{k=1}^{\lfloor Rt \rfloor} (2 + \epsilon_1)^{-k} \right) \right) \\ &> \gamma \lambda^t 2 \left(1 - \sum_{k=1}^{\infty} (2 + \epsilon_1)^{-k} \right) \\ &= \gamma \lambda^t 2 \left(1 - \frac{1}{1 + \epsilon_1} \right) \\ &= \lambda^t \left(\frac{2\epsilon_1 \gamma}{1 + \epsilon_1} \right) \end{aligned}$$




Fig. 6. The data bits are used to sequentially refine a point on a Cantor set. Its natural tree structure allows bits to be encoded sequentially. The Cantor set also has finite gaps between all points corresponding to bit sequences that first differ in a particular bit position. These gaps allow the uniformly reliable extraction of bit values from noisy observations.

Notice that this worst-case gap²¹ is a positive number that is growing exponentially in t . If the first $i - 1$ bits are the same, then both sides can be scaled by $(2 + \epsilon_1)^i = \lambda^{\frac{i}{R}}$ to get the same expressions above and so by induction, it quickly follows that the minimum gap between the encoded state corresponding to two sequences of bits that first differ in bit position i is given by $\text{gap}_i(t) =$

$$\inf_{\bar{S}: \bar{S}_i \neq S_i} |\tilde{X}_t(S) - \tilde{X}_t(\bar{S})| > \begin{cases} \lambda^{t - \frac{i}{R}} \left(\frac{2\gamma \epsilon_1}{1 + \epsilon_1} \right) & \text{if } i \leq \lfloor Rt \rfloor \\ 0 & \text{otherwise} \end{cases} \quad (11)$$

²⁰If the bits to be sent are deterministic, this is the sample space giving channel noise realizations.

²¹The typical gap is larger and so the probability of error is actually lower than this bound says it is.

Because the gaps are all positive, (11) shows that it is always possible to perfectly extract the data bits from \tilde{X}_t by using an iterative procedure.²² To extract bit information from an input I_t :

- 1) Initialize threshold $T_0 = 0$ and counter $i = 0$.
- 2) Compare input I_t to T_i . If $I_t \geq T_i$, set $\hat{S}_i(t) = +1$. If $I_t < T_i$, set $\hat{S}_i(t) = -1$.
- 3) Increment counter i and update threshold $T_i = \gamma \lambda^t \sum_{k=0}^{i-1} (2 + \epsilon_1)^{-k} \hat{S}_k$
- 4) Goto step 2 as long as $i \leq \lfloor Rt \rfloor$

Since the gaps given by (11) are always positive, the procedure works perfectly if applied to input $I_t = \tilde{X}_t$. At the decoder, apply the procedure to $I_t = -\tilde{X}_t$ instead.

With this, (9) is easy to verify by looking at the complementary event $\{\omega \mid |X_t| < \frac{\lambda^{t-\frac{i}{R}} \gamma \epsilon_1}{1 + \epsilon_1}\}$. The bound (11) thus implies that we are less than halfway across the minimum gap for bit j at time t . Consequently, there is no error in the step 2 comparison of the procedure at iterations $i \leq j$. \square

3) *Probability of error for bounded moment and other senses of stability: Proof of Theorem 3.3:* Using Markov's inequality:

$$\begin{aligned} \mathcal{P}(|X_t| > m) &= \mathcal{P}(|X_t|^\eta > m^\eta) \\ &\leq E[|X_t|^\eta] m^{-\eta} \\ &< K m^{-\eta} \end{aligned}$$

Combining with Lemma 3.1, gives:

$$\begin{aligned} \mathcal{P}(\hat{S}_1^i(t) \neq S_1^i(t)) &\leq \mathcal{P}(|X_t| \geq \lambda^{t-\frac{i}{R}} \left(\frac{\gamma \epsilon_1}{1 + \epsilon_1} \right)) \\ &< K \left(\frac{1}{\gamma} + \frac{1}{\gamma \epsilon_1} \right)^\eta \lambda^{-\eta(t-\frac{i}{R})} \\ &= \left(K \left(\frac{1}{\gamma} + \frac{1}{\gamma \epsilon_1} \right)^\eta \right) 2^{-(\eta \log_2 \lambda)(t-\frac{i}{R})} \end{aligned}$$

Since $t - \frac{i}{R}$ represents the delay between the time that bit i was ready to be sent and the decoding time, the theorem is proved. \square

All that was needed from the bounded moment sense of stability was some bound on the probability that X_t took on large values. Thus, the proof above immediately generalizes to other senses of stochastic stability if we suitably generalize the sense of anytime capacity to allow for other bounds on the probability of error with delay.

Definition 3.4: A rate R communication system achieves g -anytime reliability given by a function $g(d)$ if

$$\mathcal{P}(\hat{M}_1^{t-d}(t) \neq M_1^{t-d}(t)) < g(d)$$

$g(d)$ is assumed to be 1 for all negative values of d .

The g -anytime capacity $C_{g\text{-any}}(g)$ of a noisy channel is the least upper bound of the rates R at which the channel can be used to construct a sequential communication system that achieves g -anytime reliability given by the function $g(d)$.

Notice that for α -anytime capacity, $g(d) = K 2^{-\alpha d}$ for some K .

Theorem 3.5: For a given noisy channel and decreasing function $f(m)$, if there exists an observer \mathcal{O} and controller \mathcal{C} for the unstable scalar system that achieves $\mathcal{P}(|X_t| > m) < f(m)$ for all sequences of bounded driving noise $|W_t| \leq \frac{\Omega}{2}$, then $C_{g\text{-any}}(g) \geq \log_2 \lambda$ for the noisy channel considered with the encoder having access to noiseless feedback and $g(d)$ having the form $g(d) = f(K \lambda^d)$ for some constant K .

Proof: For any rate $R < \log_2 \lambda$,

$$\begin{aligned} \mathcal{P}(\hat{S}_1^i(t) \neq S_1^i(t)) &\leq \mathcal{P}(|X_t| \geq \frac{\lambda^{t-\frac{i}{R}} \gamma \epsilon_1}{1 + \epsilon_1}) \\ &= f\left(\frac{\gamma \epsilon_1}{1 + \epsilon_1} \lambda^{t-\frac{i}{R}}\right) \end{aligned}$$

Since the delay $d = t - \frac{i}{R}$, the theorem is proved. \square

C. Implications

At this point, it is interesting to consider a few implications of Theorem 3.5.

1) *Weaker senses of stability than η -moment:* There are senses of stability weaker than specifying a specific η -th moment or a specific tail decay target $f(m)$. An example is given by the requirement $\lim_{m \rightarrow \infty} \mathcal{P}(|X_t| > m) = 0$ uniformly for all t . This can be explored by taking the limit of $C_{\text{any}}(\alpha)$ as $\alpha \downarrow 0$. We have shown elsewhere[33], [23] that:

$$\lim_{\alpha \downarrow 0} C_{\text{any}}(\alpha) = C$$

where C is the Shannon classical capacity. This holds for all discrete memoryless channels since the α -anytime reliability goes to zero at Shannon capacity but is > 0 for all lower rates even without feedback being available at the encoder. Thus, classical Shannon capacity is the natural candidate for the relevant figure of merit.

To see why Shannon capacity can not be beaten, it is useful to consider an even more lax sense of stability. Suppose the requirement were only that $\lim_{m \rightarrow \infty} \mathcal{P}(|X_t| > m) = 10^{-5} > 0$ uniformly for all t . This imposes the constraint that the probability of a large state stays below 10^{-5} for all time. Theorem 3.5 would thus only require the probability of decoding error to be less than 10^{-5} . However, Wolfowitz' strong converse to the coding theorem[1] implies that since the block-length in this case is effectively going to infinity, the Shannon capacity of the noisy channel still must satisfy $C \geq \log_2 \lambda$. Adding a finite tolerance for unboundedly large states does not get around the need to be able to communicate $\log_2 \lambda$ bits reliably.

2) *Stronger senses of stability than η -moment:* Having f decrease only as a power law might not be suitable for certain applications. Unfortunately, this is all that can be hoped for in generic situations. Consider a DMC with no zero entries in its transition matrix. Define $\rho = \min_{i,j} p(i,j)$. For such a channel, with or without feedback, the probability of error after d time steps is lower bounded by ρ^d since that lower bounds the probability of all channel output sequences of length d . This implies that the probability of error can drop no more than exponentially in d for such DMCs. Tighter upper-bounds

²²This is a minor twist on the procedure followed by serial A/D converters.

on anytime reliability with feedback are available in [34] and [6].

Theorem 3.5 therefore implies that the only f -senses of stability which are possible over such channels are those for which:

$$\begin{aligned} f(K\lambda^d) &\geq \rho^d \\ f(m) &\geq \rho^{\frac{\log_2(\frac{m}{K})}{\log_2 \lambda}} \\ f(m) &\geq K'm^{-\frac{\log_2 \frac{1}{\rho}}{\log_2 \lambda}} \end{aligned}$$

which is a power law. This rules out the ‘‘risk sensitive’’ sense of stability in which f is required to decrease exponentially. In the context of Theorem 3.3, this also implies that there is an η beyond which all moments must be infinite!

Corollary 3.1: If any unstable process is controlled over a discrete memoryless channel with no feedback zero-error capacity, then the resulting state can have at best a power-law bound (Pareto distribution) on its tail.

This is very much related to how sequential decoding must have computational effort distributions with at best a Pareto distribution[35]. In both cases, the result follows from the interaction of two exponentials. The difference is that the computational search effort distributions assumed a particular structure on the decoding algorithm while the bound here is fundamental to the stabilization problem regardless of the observers or controllers.

Thus for DMCs and a given λ , we are either limited to a power-law tail for the controlled state because of an anytime reliability that is at most singly exponential in delay or it is possible to hold the state inside a finite box since there is adequate feedback zero-error capacity. Nothing in between can happen with a DMC.

3) *Limiting the controller effort or memory:* If there was a hard limit on actuator effort ($|U| \leq \mathcal{U}$ for some $\mathcal{U} > 0$), then the only way to maintain stability is to also have a hard limit on how big the state X can get. Theorem 3.5 immediately gives a fundamental requirement for feedback zero-error capacity $\geq \log_2 \lambda$ since $g(d) = 0$ for sufficiently large d .

Similarly, consider limited-memory time-invariant controllers which only have access to the past k channel outputs. If the channel has a finite output alphabet and no randomization is permitted at the controller, limited memory immediately translates into only a finite number of possible control inputs. Since there must be a largest one, it reduces to the case of having a hard limit on actuator effort.

We conjecture that even with randomization and time-variation, finite memory at the controller implies that the channel must have feedback zero-error capacity $\geq \log_2 \lambda$. Intuitively, if the channel has zero-error capacity $< \log_2 \lambda$, it can misbehave for arbitrarily long times and build up a huge ‘‘backlog’’ of uncertainty that can not be resolved at the controller. With finite memory, the controller has no way of knowing what uncertainty it is actually facing and so is unable to properly interpret the channel outputs to devise the proper control signals.

4) *The AWGN case with an average input power constraint:* The tight relationship between control and communication

established in Theorem 3.5 allows the construction of sequential codes for noisy channels with noiseless feedback if we know how to stabilize linear plants over such channels. Consider the problem of stabilizing an unstable plant driven by finite variance driving noise over an AWGN channel. A linear observer and controller strategy achieve mean-square stability for such systems since the problem fits into the standard LQG framework [14].

By looking more closely at the actual tail probabilities achieved by the linear observer/controller strategy, we obtain a natural anytime generalization of Schalkwijk and Kailath’s scheme[36], [37] for communicating over the power constrained additive white Gaussian noise channel with noiseless feedback. Its properties are summarized in Figure 7, but the highlight is that it achieves doubly exponential reliability with delay, universally over all sufficiently long delays.

Theorem 3.6: It is possible to communicate bits reliably across a discrete-time average-power constrained AWGN channel with noiseless feedback at any rate $R < \frac{1}{2} \log_2(1 + \frac{P}{\sigma^2})$ while achieving a g -anytime reliability of at least

$$g(d) = 2e^{-K(4^{Rd} - O(2^{Rd}))} \quad (12)$$

for some constant K that depends only on the rate R , power constraint P , and channel noise power σ^2 .

Proof: To avoid having to drag σ^2 around, just normalize units so as to consider power constraint $P' = \frac{P}{\sigma^2}$ and a channel with iid unit variance noise N_t . Choose the λ for the simulated (1) so that $R < \log_2 \lambda < \frac{1}{2} \log_2(1 + P')$.

The observer/encoder used is a linear map:

$$a_t = \beta X_t \quad (13)$$

so the channel output $B_t = \beta X_t + N_t$. Use a linear controller:

$$U_t = -\lambda \phi B_t \quad (14)$$

giving the closed-loop system:

$$X_{t+1} = \lambda(1 - \beta\phi)X_t + W_t - \lambda\phi N_t \quad (15)$$

where the β, ϕ are constants to be chosen. For the closed-loop system to be stable:

$$0 < \lambda(1 - \beta\phi) < 1 \quad (16)$$

Thus $\beta\phi \in (1 - \frac{1}{\lambda}, 1)$. Assuming (16) holds and temporarily setting the $W_t = 0$ for analysis, it is clear that the closed-loop X_t is Gaussian with a growing variance asymptotically tending to

$$\sigma_x^2 = \frac{\lambda^2 \phi^2}{1 - \lambda^2(1 - \beta\phi)^2} \quad (17)$$

The channel input power satisfies:

$$E[a_t^2] \leq \frac{\lambda^2(\beta\phi)^2}{1 - \lambda^2(1 - \beta\phi)^2}$$

Since $\lambda^2 < 1 + P'$, define $P'' = \lambda^2 - 1 < P'$ and substitute to get:

$$E[a_t^2] \leq \frac{(P'' + 1)(\beta\phi)^2}{1 - (P'' + 1)(1 - \beta\phi)^2} \quad (18)$$

By setting $\beta\phi = \frac{P''}{P''+1}$, the left hand side of (18) is identically P'' as desired. All that remains is to verify the stability condition (16):

$$\begin{aligned} \lambda(1 - \beta\phi) &= \frac{\lambda}{P'' + 1} \\ &= \frac{\sqrt{P'' + 1}}{P'' + 1} \\ &= \frac{1}{\sqrt{P'' + 1}} \\ &< 1 \end{aligned}$$

So the closed loop system is stable and the channel noise alone results in an average input power of at most $P'' < P'$.

Rather than optimizing the choice of β and ϕ to get the best tradeoff point, just set $\beta = 1$ and $\phi = \frac{P''}{P''+1}$ for simplicity. In that case, $\sigma_x^2 = P''$.

Now consider the impact of the W_t alone on the closed-loop control system. These are going through a stable system and so by expanding the recursion (15) and setting $N_t = 0$,

$$\begin{aligned} |X_t^w| &\leq \sum_{i=0}^{\infty} (\lambda(1 - \beta\phi))^i \frac{\Omega}{2} \\ &= \sum_{i=0}^{\infty} \left(\frac{1}{\sqrt{P'' + 1}} \right)^i \frac{\Omega}{2} \\ &= \frac{\Omega}{2(1 - \frac{1}{\sqrt{P'' + 1}})} \\ &= \frac{\Omega\sqrt{P'' + 1}}{2(\sqrt{P'' + 1} - 1)} \end{aligned}$$

which is a constant that can be made as small as desired by choice of Ω . Assume that the data stream S to be transmitted is independent of the channel noise N . Then, the total average input power is bounded by:

$$\begin{aligned} \sigma_x^2 + \beta^2 |X_t^w|^2 &\leq P'' + \left(\frac{\Omega\sqrt{P'' + 1}}{2(\sqrt{P'' + 1} - 1)} \right)^2 \\ &\leq P'' + \Omega^2 \frac{P'' + 1}{4(P'' + 2(1 - \sqrt{P'' + 1}))} \end{aligned}$$

Since $P'' < P'$, we can choose an Ω small enough so that the channel input satisfies the average power constraint regardless of the message bits to be sent.

All that remains is to see what $f(m)$ this control system meets for such arbitrary, but bounded, disturbances. X_t is asymptotically the sum of a Gaussian with zero mean and variance P'' together with the closed-loop impact of the disturbance $X^w(t)$. Since the total impact of the disturbance part is bounded:

$$\begin{aligned} \mathcal{P}(|X_t| > m) &\leq \mathcal{P}(|N_{\sigma_x^2}| > m - \frac{\Omega\sqrt{P'' + 1}}{2(\sqrt{P'' + 1} - 1)}) \\ &= \mathcal{P}(|N| > \frac{1}{\sqrt{P''}}(m - \frac{\Omega\sqrt{P'' + 1}}{2(\sqrt{P'' + 1} - 1)})) \\ &\leq 2e^{-\frac{1}{2P''}(m - \frac{\Omega\sqrt{P'' + 1}}{2(\sqrt{P'' + 1} - 1)})^2} \end{aligned}$$

Ignoring the details of the constants, this gives an $f(m) = 2e^{-K_1(m-K_2)^2} = 2e^{-K_1(m^2-2K_2m-K_3)}$. Applying Theo-

Scheme:	Schalkwijk [37]	Theorem 3.6
Message:	known in advance	streams in
Delay:	prespecified	universal
Error exponent:	double-exponential	double-exponential
Channel:	known AWGN	known AWGN
Constraint:	average power	average power
Noiseless Feedback:	required	required
Initialization:	2^{nR} -PAM + ML feedback	none
Ongoing:	MMSE feedback	MMSE feedback + small R -PAM perturbations
Channel input:	Gaussian	Perturbed Gaussian
Decoding:	Minimum distance	Minimum distance
Equivalent Plant:	unstable	unstable
	$R < \log_2 \lambda < C$	$R < \log_2 \lambda < C$
Initial condition:	bounded	zero
Disturbance:	zero	bounded
Stability sense:	almost-sure [18]	exponential tail

Fig. 7. Quick comparison of the Schalkwijk/Kailath scheme to the anytime generalization in this paper.

rem 3.5 immediately gives (12) since $\lambda^d > 2^{Rd}$. \square

Since the convergence is double exponential, it is faster than any exponential and hence

$$C_{\text{any}}(\alpha) = \frac{1}{2} \log_2(1 + \frac{P}{\sigma^2})$$

for all $\alpha > 0$ on the AWGN channel. If the additive channel noise were not Gaussian, but had bounded support with the same variance, then this proof immediately reveals that the zero-error capacity of such a bounded noise channel with feedback satisfies: $C_0 \geq \frac{1}{2} \log_2(1 + \frac{P}{\sigma^2})$.

In the Gaussian case, it is not immediately clear whether there are ideas analogous to those in [38] that can be used to further boost the g -anytime reliability beyond double exponential. It is clear that if it were possible, it would require nonlinear control strategies.

The AWGN case is merely one example. Theorem 3.5 gives a way to lower-bound the anytime capacity for channels with feedback in cases where the optimal control behavior is easy to see. The finite moments of the closed-loop state reveal what anytime reliability is being achieved. Often, there is a simple upper-bound that matches up with the lower-bound thereby giving the anytime capacity itself. The BEC case discussed in [16], [33], [6] is such an example. In addition, Theorem 3.5 gives us the ability to mix and match communication and control tools to study a problem. This is exploited in [30], [39] to understand the feedback anytime capacity of constrained packet erasure channels and the power constrained AWGN+erasure channel. In [40], these results are extended to the Gilbert-Eliot channel with feedback. It is also exploited in [34] to lower bound the anytime reliability achieved by a particular code for the BSC with feedback.

IV. THE SUFFICIENCY OF ANYTIME CAPACITY

A. Overview

When characterizing a noisy channel for control, the choice of information pattern[41] can be critical [14]. The sufficiency result is first established for cases with an explicit noiseless feedback path from the channel outputs back to the observer. Section IV-E takes a quick look at the simpler problem of almost-sure stabilization when the system is undisturbed and all the uncertainty comes from either the channel or the initial condition. Then, in Section IV-F, the impact of viewing time in blocks of size n and only acting on the slower time-scale is examined. Finally, Sections IV-G and IV-H give models for boundedly noisy or quantized controls and/or observations and show that such bounded noise can be tolerated.

To prove the sufficiency theorem addressing the situation illustrated in figure 2, we need to design an observer/controller pair that deals with the analog plant and communicates across the channel by using an anytime communication system. The anytime communication system works with noiseless feedback from the channel output available at the bit encoder and is considered a “black box.”

Theorem 4.1: For a given noisy channel, if there exists an anytime encoder/decoder pair with access to noiseless feedback that achieves $C_{g\text{-any}}(g) \geq \log_2 \lambda$, then it is possible to stabilize an unstable scalar plant with parameter λ that is driven by bounded driving noise through the noisy channel by using an observer that has noiseless access to the noisy channel outputs. Furthermore, there exists a constant K so that $\mathcal{P}(|X_t| > m) \leq g(K + \log_\lambda m)$.

To prove this theorem, explicit constructions are given for the observer and controller in the next sections.

B. Observer

Since the observer has access to the channel outputs, it can run a copy of the controller and hence has access to the control signals U_t . Since $W_t = X_{t+1} - \lambda X_t - U_t$, and the observer receives X_t from the plant, the observer also effectively has access to the W_t . However, it is not sufficient to merely encode the W_t independently to some precision.²³ Instead, the observer will act as though it is working with a virtual controller through a noiseless channel of finite rate R in the manner of example 2.1. The resulting bits will be sent through the anytime code.

The observer is constructed to keep the state uncertainty at the virtual controller inside a box of size Δ by using bits at the rate R . It does this by simulating a virtual process \bar{X}_t governed by:

$$\bar{X}_{t+1} = \lambda \bar{X}_t + W_t + \bar{U}_t \quad (19)$$

where the \bar{U}_t represent the computed actions of the virtual controller. This gives rise to a virtual counterpart of \tilde{X}_t

$$X_{t+1}^{\bar{U}} = \lambda X_t^{\bar{U}} + \bar{U}_t \quad (20)$$

²³This is because the unstable plant will eventually blow up even tiny uncorrected discrepancies between the encoded and actual W_t .

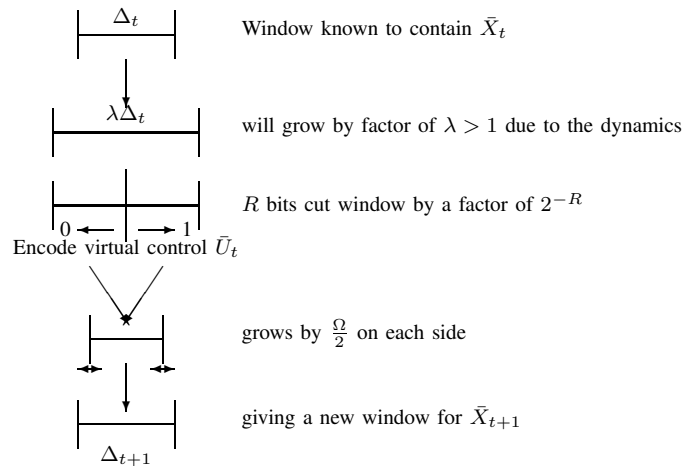


Fig. 8. Virtual controller for $R=1$. How the virtual state \bar{X} evolves.

which satisfies the relationship $\bar{X}_t = \tilde{X}_t + X_t^{\bar{U}}$. Because \bar{X}_t will be kept within a box, it is known that $-X_t^{\bar{U}}$ is close to \tilde{X}_t . The actual controller will pick controls designed to keep \tilde{X}_t close to $X_t^{\bar{U}}$.

Because of the rate constraint, the virtual control \bar{U}_t takes on one of $2^{\lfloor R(t+1) \rfloor - \lfloor Rt \rfloor}$ values. For simplicity of exposition, we ignore the integer effects and consider it to be one of 2^R values²⁴ and proceed by induction. Assume that \bar{X}_t is known to lie within $[-\frac{\Delta}{2}, \frac{\Delta}{2}]$. Then $\lambda \bar{X}_t$ will lie within $[-\frac{\lambda \Delta}{2}, \frac{\lambda \Delta}{2}]$. By choosing 2^R control values uniformly spaced within that interval, it is guaranteed that $\lambda \bar{X}_t + \bar{U}_t$ will lie within $[-\frac{\lambda \Delta}{2^{R+1}}, \frac{\lambda \Delta}{2^{R+1}}]$. Finally, the state will be disturbed by W_t and so \bar{X}_{t+1} will be known to lie within $[-\frac{\lambda \Delta}{2^{R+1}} - \frac{\Omega}{2}, \frac{\lambda \Delta}{2^{R+1}} + \frac{\Omega}{2}]$.

Since the initial condition has no uncertainty, induction will be complete if

$$\frac{\lambda}{2^R} \Delta + \Omega \leq \Delta \quad (21)$$

To get the minimum Δ required as a function of R , we can solve for (21) being an equality. This occurs²⁵ when $\Delta = \frac{\Omega}{1 - \lambda 2^{-R}}$ for every case where $R > \log_2 \lambda$. Since the slope $\frac{\lambda}{2^R}$ on the left hand side of (21) is less than 1, any larger Δ also works.

Since they arose from dividing the uncertainty window to 2^R disjoint segments, it is clear that the virtual controls \bar{U}_t can be encoded causally using R bits per unit time. These bits are sent to the anytime encoder for transport over the noisy channel.

C. Controller

The controller uses the updated bit estimates from the anytime decoder to choose a control to attempt to make the true state X_t stay close to the virtual state \bar{X}_t . It does this by having a pair of internal models as shown in figure 9.

The first, \tilde{X}_t from (3), models the unstable system driven only by the actual controls. The second is its best estimate \hat{X}_t ,

²⁴For the details of how to deal with fractional R , please see the causal source code discussion in [33].

²⁵In reality, the uncertainty approaches this from below since the system starts at the known initial condition 0.

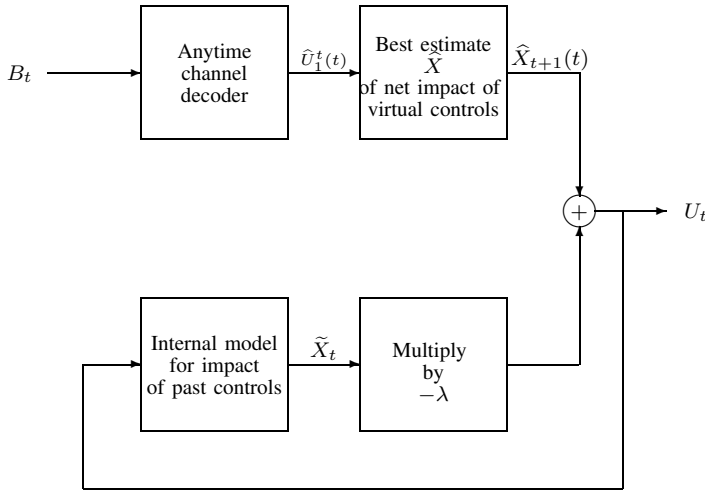


Fig. 9. The controller remembers what it did in the past and uses the anytime decoder to get an updated sense of where the observer wants it to go. It then applies a control designed to correct for any past errors and move the state to be close to the virtual state controlled by the observer.

based on the current bit estimates from the anytime decoder, of where the unstable system should be driven only by the virtual controls \bar{U}_t . Of course, the controller does not have the exact virtual controls, only its best estimates $\hat{U}_1^t(t)$ for them.

$$\hat{X}_{t+1}(t) = \sum_{i=0}^t \lambda^i \hat{U}_{t-i}(t) \quad (22)$$

This is not given in recursive form since all of the past estimates for the virtual controls are subject to re-estimation at the current time t . The control U_t is chosen to make $\tilde{X}_{t+1} = \hat{X}_{t+1}(t)$.

$$U_t = \hat{X}_{t+1}(t) - \lambda \tilde{X}_t \quad (23)$$

D. Evaluating stability

Proof of Theorem 4.1: With controls given by (23), the true state X_t can be written as:

$$\begin{aligned} X_t &= \tilde{X}_t + \tilde{X}_t = \tilde{X}_t + \hat{X}_t(t-1) \\ &= \sum_{i=0}^{t-1} \lambda^i (W_{t-i} + \hat{U}_{t-i}(t-1)) \end{aligned}$$

Notice that the actual state X_t differs from the virtual state \bar{X}_t only due to errors in virtual control estimation due to channel noise. If there were no errors in the prefix \hat{U}_1^{t-d} and arbitrarily bad errors for \hat{U}_{t-d+1}^t , then we could start at \bar{X}_{t-d} and see how much the errors could have propagated since then:

$$X_t = \lambda^d \bar{X}_{t-d} + \sum_{i=0}^{d-1} \lambda^i (W_{t-i} + \hat{U}_{t-i}(t-1))$$

Comparing this with \bar{X}_t , and noticing that the maximum

possible difference between two virtual controls is $\lambda\Delta$ gives:

$$\begin{aligned} |X_t - \bar{X}_t| &= \left| \sum_{i=0}^{d-1} \lambda^i (\bar{U}_{t-i} - \hat{U}_{t-i}(t-1)) \right| \\ &\leq \sum_{i=0}^{d-1} \lambda^i |\bar{U}_{t-i} - \hat{U}_{t-i}(t-1)| \\ &\leq \sum_{i=0}^{d-1} \lambda^{i+1} \Delta \\ &< \Delta \lambda^d \sum_{i=0}^{\infty} \lambda^{-i} \\ &= \lambda^d \frac{\Delta}{1 - \lambda^{-1}} \end{aligned}$$

Since $|\bar{X}_t| \leq \frac{\Delta}{2}$, if we know that there were no errors in the prefix of estimated virtual controls until d time steps ago, then

$$\{\hat{U}_0^{t-d}(t-1) = \bar{U}_0^{t-d}\} \Rightarrow \{|X_t| < \lambda^d \frac{2\Delta}{1 - \lambda^{-1}}\} \quad (24)$$

(24) immediately gives:

$$\begin{aligned} &\mathcal{P}(|X_t| \geq m) \\ &= \mathcal{P}(|X_t| \geq \lambda^{\frac{\log_2 m}{\log_2 \lambda}} \lambda^{\frac{\log_2(1-\lambda^{-1}) - \log_2(2\Delta)}{\log_2 \lambda}} \left(\frac{2\Delta}{1 - \lambda^{-1}} \right)) \\ &\leq \mathcal{P}(|X_t| \geq \lambda^{\lfloor \frac{\log_2 m + \log_2(1-\lambda^{-1}) - \log_2(2\Delta)}{\log_2 \lambda} \rfloor}) \left(\frac{2\Delta}{1 - \lambda^{-1}} \right) \\ &\leq g\left(\frac{\log_2 m + \log_2(1 - \lambda^{-1}) - \log_2(2\Delta)}{\log_2 \lambda} \right) \\ &\leq g\left(K'' + \frac{\log_2 m}{\log_2 \lambda} \right) \end{aligned}$$

where g bounds the probability of error for the g -anytime code and K'' is some constant. \square

Specializing to the case of α -anytime capacity, it is clear that:

$$\begin{aligned} \mathcal{P}(|X_t| \geq m) &\leq K''' 2^{-\alpha \frac{\log_2 m}{\log_2 \lambda}} \\ &= K''' m^{-\frac{\alpha}{\log_2 \lambda}} \end{aligned}$$

which gives a power-law bound on the tail. If the goal is a finite η -th moment,

$$\begin{aligned} E[|X_t|^\eta] &= \int_0^\infty \mathcal{P}(|X_t|^\eta \geq m) dm \\ &= \int_0^\infty \mathcal{P}(|X_t| \geq m^{\frac{1}{\eta}}) dm \\ &\leq 1 + K''' \int_1^\infty m^{-\frac{\alpha}{\eta \log_2 \lambda}} dm \end{aligned}$$

As long as $\alpha > \eta \log_2 \lambda$, the integral above converges and hence the controlled process has a bounded η -moment.

Theorem 4.2: It is possible to control an unstable scalar process driven by a bounded disturbance over a noisy channel so that the η -moment of $|X_t|$ stays finite for all time if the channel has feedback anytime capacity $C_{\text{any}}(\alpha) > \log_2 \lambda$ for some $\alpha > \eta \log_2 \lambda$ and the observer is allowed to observe the noisy channel outputs and the state exactly.

Aside from the usual gap between $>$ and \geq , this shows that the necessity condition in Theorem 3.3 is tight. Since there are no assumptions on the disturbance process except for its boundedness, the sufficiency theorems here automatically cover the case of stochastic disturbances having any sort of memory structure as long as they remain bounded in support.

E. Almost-sure stability

Control theorists are sometimes interested in an even simpler problem for which there is no disturbance (i.e. $W_t = 0$ for all t) but the initial condition X_0 is unknown to within some bound Ω . For this problem, the goal is ensuring that the state X_t tends to zero almost surely. This short section constructively shows that any sufficiency result for η -stability also extends to almost-sure stabilization. To do this, we consider the system:

$$X'_{t+1} = \lambda' X'_t + U'_t + W'_t \quad (25)$$

and use it to prove a key lemma:

Lemma 4.1: If it is possible to η' -stabilize a persistently disturbed system from (25) when driven by any driving noise W' bounded by Ω , then there exists a time-varying observer with noiseless access to the state and a time-varying controller so that any undisturbed system (1) with initial condition $|X_0| \leq \frac{\Omega}{2}$, $W_t = 0$, and $0 < \lambda < \lambda'$ can be stabilized in the sense that there exists a K so that:

$$E[|X_t|^{\eta'}] \leq K \left(\frac{\lambda}{\lambda'}\right)^{\eta' t} \quad (26)$$

Proof: Since $W_t = 0$ for $t > 0$, it is immediately clear that the system of (25) can be related to the original system of (1) by the following scaling relationships:

$$\begin{aligned} W'_0 &= X_0 \\ W'_t &= 0 \text{ if } t > 0 \\ X'_0 &= 0 \\ X'_t &= \left(\frac{\lambda'}{\lambda}\right)^{t-1} X_{t-1} \text{ if } t > 0 \\ U'_t &= \left(\frac{\lambda'}{\lambda}\right)^{t-1} U_{t-1} \end{aligned}$$

It is possible to use an observer/controller design for the system of (25) to construct one for the original system (1) through the same mapping. The input to the observer constructed with X' in mind will just be $(\frac{\lambda'}{\lambda})^t X_t$ and the controls U' just need to be scaled down by a factor $(\frac{\lambda'}{\lambda})^t$ so that they will properly apply to the X_t system.

Since (25) can be η' -stabilized, there exists a K' so that for all $t \geq 0$,

$$\begin{aligned} K' &\geq E[|X'_t|^{\eta'}] \\ &= E\left[\left(\frac{\lambda'}{\lambda}\right)^{(t-1)\eta'} |X_{t-1}|^{\eta'}\right] \\ &= \left(\frac{\lambda'}{\lambda}\right)^{\eta'(t-1)} E[|X_{t-1}|^{\eta'}] \end{aligned}$$

which immediately yields (26). \square

Lemma 4.1 can be used to get almost-sure stability by noticing that:

$$\begin{aligned} E\left[\sum_{t=0}^{\infty} |X_t|^{\eta'}\right] &= \sum_{t=0}^{\infty} E[|X_t|^{\eta'}] \\ &\leq \sum_{t=0}^{\infty} K \left(\frac{\lambda}{\lambda'}\right)^{\eta' t} \\ &\leq \frac{K}{1 - \left(\frac{\lambda}{\lambda'}\right)^{\eta'}} \end{aligned}$$

which is bounded. It immediately follows that:

$$\begin{aligned} \lim_{t \rightarrow \infty} |X_t|^{\eta'} &= 0 \text{ almost surely} \\ \lim_{t \rightarrow \infty} X_t &= 0 \text{ almost surely} \end{aligned}$$

which is summarized in the following theorem:

Theorem 4.3: If it is possible to η' -stabilize a persistently disturbed system from (25) when driven by any driving noise W' bounded by Ω , then there exists a time-varying observer with noiseless access to the state and a time-varying controller so that any undisturbed system (1) with initial condition $|X_0| \leq \frac{\Omega}{2}$, $W_t = 0$, and $0 < \lambda < \lambda'$ can be stabilized in the almost-sure²⁶ sense:

$$\lim_{t \rightarrow \infty} X_t = 0 \text{ almost surely}$$

The important thing to notice about Lemma 4.1 and Theorem 4.3 is that they do not depend on the detailed structure of the original problem except for the need to observe the state perfectly at the encoder and to be able to apply controls with perfect precision. It is clear that if either the state observation or the control application was limited in precision, then there would be no way to drive the state to zero almost surely.

Theorem 4.3 is used in Section V to get Corollary 5.3 which says that for almost-sure stabilization of an undisturbed plant across a discrete memoryless channel (DMC), Shannon capacity larger than $\log_2 \lambda$ suffices regardless of the information pattern.

F. Time in blocks and delayed observations

In the discussion so far, time has operated at the same scale for channel uses, system dynamics, plant observations, and control application. Furthermore, the only structural delay in the system was the one-step-delay across the noisy channel needed to allow the interconnection of the controller, observer, channel, and plant to make sense. It is interesting to consider different parts of the system operating at slightly different time scales and to see the impact of fixed and known delays in the system.

1) *Observing and controlling the plant on a slower time scale:* In the control context, it is natural to consider cases where the plant evolves on a slower time scale than communication. Formally, suppose that time is grouped into blocks of size n and the observer is restricted to only encode the

²⁶Here, the probability is over the channel's noisy actions and any randomness present at the observer and controller. The convergence holds for every possible initial condition and so it does not matter if the initial condition is included in the probability model.

value of X_t at times that are integer multiples of n . Similarly, suppose that the controller only takes an action²⁷ immediately before the observer will sample the state. The effective system dynamics change to

$$X_{n(k+1)} = \lambda^n X_{nk} + U_{n(k+1)-1} + W'_k \quad (27)$$

where $W'_k = \sum_{j=0}^{n-1} \lambda^{n-1-j} W_{nk+j}$. Observe that $|W'_k|$ is known to be bounded within an interval of size $\Omega' < \lambda^n \frac{\Omega}{\lambda-1}$. Essentially, everything has just scaled up by a factor of λ^n . Thus all the results above continue to hold above for a system described by (27) at times which are integer multiples of n . The rate must be larger than $\log_2 \lambda^n = n \log_2 \lambda$ bits per n time steps which translates to $\log_2 \lambda$ bits per time step. The anytime reliability $\alpha > \eta \log_2 \lambda^n = n(\eta \log_2 \lambda)$ for delay measured in units of n time-steps translates into $\alpha > \eta \log_2 \lambda$ for delay measured in unit time steps. This is the same as it was for the system described by (1).

The only remaining question is what happens to the state at times within the blocks since no controls are being applied while the state continues to grow on its own. At such times, the state has just grown by a factor of at most λ^n with an additive term of at most $\lambda^n \frac{\Omega}{\lambda-1}$.

$$\begin{aligned} & E[(\lambda^n (X_{nk} + \frac{\Omega}{\lambda-1}))^\eta] \\ &= \lambda^{\eta n} E[(X_{nk} + \frac{\Omega}{\lambda-1})^\eta] \\ &\leq \lambda^{\eta n} E\left[\left(2 \max(|X_{nk}|, \frac{\Omega}{\lambda-1})\right)^\eta\right] \\ &= \lambda^{\eta n} 2^\eta \int_0^\infty \mathcal{P}(\max(|X_{nk}|^\eta, (\frac{\Omega}{\lambda-1})^\eta) \geq \tau) d\tau \\ &= \lambda^{\eta n} 2^\eta \left((\frac{\Omega}{\lambda-1})^\eta + \int_{(\frac{\Omega}{\lambda-1})^\eta}^\infty \mathcal{P}(|X_{nk}|^\eta \geq \tau) d\tau \right) \\ &< \lambda^{\eta n} 2^\eta \left((\frac{\Omega}{\lambda-1})^\eta + \int_0^\infty \mathcal{P}(|X_{nk}|^\eta \geq \tau) d\tau \right) \\ &= \lambda^{\eta n} 2^\eta \left((\frac{\Omega}{\lambda-1})^\eta + E[|X_{nk}|^\eta] \right) \end{aligned}$$

which is finite since the original is finite. Thus:

Theorem 4.4: If for all $\Omega > 0$, it is possible to stabilize a particular unstable scalar system with gain λ^n and arbitrary disturbance signal bounded by Ω when we are allowed n uses of a particular channel between when the control-system evolves, then for any $\Omega > 0$ it is also possible to stabilize an unstable scalar system with gain λ that evolves on the same time scale as the channel using an observer restricted to only observe the system every n time steps.

By simple application of Theorem 4.4, it is known that Theorem 4.2 and similarly Theorem 3.3 continue to hold even if the observers/controllers only get access to the analog system at timesteps that are integer multiples of some n . This is used when considering noisy observations in Section IV-H and in the context of vector-valued states in Part II.

2) *Known fixed delays:* Similarly, we can study cases where the assumed “round trip delay” is larger than one. Suppose the control signal applied at time t depends only on channel outputs up to time $t - v$ for some $v > 0$.

It is easy to see that while this sort of deterministic delay does degrade performance, it does not change stability. The proof of Theorem 4.1 goes through as before. Specifically, in Section IV-C, (22) will change to:

$$\hat{X}_{t+1}(t) = \sum_{i=0}^t \lambda^i \hat{U}_{t-i}(t-v) \quad (28)$$

Everything else proceeds as before, just that in place of d for the probability of error we will have $d + v$. Specifically, in place of (24), we now know only that:

$$\begin{aligned} |X_t| &< \lambda^{d+v} \frac{2\Delta}{1-\lambda^{-1}} \\ &= \lambda^d \frac{2\Delta\lambda^v}{1-\lambda^{-1}} \end{aligned} \quad (29)$$

This is just a change in the constant factor and results in a smaller (more negative) constant K to deal with the larger uncertainty. This change of constant does not make a bounded η -moment become unbounded. The result is summarized in the following theorem:

Theorem 4.5: Theorems 4.1 and 4.2, continue to hold if the control signal U_t is required to depend only on the channel outputs up through time $t-v$ where $v \geq 0$. Only the constants grow larger.

G. Noisy or quantized controls

The control signals U_t may not be able to be set by the controller to infinite precision. The applied control \hat{U}_t at the plant might be different from the intended control U_t^i generated at the controller. This section considers the case of Γ_c -precise controls where the difference is bounded so $|U_t - \hat{U}_t^i| \leq \frac{\Gamma_c}{2}$ for some constant Γ_c to reflect the noise at the controller. It is easy to see that the plant dynamics now effectively change from 1 to:

$$X_{t+1} = \lambda X_t + U_t^i + (W_t + (U_t - U_t^i))$$

where the term $(W_t + (U_t - U_t^i))$ can be considered the new bounded disturbance for the system. So in place of Ω , we simply use the new bound $\Omega + \Gamma_c$. Thus, all the previous results continue to hold in the case of boundedly noisy control signals.

Theorem 4.6: If for all $\Omega > 0$, it is possible to stabilize a particular unstable scalar system with arbitrary disturbance signal bounded by Ω given the ability to apply precise control signals, then for all $\Gamma_c > 0$ and $\Omega > 0$, it remains possible to stabilize the same unstable scalar system with arbitrary disturbance signal bounded by Ω given the ability to apply only Γ_c -precise control signals.

²⁷The controller can take “no action” by setting $U_t = 0$.

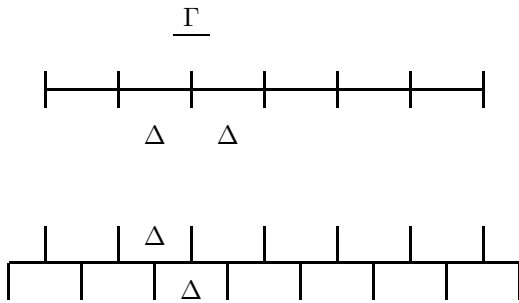


Fig. 10. With noisy observations, no strict partition of the line can adequately capture the uncertainty since it can straddle the boundary of two regions. By doubling the number of bins, it is guaranteed that the uncertainty arising from observation noise can be contained inside a single bin.

H. Noisy or quantized observations

The observer of Section IV-B has exact knowledge of the state X_t . Suppose that the observation is instead $X_{noisy}(t) = X_t + N_t$ where N_t is known to be within a bound $(-\frac{\Gamma}{2}, \frac{\Gamma}{2})$. For example, this models situations where the input to the encoder has already been quantized to some resolution.²⁸

The observer needs to ensure that the virtual state \bar{X} is within an interval of size Δ . To do this, just choose a large enough $\Delta > 2\Gamma$ so that $X_{noisy}(t)$ and X_t both pick out the same interval for the state. As figure 10 illustrates, this is not quite enough since the intervals used in Section IV-B are partitions of the real line. Meanwhile, each observation of $X_{noisy}(t)$ gives rise to an uncertainty window for $X_t \in (X_{noisy}(t) - \frac{\Gamma}{2}, X_{noisy}(t) + \frac{\Gamma}{2})$ that might straddle a boundary of the partition.²⁹ Doubling the number of intervals and having them overlap by half ensures that the uncertainty window can always fit inside a single interval. Such a doubling increases the data rate by at most an additional bit. To amortize this additional bit, Theorem 4.4 from Section IV-F is used and time is considered in blocks of size n . Then, the required rate for achievability with blocked time is $R > 1 + \log_2 \lambda^n$ bits per n time-steps or $R > \frac{1}{n} + \log_2 \lambda$ bits per time step. Since n can be large enough, $R > \log_2 \lambda$ is good enough. Delayed control actions also causes no new concerns. Thus, we get the following corollary to Theorems 4.2 and 4.5:

Corollary 4.1: It is possible to control an unstable scalar process driven by a bounded disturbance over a noisy channel so that the η -moment of $|X_t|$ stays finite for all time if the channel has feedback anytime capacity $C_{any}(\alpha) > \log_2 \lambda$ for some $\alpha > \eta \log_2 \lambda$ and the observer is allowed to observe the noisy channel outputs exactly and has a boundedly noisy view of the state.

This is true even if the control U_t is only allowed to depend on channel outputs up through time $t - v$ where $v \geq 0$.

²⁸The quantization is assumed to be coarse, but with infinite dynamic range. Section III-C tells us that finite dynamic range will impose the requirement of zero-error capacity on the link.

²⁹This will not arise for statically quantized states since those will have fixed boundaries. In that case, nothing needs to be done except ensuring that the partitions respect those boundaries.

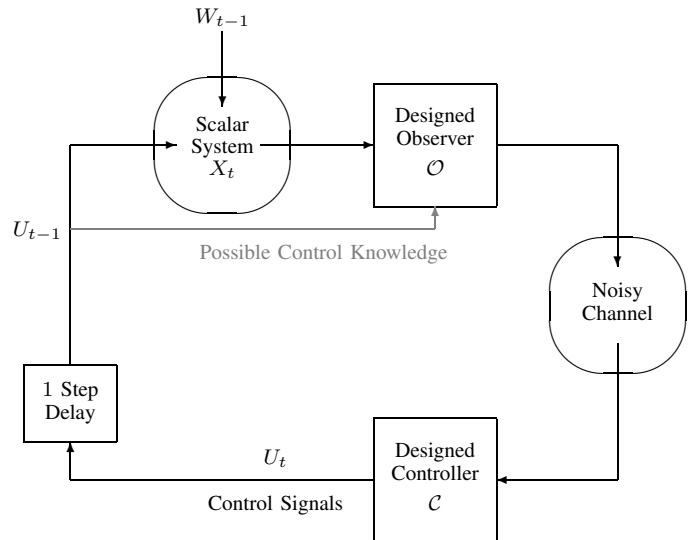


Fig. 11. Control over a noisy communication channel without explicit feedback of channel outputs.

V. RELAXING FEEDBACK

In this section, we relax the (unrealistic) assumption that the observer can observe the outputs of the noisy channel directly. This change of information pattern has the potential to make the problem more difficult. In distributed control, this was first brought out in [42] by the famous Witsenhausen counterexample. This showed that even in the case of LQG problems, nonlinear solutions can be optimal when the information patterns are not classical. This same example also showed how the “control” signals can start to play a dual role — simultaneously being used for control and to communicate missing information from one party to another [43]. Information theory also has experience with the new challenges that arise in distributed problems of source and channel coding [44].

This section restricts the information pattern in stages. First, we consider the problem of Figure 11 in which the observer can see the controls but not the channel outputs. Then, we consider the problem of Figure 12 that restricts the observer to only see the states X_t . This section is divided based on the approach rather than the problem.

In Section V-A, the solutions are based on anytime codes without feedback. These give rise to sufficient conditions that are more restrictive than the necessary conditions of Theorem 3.3. The main result is Theorem 5.2 — a random construction that shows it is possible, in the case of DMCs, to have nearly memoryless time-varying observers and still achieve stability without any feedback. All the complexity can in principle be shifted to the controller side.

In Section V-B, the solutions are based on explicitly communicating the channel outputs back to the observer through either the control signals or by making the plant itself “dance” in a stable way that communicates limited information noiselessly with no delay. Such solutions give rise to tight sufficient conditions. These are not as constructive, but serve to establish the fundamental connection between stabilization

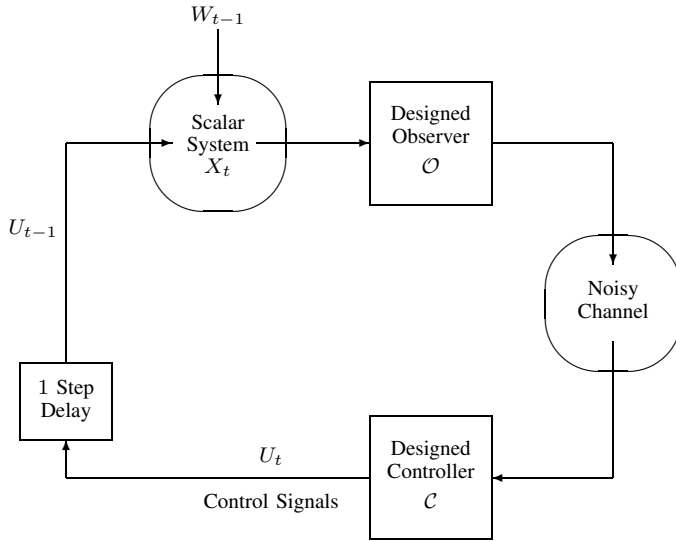


Fig. 12. Control over a noisy communication channel without any explicit feedback path from controller to observer except through the plant.

and communication with noiseless feedback.

A. Using anytime codes without feedback

Noisefree access to the control signals is not problematic in the case of Corollary 4.1 since the control signals are calculated from the perfect channel feedback. Without such perfect feedback, it is more realistic to consider only noisy access to the control signals. Furthermore, observe that in Section IV-B, knowledge of the actual applied controls is used to calculate W_t from the observed X_{t+1}, X_{t-1}, U_t . Thus, any bounded observation noise on the control signals U_t just translates into an effectively larger Γ bound on the state observation noise. By Corollary 4.1, any finite Γ can be dealt with and thus:

Corollary 5.1: It is possible to control an unstable scalar process driven by a bounded disturbance over a noisy channel so that the η -moment of $|X_t|$ stays finite for all time if the channel without feedback has $C_{\text{any}}(\alpha) > \log_2 \lambda$ for some $\alpha > \eta \log_2 \lambda$ and the observer is allowed noisy access to the control signals and the state process as long as the noise on both is bounded.

As discussed in [6], without noiseless feedback the anytime capacity will tend to be considerably lower for a given α , and so there will be a gap between the necessary condition established in Theorem 3.3 and the sufficient condition in Corollary 5.1.

Next, consider the problem of figure 12 that restricts the observer to only see the states X_t . The challenge is that the observer of Section IV-B needs to know the controls in order to remove their effect so as to focus only on encoding the virtual process \bar{X}_t . As such, a new type of observer is required:

Definition 5.1: A Δ -lattice based quantizer is a map (depicted in Figure 13 that maps inputs X to integer bins j . The j -th bin spans $(\Delta \frac{j}{2}, \Delta(\frac{j}{2} + 1)]$ and is assigned to $X \in (\Delta(\frac{j}{2} + \frac{1}{4}), \Delta(\frac{j}{2} + \frac{3}{4}))$ near the center of the bin.

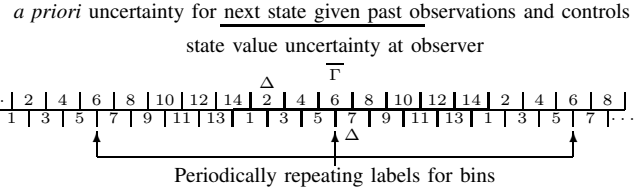


Fig. 13. A 15-regularly-labeled lattice based quantizer. If the observer had known the controls, it would have centered the lattice to cover the top bar exactly. Because it does not, one additional quantization bin must be added at the end so that the uncertainty never covers two bins bearing the same label.

A L -regularly-labeled Δ -lattice based quantizer is one which outputs $j \bmod L$ when the input is assigned to bin j — one for which the L bin labels repeat periodically.

A randomly-labeled Δ -lattice based quantizer is one which outputs A_j when the input it assigned to bin j where the A_j are drawn iid from a specified distribution.

Lattice based quantizers have some nice properties:

- Lemma 5.1:*
- If $X_{\text{noisy}}(t) = X_t + N_t$ with observation noise $N_t \in (-\frac{\Gamma}{2}, \frac{\Gamma}{2})$, then as long as $\Delta > 2\Gamma$, the bin j selected by a Δ -lattice based quantizer facing input $X_{\text{noisy}}(t)$ is guaranteed to contain X_t .
 - There exists a constant K depending only on λ, Δ, Ω so that if X_t is within a single particular bin, then X_{t+n} can be in no more than $K\lambda^n$ possible adjacent bins whose positions are a function of the control inputs applied during those n time periods as well as the original bin index for X_t .
 - If $L > K\lambda^n$ then knowing the L -regular label assigned to $X_{\text{noisy}}(t+n)$ is enough to determine a bin guaranteed to contain X_{t+n} assuming knowledge of a bin containing X_t as well as the control inputs applied during those n time periods.

Proof of [a]: $X_{\text{noisy}}(t) \in (\Delta(\frac{j}{2} + \frac{1}{4}), \Delta(\frac{j}{2} + \frac{3}{4}))$ implies $X_t \in (\Delta(\frac{j}{2} + \frac{1}{4}) - \frac{\Gamma}{2}, \Delta(\frac{j}{2} + \frac{3}{4}) + \frac{\Gamma}{2})$. But $\frac{\Gamma}{2} < \frac{\Delta}{4}$ by assumption and hence $X_t \in (\Delta(\frac{j}{2} + \frac{1}{4}) - \frac{\Delta}{4}, \Delta(\frac{j}{2} + \frac{3}{4}) + \frac{\Delta}{4}) = (\Delta \frac{j}{2}, \Delta(\frac{j}{2} + 1))$ which is the extent of the bin j .

Proof of [b]: First, suppose that the control actions were all zero during the interval in question. Because the system is linear, without loss of generality, assume that we start in the $j = 0$ bin, $[0, \Delta]$. After n time-steps, this can reach at most $[0, \lambda^n \Delta]$ without disturbances. The bounded disturbances can contribute at most

$$\begin{aligned} \sum_{i=0}^{n-1} \lambda^i \frac{\Omega}{2} &< \lambda^n \frac{\Omega}{2} \sum_{i=1}^{\infty} \lambda^{-i} \\ &= \lambda^n \frac{\Omega}{2(\lambda-1)} \end{aligned}$$

to each side, resulting in an interval of with total length $\lambda^n(\Delta + \frac{\Omega}{\lambda-1})$.

By linearity, the effect of any control inputs is a simple translation and is therefore just translates the interval by some positive or negative amount. Because of the overlapping nature of the bins, a single interval can overlap with at most 2 additional partial bins at the boundaries.

Since the bins are spaced by $\frac{\Delta}{2}$, the number of possible bins the state can be in is bounded by $2 + \lambda^n(2 + \frac{2\Omega}{\Delta(\lambda-1)})$ and so

$K = 4 + \frac{2\Omega}{\Delta(\lambda-1)}$ makes property [b] true.

Proof that [a],[b] \Rightarrow [c]: [a] guarantees that the bin corresponding to $X_{noisy}(t+n)$ is guaranteed to contain X_{t+n} . [b] guarantees there are only at most $K\lambda^n < L$ adjacent bins that the state could be in. Since the modulo operation used to assign regular labels only assigns the same label to a bin L positions away or further, all of the $K\lambda^n$ positions have distinct labels and hence the labeling of $X_{noisy}(t+n)$ picks out the unique correct bin. \square

Lemma 5.1 allows the observer to just use regular Δ -lattice quantizer to translate the state positions into bins since the control actions are side-information that is known perfectly at the intended recipient (the controller). The overhead implied by the constant K can be amortized by looking at time in blocks of n and so does not asymptotically cost any rate. This can be used to extend Corollary 5.1 to cases without any access to the control. Every n time-units, the observer can just apply the appropriate regular Δ -lattice quantizer and send the bin labels through an anytime code that operates without feedback. However, anytime codes without feedback have a natural tree structure since the impact of the distant past must never die out. In the stabilization context, this tree structure forces the observer/encoder to remember the bin sequence corresponding to all the past states. This seems wasteful since closed-loop stability implies that the plant state will keep returning to the bins in the neighborhood of the origin. This suggests that this memory at the observer is not necessary.

Theorem 5.2: It is possible to control an unstable scalar process driven by a bounded disturbance over a DMC so that the η -moment of $|X_t|$ stays finite for all time if the channel without feedback has random coding error exponent $E_r(R) > \eta \log_2 \lambda$ for some $R > \log_2 \lambda$ and the observer is allowed boundedly noisy access to the state process.

Furthermore, there exists an $n > 0$ so this is possible by using an observer consisting of a time-varying randomly-labeled Δ -lattice based quantizer that samples the state every n time steps and outputs a random label for the bin index. The random labels are chosen iid from \mathcal{A}^n according to the distribution that maximizes the random coding error exponent at R . The controller must have access to the common randomness used to choose the random bin labels.

Proof: Fix a rate $R > \log_2 \lambda$ for which $E_r(R) > \eta \log_2 \lambda$. Lemma 5.1 applies to our quantizer. Pick n, Δ large enough so that $2^{nR} > K\lambda^n$ where the K comes from property [b] above. This gives:

- d. Conditioned on actual past controls applied, the set of possible paths that the states X_0, X_n, X_{2n}, \dots could have taken through the quantization bins is a subset of a trellis that has a maximum branching factor of 2^{nR} . Furthermore, the total length covered by the d -stage descendants of any particular bin is bounded above by $K\lambda^{dn}$.

Not all such paths through the trellis are necessarily possible, but all possible paths do lie within the trellis. Figure 14 shows what such a trellis looks like and Figure 15 shows its tree like local property. Furthermore, the labels on each bin are iid through both time and across bins.

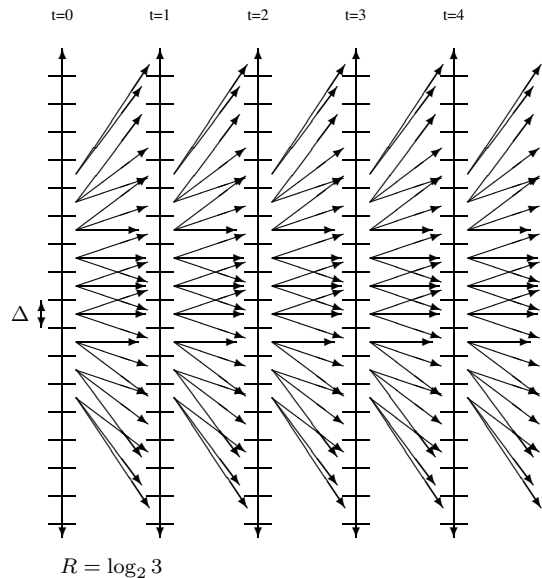


Fig. 14. A short segment of the randomly labeled regular trellis from the point of view of the controller that knows the actual control signals applied in the past. The example has $R = \log_2 3$ and $\lambda \approx 2.4$ with Δ large.

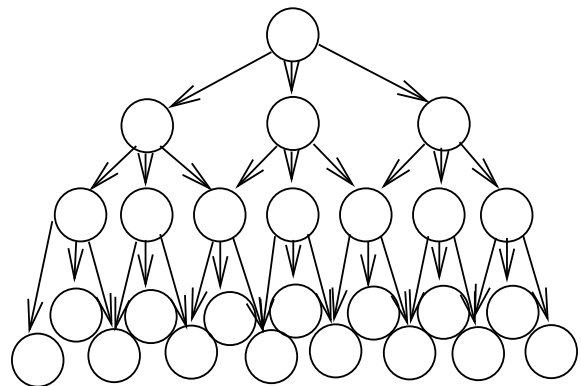


Fig. 15. Locally, the trellis looks like a tree with the nodes corresponding to the intervals where the state might have been and the levels of the tree correspond to the time. It is not a tree because paths can remerge, but all labels on disjoint paths are chosen so that they are independent of each other.

Call two paths of length t through the trellis disjoint with depth d if their last common node was at depth $t - d$ and the paths are disjoint after that. Consequently:

- e. If two paths are disjoint in the trellis at a depth of d , then the channel inputs corresponding to the past dn channel uses are independent of each other.

The suboptimal controller just searches for the ML path through the trellis. The trellis itself is constructed based on the controller's memory of all past applied controls. Once an ML path has been identified, a control signal is applied based on the bin estimate at the end of the ML path. The control signal just attempts to drive the center of that bin to zero.

Consider an error event at depth d . This represents the case that the maximum likelihood path last intersected with the true path dn time steps ago. By property [d] above, the control will be based on a state estimate that can be at most $K\lambda^{dn}$ bins away from the true state. Thus:

- f. If an error event at depth d occurs at time t , the state $|X_{t+n}|$ can be no larger than $K'\lambda^{(d+1)n}$ for some constant $K' = 2\Delta K$ that does not depend on d or t .

Property [f] plays the role of (24) in this proof.

By property [d], there are no more than 2^{dnR} possible false paths that last intersected the true path d stages ago. By the memorylessness of the channel, the log-likelihood of each path is the sum of the likelihood of the “prefix” of the path leading up to d stages ago and the “suffix” of the path from that point onward. For a path that is disjoint from the true path at a depth of d to beat all paths that end up at the true final state, the false path must have a suffix log-likelihood that beats the suffix log-likelihood of at least the true path. Property [e] guarantees that the channel inputs corresponding to the false paths are pairwise independent of the true inputs for the past dn channel uses.

All that is required to apply Gallager’s random block-coding analysis of Chapter 5 in [1] is such a pairwise independence³⁰ between the true and false codewords for a code of length dn .

- g. The probability that the ML path diverges from the true path at depth d is no more than $2^{-dnE_r(R)}$.

All that remains is to analyze the η -moment by combining [g] and [f] and using the union bound to compute the expectation.

$$\begin{aligned} E[|X_{t+n}|^\eta] &\leq \sum_{d=0}^{\frac{t}{n}} 2^{-dnE_r(R)} (K'\lambda^{(d+1)n})^\eta \\ &< (K'\lambda^n)^\eta \sum_{d=0}^{\infty} 2^{-dnE_r(R)} \lambda^{\eta dn} \\ &= (K'\lambda^n)^\eta \sum_{d=0}^{\infty} 2^{-dn(E_r(R) - \eta \log_2 \lambda)} \\ &= K'' < \infty \end{aligned}$$

where the final geometric sum converges since $E_r(R) > \eta \log_2 \lambda$. \square

Although the condition in Theorem 5.2 is not tight, the result has several nice features. First, it allows easy verification of sufficiency for a good channel since $E_r(R)$ is easy to calculate. Structurally, it demonstrates that there is no need to use very complex observers. The intrinsic memory in the plant can play the role of the memory that would otherwise need to be implemented in a channel code. The complexity can be shifted to the controller, and even that complexity is not too bad. Sequential decoding can be used at the controller since it is known to have the same asymptotic performance with respect

³⁰Notice that pairwise independence is also obtained if the random labels were assigned using an appropriate random time-varying infinite constraint-length convolutional code (with the symbol-merging tricks of Figure 6.2.1 of [1] to match the desired channel input-distribution) applied to the binary expansion of the integer j corresponding to the selected bin at each stage. Since the closed-loop system is stable, the state is presumably small and the bin is close to 0. As such, all of the higher-order bits in the binary expansion of the bin label are zeros and do not cause any computational burden when operating the convolutional code. This is related to the feedback convolutional codes with variable constraint-lengths discussed further in [6]. Because of this, the computational burden of running this observer is non-increasing with time.

to delay as the ML decoder[45], [46]. Because the closed-loop system is stable and thereby renews itself constantly, the computational burden of running sequential decoding (and hence the controller) does not grow unboundedly with time [47].

Since $E_r(R, Q) > 0$ for all $R < C$ and the capacity-achieving distribution Q , Theorem 5.2 can also be recast in a weaker Shannon capacity-centric form:

Corollary 5.2: If the observer is allowed boundedly noisy access to the plant state, and the noisy channel is a DMC with Shannon capacity $C > \log_2 \lambda$, then there exists some $\eta > 0$ and an observer/controller pair that stabilizes the system in closed loop so that the η -moment of $|X_t|$ stays finite for all time.

Furthermore, there exists an $n > 0$ so this is possible by using an observer consisting of a time-varying randomly-labeled Δ -lattice based quantizer that samples the state every n time steps and outputs a random label for the bin index. This random labels are chosen iid from the \mathcal{A}^n according to the capacity-achieving input distribution. The controller must have access to the common randomness used to choose the random bin labels.

Applying Theorem 4.3 to Corollary 5.2 immediately results in the following new corollary:

Corollary 5.3: If the observer is allowed perfect access to the plant state, and the noisy channel is a DMC with Shannon capacity $C > \log_2 \lambda$, then there exists an observer/controller pair that stabilizes the system (1) in closed loop so that:

$$\lim_{t \rightarrow \infty} X_t = 0 \text{ almost surely}$$

as long as the initial condition $|X_0| \leq \frac{\Omega}{2}$ and the disturbances $W_t = 0$.

Furthermore, there exists an $n > 0$ so this is possible by using an observer consisting of a time-varying randomly-labeled Δ_t -lattice based quantizer that samples the state every n time steps and outputs a random label for the bin index. The Δ_t shrink geometrically with time, and the random labels are chosen iid from the \mathcal{A}^n according to the capacity-achieving input distribution. The controller must have access to the common randomness used to choose the random bin labels.

B. Communicating the channel outputs back to the observer

In this section, the goal is to recover the tight condition on the channel from Theorem 4.2. To do this, we construct a controller that explicitly communicates the noisy channel outputs to the observer using whatever “channels” are available to it. First we consider using a noiseless control signal to embed the feedback information. This motivates the technique used to communicate the feedback information by making the plant itself dance in a stable way that tells the observer the channel output.

1) *Using the controls to communicate the channel outputs:* The idea is to “cheat”³¹ and communicate the channel outputs through the controls. The control signal is thus serving dual purposes — stabilization of the system and the communication of channel outputs. Suppose the observer had noiseless access to the control signals. The controller can choose to quantize its real-valued controls to some suitable level and then use the infinite bits remaining in the fractional part to communicate the channel outputs to the observer. The observer can then extract these bits noiselessly and give them to the anytime encoder as noiseless channel feedback.

Of course, this additional fractional part will introduce an added disturbance to the plant. One approach is to just consider the quantization and channel output communication terms together as a bounded noise on the control signals considered in Section IV-G. This immediately yields:

Corollary 5.4: It is possible to control an unstable scalar process driven by a bounded disturbance over a noisy channel so that the η -moment of $|X_t|$ stays finite for all time if the channel has feedback anytime capacity $C_{\text{any}}(\alpha) > \log_2 \lambda$ for some $\alpha > \eta \log_2 \lambda$ and the observer is allowed to observe the control signals perfectly.

However, the additional disturbance introduced by the quantization of the original control signal and the introduction of the new fractional part representing the channel output is known perfectly at the controller end. Meanwhile, the output of the virtual-process based observer does not depend on the actual applied controls anyway since it subtracts them off. So rather than compensating for this quantization+signaling by expanding the uncertainty Ω and thus changing the Δ at the observer, the controller can just clean up after itself. This idea allows us to eliminate all access to the control signals at the observer and generalizes to many cases of countably large channel output alphabets.

2) *Removing noiseless access to the controls at the observer:* There are two tricks involved. The first is the idea of making the plant “dance” appropriately and using the moves in the dance to communicate the channel outputs. The second idea is to introduce an artificial delay of 1 time step in the determination of the “non-dance” component of the control signals. This makes the non-dance component completely predictable by the observer and allows the observer to clearly see the dance move corrupted only by the bounded process disturbance. Putting it together gives:

Theorem 5.3: Given a noisy channel with a countable alphabet, identify the channel output alphabet with the integers and suppose that there exist $K > 0, \beta > \eta$ so that the channel

³¹We call this “cheating” since it violates the spirit of the requirement against access to the channel outputs. However, it is important to establish this result because it points out the need for a serious future study where the communication constraints back from the controller to the observer are modeled more carefully. A more realistic model for the problem should have a sensor observing the plant connected via a communication channel to the controller. The controller is then connected to an actuator through another communication channel. The actuator finally acts upon the plant itself. With no complexity constraints, this reduces to the case studied here with the controller merely playing the role of a relay bridging together two communication channels. The relay anytime reliability will become the relevant quantity to study.

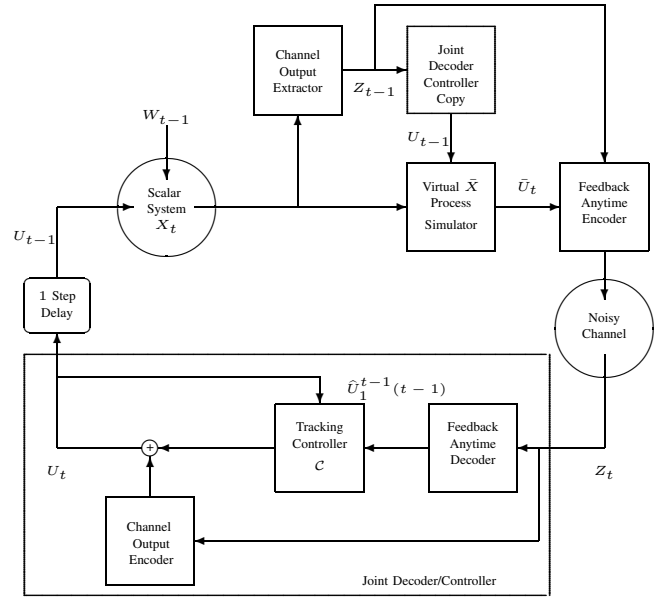


Fig. 16. Overlaying messages onto the control signal and recovering the messages at the observer. The control signal is generated based on unit-delayed channel outputs with the current output being communicated back.

outputs B_t satisfy: $\mathcal{P}(|B_t| \geq i) \leq Ki^{-\beta}$ for all t regardless of the channel inputs.

Then, it is possible to control an unstable scalar plant driven by a bounded disturbance over that channel so that the η -moment of $|X_t|$ stays finite for all time if the channel has feedback anytime capacity $C_{\text{any}}(\alpha) \geq \log_2 \lambda$ for some $\alpha > \eta \log_2 \lambda$ even if the observer is only allowed to observe the state X_t corrupted by bounded noise.

Proof: The overall strategy is illustrated in Figure 16. The channel output extraction at the observer is illustrated in Figure 17 in the context of a channels with output alphabet size $|\mathcal{B}| = 5$.

Let $U_t(b_0^{t-1})$ be the control that would be applied from Theorem 4.5 as transformed by the action of Theorem 4.6 if necessary. It only depends on the strictly past channel outputs.

Let b_t be the current channel output. The control applied is:

$$U'_t(b_0^t) = U_t(b_0^{t-1}) + F(b_t) - \lambda (U'_{t-1}(b_0^{t-1}) - U_{t-1}(b_0^{t-2})) \quad (30)$$

where the function $F(b_t)$ is the “dance move” corresponding to the channel output.

First consider the case that perfect state observations X_t are available at observer. At time t the observer can see the control signal only as it is corrupted by the process disturbance since $U_{t-1} + W_{t-1} = X_t - \lambda X_{t-1}$. By observing X perfectly, the observer has in effect gained boundedly noisy access to the U with $\Gamma_u = \Omega$. Now suppose that the observations of X were boundedly noisy with some Γ . In that case:

$$\begin{aligned} & |U_{t-1} - (X_{\text{noisy}}(t) - \lambda X_{\text{noisy}}(t-1))| \\ &= |U_{t-1} - (X_t - \lambda X_{t-1}) + (\lambda N_{t-1} - N_t)| \\ &= |-W_{t-1} + (\lambda N_{t-1} - N_t)| \\ &\leq \Omega + (\lambda + 1)\Gamma \end{aligned}$$

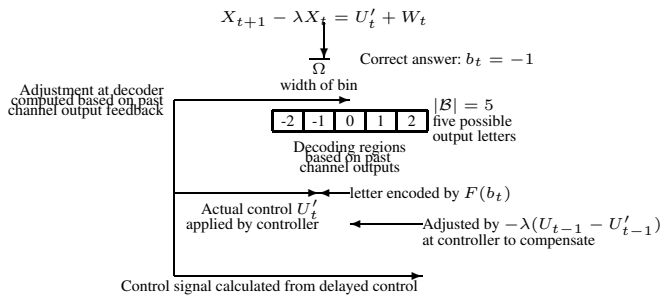


Fig. 17. How to communicate the channel outputs through the plant with state observations only. The controller restricts its main control signal to be calculated with an extra delay of 1 time unit and then adjusts it by $-\lambda(U_{t-1} - U'_{t-1})$ to eliminate the effect of the past communication. The final control signal applied is shifted slightly to encode which b_t was received. The decoder uses the past b_0^{t-1} to align its decoding regions and then reads off b_t by using $X_{t+1} - \lambda X_t$.

In this case, the effective observation noise on the controls is bounded by $\Gamma_u = \Omega + (\lambda + 1)\Gamma$.

Just by looking at the state and its history, the observer has access to U'_t with the property that $|U'_t - U_t^o| \leq \Gamma_u$. To ensure decodability of b_t , set $F(b_t) = 3\Gamma_u b_t$ so the channel outputs are modulated to be integer multiples of $3\Gamma_u$.

At time $t = 0$, the observer is unchanged since there is nothing for it to learn and no applied controls. At time $t = 1$, because of the induced delay of 1 extra time step, there are no delayed controls ready to apply either and so the applied control only consists of $3\Gamma_u b_0$. This is observed up to precision Γ_u and so the observer can uniquely recover b_0 and feed it to its anytime encoder.

Assume now that the observer was successful in learning b_0^{t-1} in the past. Then it can compute the $U_t(b_0^{t-1})$ term as well as the $U'_{t-1}(b_0^{t-2}) - U_{t-1}(b_0^{t-1})$ using this knowledge and can subtract both of them from its observed U'_t . This leaves only the $3\Gamma_u b_t$ term which can be uniquely decoded given that the observation noise is no more than Γ_u in either direction. By induction, the observer can effectively recover the past channel outputs from its noiseless observations of the control signal and can thereby operate the feedback anytime-encoder successfully.

The communication of each channel output b_t only impacts the very next state by shifting it by $3\Gamma_u b_t$. At the next time, it is canceled out by the correction term $-\lambda(U'_{t-1}(b_0^{t-1}) - U_{t-1}(b_0^{t-2}))$. The non-dancing controlled state $X'_{t+1} = (X_{t+1} - 3\Gamma_u b_t)$ has at least a power-law tail $\mathcal{P}(X'_{t+1} \geq x) \leq K' x^{-(\eta+\epsilon)}$ for some K' and $\epsilon > 0$. Then

$$\begin{aligned}
 E[|X|^\eta] &= \\
 &= \int_0^\infty \mathcal{P}(|X' + 3\Gamma_u B| \geq m^{\frac{1}{\eta}}) dm \\
 &\leq \int_0^\infty \mathcal{P}\left(2 \max(|X'|, 3\Gamma_u |B|) \geq m^{\frac{1}{\eta}}\right) dm \\
 &\leq \int_0^\infty \mathcal{P}\left(|X'| \geq \frac{1}{2} m^{\frac{1}{\eta}}\right) dm + \\
 &\quad + \mathcal{P}\left(|B| \geq \frac{1}{6\Gamma_u} m^{\frac{1}{\eta}}\right) dm \\
 &\leq \int_0^\infty K' \left(\frac{1}{2} m^{\frac{1}{\eta}}\right)^{-(\eta+\epsilon)} + K \left(\frac{1}{6\Gamma_u} m^{\frac{1}{\eta}}\right)^{-\beta} dm
 \end{aligned}$$

Since $\beta > \eta$, this converges and so the η -moment of X also exists. \square

The channel output condition in 5.3 is clearly satisfied whenever the channel has a finite output alphabet. Beyond that case, it is satisfied in generic situations when the input alphabet is finite and the transition probabilities $p(b|a)$ individually have an light enough tail for each one of the finite a values.³² When the channel input alphabet is itself countable, the condition is harder to check.

If information must flow noiselessly from the controller to the observer, the key question is to quantify the instantaneous zero-error capacity of the effective channel *through the plant*. Here, the bounded support of W and the unconstrained nature of U are critical since they allow the instantaneous zero-error capacity of that effective channel to be infinite. Of course, there remains the problem of the dual-nature of the control signal — it is simultaneously being asked to stabilize the plant as well as to feedback information about the channel outputs. The theorem shows that the ability of the controller to move the plant provides enough feedback to the encoder in the case of finite channel output alphabets or channels with uniformly exponentially bounded output statistics.

At an abstract level, the controller is faced with the problem of causal “writing on dirty paper”[48] where the information it wishes to convey in one time step is the channel output and the dirty paper consists of the control signals it must apply to keep the system stable and to counteract the effect of the writing it did in previous time steps. Here, the problem is finessed by introducing the artificial delay at the controller to ensure that the “dirt” is side-information known both to the transmitter and the receiver. For finite output alphabets, it is also possible to take a direct “precoding” approach to do this by encoding the channel outputs by placing the control to the appropriate value modulo $3\Gamma_u(|\mathcal{B}| + 1)$. This is a bounded perturbation of the control inputs and Theorem 4.6 tells us that this does not break stability if the Δ is adjusted appropriately.

Finally, it might seem that this particular “dance” by the plant will be a disaster for performance metrics beyond stabilization. This is probably true, but we conjecture that such implicit feedback through the plant will be usable without much loss of performance. If it has memory, the observer can notice when and how the channel has misbehaved since the

³²For example, an AWGN channel with a hard-input constraint and quantized outputs.

plant's state will start growing rather than staying near 0. The Δ -lattice based quantizer used in the observer for Theorem 5.2 could not exploit this because it was memoryless and used uniformly sized bins regardless of whether the state was large or small.

VI. CONTINUOUS TIME SYSTEMS

A. Overview

So far, we have considered a discrete-time model (1) for the dynamic system that must be stabilized over the communication link. This has simplified the discussion by having a common clock that drives both the system and the uses of the noisy channel. In general, there will be a τ_c that represents the time between channel uses. This allows translating everything into absolute time units.

$$\dot{X}(t) = \lambda X(t) + U(t) + W(t), \quad t \geq 0 \quad (31)$$

where the bounded disturbance $|W(t)| \leq \frac{\Omega}{2}$ and there is a known initial condition $X(0) = 0$. If the open-loop system is unstable, then $\lambda > 0$.

Sampling can be used to extend both the necessity and sufficiency results to the continuous time case. The basic result is that stability requires an anytime capacity greater than λ nats per second.

B. Necessity

For necessity, we are free to choose the disturbance signal $W(t)$ and consequently can restrict ourselves to piecewise constant signals³³ that stay constant for time τ . By sampling at the rate $\frac{1}{\tau}$, the sampled state evolves as $X(\tau(i+1)) =$

$$e^{\lambda\tau} X(\tau i) + \left(\frac{e^{\lambda\tau} - 1}{\lambda} W_i \right) + \int_{i\tau}^{(i+1)\tau} U(s) e^{\lambda(\tau(i+1)-s)} ds \quad (32)$$

Notice that (32) is just a discrete time system with $\lambda' = e^{\lambda\tau}$ taking the role of λ in (1), and the disturbance is bounded by $\Omega' = \frac{\Omega(e^{\lambda\tau}-1)}{\lambda}$. All that remains is to reinterpret the earlier theorem.

By setting $\tau = \tau_c$ to match up the sampling times to the channel use times, it is clear that the appropriate anytime capacity must exceed $\log_2 \lambda' = \tau_c \lambda \log_2 e$ bits per channel use. By converting units to nats per second³⁴, we get the intuitively appealing result that the anytime capacity must be greater than λ nats/sec.³⁵ Similarly, to hold the η -th moment constant, the probability of error must drop with delay faster than $K2^{-(\eta\lambda \log_2 e)d\tau_c}$ where d is in units of channel uses and thus $d\tau_c$ has units of seconds. Thus, we get the following pair of theorems:

Theorem 6.1: For a given noisy channel and $\eta > 0$, if there exists an observer \mathcal{O} and controller \mathcal{C} for the unstable scalar continuous time system that achieves $E[|X(t)|^\eta] < K$ for all t and bounded driving noise signals $|W(t)| \leq \frac{\Omega}{2}$, then

the channel's feedback anytime capacity $C_{\text{any}}(\eta\lambda \log_2 e) \geq \lambda$ nats per second.

Theorem 6.2: For a given noisy channel and decreasing function $f(m)$, if there exists an observer \mathcal{O} and controller \mathcal{C} for the unstable continuous-time scalar system that achieves $\mathcal{P}(|X(t)| > m) < f(m)$ for all t and all bounded driving noise signals $|W(t)| \leq \frac{\Omega}{2}$, then $C_{\text{g-any}}(g) \geq \lambda$ nats per second for the noisy channel considered with the encoder having access to noiseless feedback and $g(d)$ having the form $g(d) = f(Ke^{\lambda d})$ for some constant K .

C. Sufficiency

For sufficiency, the disturbance is arbitrary but we are free to sample the signal as desired at the observer and apply piecewise constant control signals. Sampling every τ units of time gives rise to (32) only with the roles of W and U reversed. It is clear that $W_i = \int_{i\tau}^{(i+1)\tau} W(s) e^{\lambda(\tau(i+1)-s)} ds$ is still bounded by substituting in the upper and lower bounds and then noticing that $|W_i| \leq \frac{\Omega(e^{\lambda\tau}-1)}{2\lambda}$.

Thus, the same argument above holds and the sufficiency Theorems 4.1, 4.2, and 5.3 as well as Corollaries 5.4 and 5.1 translate cleanly into continuous time. In each, the relevant anytime capacity must be greater than λ nats per second. Since the necessary and sufficient conditions are right next to each other, it is clear that the choice of sampling time does not impact the sense of stability that can be achieved. Of course, this need not be optimal in terms of performance.

Finally, if the channel we face is an input power-constrained ∞ -bandwidth AWGN channel, more can be said. Section III-C.4 makes it clear that nothing special is required in this case: using linear controllers and observers is good enough if the average power constraint is high enough. But what if the channel had a hard amplitude constraint that allowed the encoder no more than P power per unit time? In this case, it is possible to generalize Theorem 5.2 in an interesting way.

In [49] we give an explicit construction of a feedback-free anytime code for the infinite bandwidth AWGN channel that uses a sequential form of orthogonal signaling. In the ∞ -bandwidth AWGN channel, pairwise orthogonality between codewords plays the role that pairwise independence does for DMCs. Applying that principle through the proof of Theorem 5.2, the observer/encoder can simply be a *time-invariant* regular partition of the state space with the bins being labeled with orthogonal pulses, each with an energy equal to the hard limit for the channel.³⁶ The encoder just pieces together pulses with shapes corresponding to where the state is at the sampling times. The controller then searches for the most likely path based on the channel output signal as well as the past control values, and then applies a control based on the current estimate. This approach allows the use of occasional bandwidth expansion to deal with unlucky streaks of channel

³⁶In particular, the following sequence of pulses work with an appropriate scaling. For $0 \leq t \leq \tau$, set $g_{i,\tau}(t) = \frac{1}{\tau} \text{sgn}(\sin(\frac{4\pi i}{\tau}t))$ and $g_{-i,\tau}(t) = \frac{1}{\tau} \text{sgn}(\sin(\frac{2\pi(2i-1)}{\tau}t))$ and zero everywhere else. Here τ is the time between taking samples of the state. The $g_{i,\tau}$ functions are orthogonal, and the i -th function is the channel input corresponding to the i -th lattice bin for the plant state observation.

³³zero order hold

³⁴Assuming that \dot{X} is in per second units.

³⁵This truly justifies nats as the "natural" unit of information!

noise while keeping the channel input power constant. The details of this approach are given in [50].

VII. A HIERARCHY OF COMMUNICATION PROBLEMS

In this final section, we interpret some of the results in a different way inspired by the approach used in computational complexity theory. There, the scarce resource is the time and space available for computation and the asymptotic question is whether or not a certain family of problems (indexed by n) can be solved using the limited amount of resource available. While explicit algorithms for solving problems do play a role, “reductions” from one problem to another also feature prominently in relating the resource requirements among related problems [51].

In communication, the scarce resource can be thought of as being the available channel.³⁷ Problems should be ordered by what channels are good enough for them. We begin with some simple definitions and then see how they apply to classical results from information theory. Finally, we interpret our current results in this framework.

Definition 7.1: A communication problem is a partially specified random system together with an information pattern and a performance objective. This is specified by a triple: $(\mathcal{S}, \mathcal{I}, \mathcal{V})$. The partially specified random system $\mathcal{S} = (S_0, S_1, \dots)$ in which S_i are real valued functions on $[0, 1]^{i+1} \times \mathbb{R}^i$. The output of the S_i function is denoted X_i . The information pattern \mathcal{I} identifies what variables each of the i -th encoders and decoders has access to. The performance objective \mathcal{V} is a statement that must evaluate to either true or false once the entire random system is specified.

As depicted in Figure 18, the communication problem is thus an open system that awaits interconnection with encoder, channel, and decoder maps. The channel is a measurable map f_c from $[0, 1] \times \mathbb{R}$ into \mathbb{R} . The encoder and decoder are both represented by a possibly time-varying sequence of real valued functions compatible with the information pattern \mathcal{I} .

Once all the maps are specified, the random system becomes completely specified by tying them to an underlying probability space consisting of three iid sequences (W_i, V_i, R_i) of continuous uniform random variables on $[0, 1]$. The W_i^i are connected to the first input of S_i while V_i is connected to the first input of the memoryless channel. As is usual, the output of the encoder is connected to the remaining input of the channel, and all the past outputs of the channel are connected to the decoding functions as per the information patterns. Finally, assume that common randomness R_i is made available to both the encoder and decoder so that they may do random coding if desired. Once everything is connected, it is possible to evaluate the truth or falsehood of \mathcal{V} .

Definition 7.2: A channel is said to solve the problem if there exist suitable encoder and decoder maps compatible with the given information pattern so that the combined random system satisfies the performance objective \mathcal{V} .

Communication problem A is *harder than* problem B if any channel f_c that solves A also solves B .

³⁷This might in turn be related to other more primitive scarce resources like power or bandwidth available for communication.

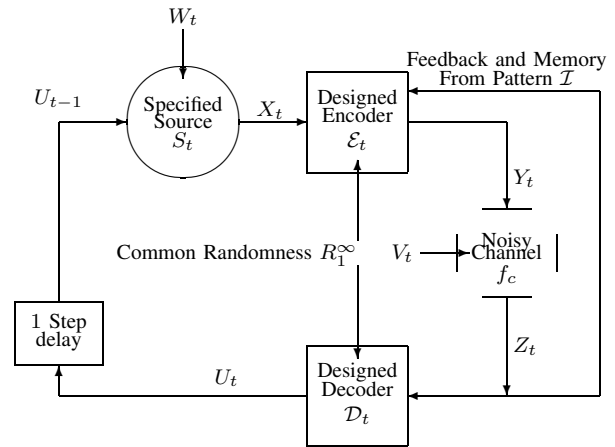


Fig. 18. Abstractly, a communication problem consists of a partially specified random system consisting of a known and possibly interactive source together with an information pattern. The noisy channel and encoder/decoders need to be specified before all the random variables become properly defined.

Each particular communication problem therefore divides channels into two classes: those that solve it and those that do not. Suitable families of communication problems, ordered by hardness, can then be used to sort channels as well. Channels that solve harder problems are better than ones that do not. The equivalence of certain families of communication problems means that they induce the same orderings on communication channels. This will become clearer by the examples of the next few sections.

A. Classical Examples

1) *The Shannon communication problem:* Shannon identified the problem of communicating bits reliably as one of the core problems of communication. In our framework, this problem is formalized as follows:

- $X_i = 1$ if $W_i > \frac{1}{2}$ and $X_i = 0$ otherwise. The functions S_i ignore all other inputs.
- The information pattern \mathcal{I} specifies that \mathcal{D}_i has access to Z_1^i . The encoder information pattern is complete in the case of communication with feedback: \mathcal{E}_i has access to X_1^i as well as Z_1^{i-1} . Without feedback, \mathcal{E}_i has access only to X_1^i .
- The performance objective $\mathcal{V}(\epsilon, d)$ is satisfied if $\mathcal{P}(X_i \neq U_{i+d}) \leq \epsilon$ for every $i \geq 0$.

The Shannon communication problem naturally comes in a pair of families $A_{\epsilon, d}^f$ with feedback and $A_{\epsilon, d}^{n, f}$ without feedback. These families are indexed by the tolerable probability of bit error ϵ and end-to-end delay d .

To obtain other rates $R > 0$, adjust the source functions as follows:

- $X_i = \begin{cases} 1 & \text{if } W_i \in [\frac{j}{2^{\lfloor Ri \rfloor - \lfloor R(i-1) \rfloor}}, \frac{j+1}{2^{\lfloor Ri \rfloor - \lfloor R(i-1) \rfloor}}) \end{cases}$ for integer $j \geq 0$. The possibly time-varying functions S_i ignore all other inputs.

These naturally result in families $A_{R, \epsilon, d}^f$ and $A_{R, \epsilon, d}^{n, f}$ for the feedback and feedback-free cases respectively. It is immediately clear that $A_{R, \epsilon, d}^{n, f}$ is harder than $A_{R, \epsilon, d}^f$ and furthermore

problems with smaller ϵ or d are harder than those with larger ones. It is also true that $A_{R,\epsilon,d}^f$ is harder than $A_{R',\epsilon,d}^f$ whenever $R \leq R'$ in that it is more challenging to communicate reliably at a high rate rather than a low one.

The set of channels with classical Shannon feedback capacity of at least R is therefore:

$$\mathcal{C}_R^f = \bigcap_{\epsilon > 0} \bigcap_{R' < R} \bigcup_{d > 0} \{f_c | f_c \text{ solves } A_{R',\epsilon,d}^f\} \quad (33)$$

and similarly for \mathcal{C}_R^{nf} . The classical result that feedback does not increase capacity tells us that $\mathcal{C}_R^f = \mathcal{C}_R^{nf}$. Because of this, we just call them both \mathcal{C}_R .

2) *The zero-error communication problem:* A second problem is the one of zero error communication. It is defined exactly the same as the Shannon communication problem above, except that $\epsilon = 0$.

The channels that have feedback zero-error capacity of at least R with feedback are therefore:

$$\mathcal{C}_{0,R}^f = \bigcap_{R' < R} \bigcup_{d > 0} \{f_c | f_c \text{ solves } A_{R',0,d}^f\} \quad (34)$$

and similarly for $\mathcal{C}_{0,R}^{nf}$. In this case, the result with and without feedback can be different and furthermore, $\mathcal{C}_{0,R}^{nf} \subset \mathcal{C}_{0,R}^f \subset \mathcal{C}_R$ [25]. In this sense, zero-error communication is fundamentally a harder problem than ϵ -error communication.

3) *Estimation problems with distortion constraints:* Consider iid real valued sources with cumulative distribution functions $F_X(t) = \mathcal{P}(X \leq t)$.

- $X_i = F_X^{-1}(W_i)$ ignoring all the other inputs. This gives the desired source statistics.
- The information patterns remain as in the Shannon problem.
- The performance objective $\mathcal{V}(\rho, D, d)$ is satisfied if $\lim_{n \rightarrow \infty} \frac{1}{n} E[\sum_{i=1}^n \rho(X_i, U_{i+d})] \leq D$.

Call these estimation problems $A_{(F_X, \rho, D, d)}^f$ and $A_{(F_X, \rho, D, d)}^{nf}$ (for the cases with/without feedback) and once again associate them with the set of channels that solve them in the limit of large delays:

$$\mathcal{C}_{e,(F_X, \rho, D)}^f = \bigcap_{D' > D} \bigcup_{d > 0} \{f_c | f_c \text{ solves } A_{(F_X, \rho, D', d)}^f\} \quad (35)$$

and similarly for $\mathcal{C}_{e,(F_X, \rho, D)}^{nf}$. For cases where the distortion ρ is bounded, the existing separation result can be interpreted as follows:

$$\mathcal{C}_{R(D)} = \mathcal{C}_{e,(F_X, \rho, D)}^{nf} = \mathcal{C}_{e,(F_X, \rho, D)}^f \quad (36)$$

where $R(D)$ is the information-theoretic rate-distortion curve.

The interpretation of this separation theorem is that in the limit of large delays, estimation problems with a fidelity constraint are no harder or easier than Shannon communication problems dealing with bits. Both families of problems induce essentially the same partial order on channels.

B. Anytime communication problems

The anytime communication problems are natural generalizations of the binary data communication problems above. Everything remains as in the Shannon communication problem, only the performance measure changes. Let $U_t = 0.\widehat{X}_0(t), \widehat{X}_1(t), \widehat{X}_2(t), \dots$ when written out in binary notation. This can always be done and the parsing of the string is unique no matter what the rate is.

- $\mathcal{V}_{(K, \alpha)}$ is satisfied if $\mathcal{P}(X_i \neq \widehat{X}_i(i+d)) \leq K2^{-\alpha d}$ for every $i \geq 0, d \geq 0$.

Call these problems $A_{(R, \alpha, K)}^f$ when feedback is allowed and $A_{(R, \alpha, K)}^{nf}$ when it is not permitted. Once again, it is clear that the non-feedback problems are harder than the corresponding feedback problems. Furthermore, $A_{(R, \alpha, K)}$ is harder than $A_{(R, \alpha', K)}$ if $\alpha' \leq \alpha$ in addition to the usual fact of $A_{(R, \alpha, K)}$ being harder than $A_{(R', \alpha, K)}$ if $R' \leq R$. Similarly, smaller K values are harder than larger ones.

The channels with α -anytime feedback capacity of at least R are then given by:

$$\mathcal{C}_{a,(R, \alpha)}^f = \bigcap_{R' < R} \bigcap_{\alpha' < \alpha} \bigcup_{K > 0} \{f_c | f_c \text{ solves } A_{(R', \alpha', K)}^f\} \quad (37)$$

with a similar definition for $\mathcal{C}_{a,(R, \alpha)}^{nf}$. It is immediately clear that

$$\mathcal{C}_{0,R}^f \subseteq \mathcal{C}_{a,(R, \alpha)}^f \subseteq \mathcal{C}_R$$

The case of $\alpha = 0$ is defined as the limit:

$$\mathcal{C}_{a,(R, 0)}^f = \bigcup_{\alpha > 0} \mathcal{C}_{a,(R, \alpha)}^f \quad (38)$$

It turns out in this case that $\mathcal{C}_{a,(R, 0)}^f = \mathcal{C}_{a,(R, 0)}^{nf} = \mathcal{C}_R$ since infinite random tree codes can be used to communicate reliably at all rates below the Shannon capacity [23].

However, for other $\alpha > 0$,

$$\mathcal{C}_{0,R}^{nf} \subset \mathcal{C}_{a,(R, \alpha)}^{nf} \subset \mathcal{C}_{a,(R, \alpha)}^f \subset \mathcal{C}_R$$

and

$$\mathcal{C}_{0,R}^{nf} \subset \mathcal{C}_{0,R}^f \subset \mathcal{C}_{a,(R, \alpha)}^f \subset \mathcal{C}_R$$

with all of these being strict inclusion relations. $\mathcal{C}_{0,R}^f$ and $\mathcal{C}_{a,(R, \alpha)}^{nf}$ are not subsets of each other in general.

In this sense, there is a non-trivial hierarchy of problems with Shannon communication as the easiest example and zero-error communication as the hardest.

C. Control and the relation to anytime communication

The stabilization problems considered in this paper are different in that they are interactive. The formulation should be apparent by comparing Figure 18 with Figure 2.

- X_i represents the state of the scalar control problem with unstable system dynamics given by $\lambda > 1$. The W_t is the bounded disturbance and U_i represents the control signal used to generate X_{i+1} .
- The information pattern with and without feedback is as before.
- The performance objective $\mathcal{V}_{(\eta, K)}$ is satisfied if $E[|X_i|^\eta] \leq K$ for all $i \geq 0$.

Call this problem $A_{(\lambda,\eta,K)}^f$ for cases with feedback and $A_{(\lambda,\eta,K)}^{nf}$ for cases without feedback available at the encoder. The problem without feedback is harder than the problem with feedback. It is also clear that $A_{(\lambda,\eta,K)}^f$ is harder than $A_{(\lambda',\eta,K)}^f$ whenever $\lambda \geq \lambda'$ and similarly for A^{nf} . The same holds if η is made larger or K is made smaller.

$$\mathcal{C}_{c,(\lambda,\eta)}^f = \bigcap_{\lambda' < \lambda} \bigcap_{\eta' < \eta} \bigcup_{K > 0} \{f_c | f_c \text{ solves } A_{(\lambda',\eta',K)}^f\} \quad (39)$$

with a similar definition for $\mathcal{C}_{c,(\lambda,\eta)}^{nf}$. The necessity result of Theorem 3.3 establishes that

$$\mathcal{C}_{c,(\lambda,\eta)}^{nf} \subseteq \mathcal{C}_{c,(\lambda,\eta)}^f \subseteq \mathcal{C}_{a,(\log_2 \lambda, \eta \log_2 \lambda)}^f$$

while Theorem 4.2 establishes the other direction for the case of feedback:

$$\mathcal{C}_{c,(\lambda,\eta)}^f \subseteq \mathcal{C}_{c,(\lambda,\eta)}^{nf} = \mathcal{C}_{a,(\log_2 \lambda, \eta \log_2 \lambda)}^{nf} \quad (40)$$

Meanwhile without feedback and restricting to the set of finite output alphabet channels (ie. where the range of f_c has finite cardinality.) denoted \mathcal{C}_{fin} , Theorem 5.3 implies:

$$\mathcal{C}_{a,(\log_2 \lambda, \eta \log_2 \lambda)}^f \cap \mathcal{C}_{\text{fin}} \subseteq \mathcal{C}_{c,(\lambda,\eta)}^{nf} \cap \mathcal{C}_{\text{fin}}$$

Combining with (40) gives the following result for finite output alphabet channels:

$$\mathcal{C}_{c,(\lambda,\eta)}^{nf} \cap \mathcal{C}_{\text{fin}} = \mathcal{C}_{c,(\lambda,\eta)}^f \cap \mathcal{C}_{\text{fin}} = \mathcal{C}_{a,(\log_2 \lambda, \eta \log_2 \lambda)}^{nf} \cap \mathcal{C}_{\text{fin}} \quad (41)$$

Finally, notice how the mapping from (λ, η) to (R, α) is one-to-one and onto. By setting $\lambda = 2^R$ and $\eta = \frac{\alpha}{R}$ it is possible to translate in the opposite direction and this does provide some additional insight. For example, in the anytime communication problem, it is clear that increasing R from 2 to 3 while keeping α constant at 6 results in a harder problem. When translated to stabilization, without the results established here, it is far from obvious that the equivalent move from $\lambda = 4$ to $\lambda = 8$ with a simultaneous drop in the required η from 3 to 2 is also a move in a fundamentally harder direction.

D. Discussion

Traditionally, this hierarchy of communication problems had not been explored since there were apparently only two interesting levels: problems equivalent to classical Shannon communication and those equivalent to zero-error communication. Anytime communication problems are intermediate between the two. Though feedback anytime communication problems are interesting on their own, the equivalence with feedback stabilization makes them even more fundamental.

It is interesting to consider where Schulman's interactive computation problems fit in this sort of hierarchy. Because a constant factor slowdown is permitted by the asymptotics, such problems of interactive computation do not distinguish between channels of different Shannon capacity. In the language of this section, this means that Shannon communication problems are harder than those of interactive computation considered in [3].

Furthermore, the noisy channel definition given here can be extended to include channels with memory. Simply make the current channel output depend on all the current and past V_t and Y_t . In that case, (40) will continue to hold. Since the finite-output alphabet constructions never needed memorylessness, (41) will also hold.

The constructive nature of the proofs for the underlying theorems makes them akin to the "reductions" used in theoretical computer science to show that two problems belong to the same complexity class. They are direct translations at the level of problems and solutions. In contrast, the classical separation results go through the mutual information characterization of $R(D)$ and C . It would be interesting to study a suitable analog of (36) for channels with memory. Feedback can now increase the capacity so the with-feedback and feedback-free problems are no longer equivalent. However, it would be nice to see a direct reduction of Shannon's communication problem to an estimation problem that encompasses such cases as well. The asymptotic equivalence situation is likely even richer in the multiuser setting where traditional separation theorems do not hold.

ACKNOWLEDGMENTS

The authors would like to thank Mukul Agarwal, Shashibhushan Borade, Devavrat Shah, and Lav Varshney for comments on earlier versions of this paper. We thank Nicola Elia for several constructive discussions about the subject matter of this paper and Sekhar Tatikonda for many discussions over a long period of time which have influenced this work in important ways. Finally, we thank the anonymous reviewers for a careful reading of the paper and helpful feedback.

REFERENCES

- [1] R. G. Gallager, *Information Theory and Reliable Communication*. New York, NY: John Wiley, 1971.
- [2] S. Vembu, S. Verdu, and Y. Steinberg, "The source-channel separation theorem revisited," *IEEE Trans. Inform. Theory*, vol. 41, no. 1, pp. 44–54, Jan. 1995.
- [3] L. J. Schulman, "Coding for interactive communication," *IEEE Trans. Inform. Theory*, vol. 42, no. 6, pp. 1745–1756, Nov. 1996.
- [4] S. Rajagopalan, "A coding theorem for distributed computation," Ph.D. dissertation, University of California, Berkeley, 1994.
- [5] C. E. Shannon, "Coding theorems for a discrete source with a fidelity criterion," *IRE National Convention Record*, vol. 7, no. 4, pp. 142–163, 1959.
- [6] A. Sahai, "Why block length and delay are not the same thing," *IEEE Trans. Inform. Theory*, To be submitted. [Online]. Available: <http://www.eecs.berkeley.edu/~sahai/Papers/FocusingBound.pdf>
- [7] Y. C. Ho, M. P. Kastner, and E. Wong, "Teams, signaling, and information theory," *IEEE Trans. Automat. Contr.*, vol. 23, no. 2, pp. 305–312, Apr. 1978.
- [8] S. Tatikonda, "Control under communication constraints," Ph.D. dissertation, Massachusetts Institute of Technology, Cambridge, MA, 2000.
- [9] S. Tatikonda and S. K. Mitter, "Control under communication constraints," *IEEE Trans. Automat. Contr.*, vol. 49, no. 7, pp. 1056–1068, July 2004.
- [10] G. N. Nair and R. J. Evans, "Communication-limited stabilization of linear systems," in *Proceedings of the 39th IEEE Conference on Decision and Control*, Sydney, Australia, Dec. 2000, pp. 1005–1010.
- [11] —, "Stabilizability of stochastic linear systems with finite feedback data rates," *SIAM Journal on Control and Optimization*, vol. 43, no. 2, pp. 413–436, July 2004.
- [12] S. Tatikonda and S. K. Mitter, "Control over noisy channels," *IEEE Trans. Automat. Contr.*, vol. 49, no. 7, pp. 1196–1201, July 2004.

- [13] A. Sahai, S. Tatikonda, and S. K. Mitter, "Control of LQG systems under communication constraints," in *Proceedings of the 1999 American Control Conference*, San Diego, CA, June 1999, pp. 2778–2782.
- [14] S. Tatikonda, A. Sahai, and S. K. Mitter, "Stochastic linear control over a communication channel," *IEEE Trans. Automat. Contr.*, vol. 49, no. 9, pp. 1549–1561, Sept. 2004.
- [15] R. Bansal and T. Basar, "Simultaneous design of measurement and control strategies for stochastic systems with feedback," *Automatica*, vol. 45, no. 5, pp. 679–694, Sept. 1999.
- [16] A. Sahai, "Evaluating channels for control: Capacity reconsidered," in *Proceedings of the 2000 American Control Conference*, Chicago, CA, June 2000, pp. 2358–2362.
- [17] N. C. Martins, N. Elia, and M. A. Dahleh, "Stabilization of uncertain systems in the presence of a stochastic digital link," *IEEE Trans. Automat. Contr.*, submitted for publication.
- [18] N. Elia, "When Bode meets Shannon: control-oriented feedback communication schemes," *IEEE Trans. Automat. Contr.*, vol. 49, no. 9, pp. 1477–1488, Sept. 2004.
- [19] —, "Remote stabilization over fading channels," *Systems and Control Letters*, vol. 54, no. 3, pp. 239–249, Mar. 2005.
- [20] S. Yuksel and T. Basar, "Achievable rates for stability of LTI systems over noisy forward and feedback channels," in *Proceedings of the 2005 Conference on Information Sciences and Systems*, Baltimore, MD, Mar. 2005, paper 12.
- [21] R. Gray, "Information rates of autoregressive processes," *IEEE Trans. Inform. Theory*, vol. 16, no. 4, pp. 412–421, July 1970.
- [22] T. Berger, "Information rates of Wiener processes," *IEEE Trans. Inform. Theory*, vol. 16, no. 2, pp. 134–139, Mar. 1970.
- [23] A. Sahai and S. K. Mitter, "Source coding and channel requirements for unstable processes," *IEEE Trans. Inform. Theory*, To be submitted. [Online]. Available: <http://www.eecs.berkeley.edu/~sahai/Papers/anytime.pdf>
- [24] T. Berger, *Rate Distortion Theory*. Englewood Cliffs, NJ: Prentice-Hall, 1971.
- [25] C. E. Shannon, "The zero error capacity of a noisy channel," *IEEE Trans. Inform. Theory*, vol. 2, no. 3, pp. 8–19, Sept. 1956.
- [26] V. Borkar and S. K. Mitter, "LQG control with communication constraints," in *Communications, Computation, Control, and Signal Processing: a Tribute to Thomas Kailath*. Norwell, MA: Kluwer Academic Publishers, 1997, pp. 365–373.
- [27] G. D. Forney, "Convolutional codes II. maximum-likelihood decoding," *Information and Control*, vol. 25, no. 3, pp. 222–266, July 1974.
- [28] M. V. Burnashev, "Data transmission over a discrete channel with feedback, random transmission time," *Problemy Perdachi Informatsii*, vol. 12, no. 4, pp. 10–30, Oct./Dec. 1976.
- [29] R. L. Dobrushin, "An asymptotic bound for the probability error of information transmission through a channel without memory using the feedback," *Problemy Kibernetiki*, vol. 8, pp. 161–168, 1962.
- [30] A. Sahai and Q. Xu, "The anytime reliability of constrained packet erasure channels with feedback," in *Proceedings of the 42nd Allerton Conference on Communication, Control, and Computing*, Monticello, IL, Sept. 2004, pp. 200–209.
- [31] B. Sinopoli, L. Schenato, M. Franceschetti, K. Poolla, M. Jordan, and S. Sastry, "Kalman filtering with intermittent observations," *IEEE Trans. Automat. Contr.*, vol. 49, no. 9, pp. 1453–1464, Sept. 2004.
- [32] X. Liu and A. J. Goldsmith, "Kalman filtering with partial observation losses," *IEEE Trans. Automat. Contr.*, submitted for publication.
- [33] A. Sahai, "Any-time information theory," Ph.D. dissertation, Massachusetts Institute of Technology, Cambridge, MA, 2001.
- [34] H. T. Simsek, "Anytime channel coding with feedback," Ph.D. dissertation, University of California, Berkeley, 2004.
- [35] I. M. Jacobs and E. R. Berlekamp, "A lower bound to the distribution of computation for sequential decoding," *IEEE Trans. Inform. Theory*, vol. 13, no. 2, pp. 167–174, Apr. 1967.
- [36] J. P. M. Schalkwijk and T. Kailath, "A coding scheme for additive noise channels with feedback – I: No bandwidth constraint," *IEEE Trans. Inform. Theory*, vol. 12, no. 2, pp. 172–182, Apr. 1966.
- [37] J. P. M. Schalkwijk, "A coding scheme for additive noise channels with feedback – II: Band-limited signals," *IEEE Trans. Inform. Theory*, vol. 12, no. 2, pp. 183–189, Apr. 1966.
- [38] A. J. Kramer, "Improving communication reliability by use of an intermittent feedback channel," *IEEE Trans. Inform. Theory*, vol. 15, no. 1, pp. 52–60, Jan. 1969.
- [39] A. Sahai and Q. Xu, "The anytime reliability of the AWGN+erasure channel with feedback," in *Proceedings of the 42nd Allerton Conference on Communication, Control, and Computing*, Monticello, IL, Sept. 2004, pp. 300–309.
- [40] A. Sahai, S. Avestimehr, and P. Minero, "Anytime communication over the Gilbert-Eliot channel with noiseless feedback," in *Proc. IEEE International Symposium on Information Theory*, Adelaide, Australia, Sept. 2005, pp. 1783–1787.
- [41] H. S. Witsenhausen, "Separation of estimation and control for discrete time systems," *Proceedings of the IEEE*, vol. 59, no. 11, pp. 1557–1566, Nov. 1971.
- [42] —, "A counterexample in stochastic optimum control," *SIAM Journal on Control*, vol. 6, no. 1, pp. 131–147, Jan. 1968.
- [43] S. K. Mitter and A. Sahai, "Information and control: Witsenhausen revisited," in *Learning, Control and Hybrid Systems: Lecture Notes in Control and Information Sciences 241*, Y. Yamamoto and S. Hara, Eds. New York, NY: Springer, 1999, pp. 281–293.
- [44] T. Cover and J. Thomas, *Elements of Information Theory*. New York, NY: John Wiley, 1991.
- [45] G. D. Forney, "Convolutional codes III. sequential decoding," *Information and Control*, vol. 25, no. 3, pp. 267–297, July 1974.
- [46] F. Jelinek, "Upper bounds on sequential decoding performance parameters," *IEEE Trans. Inform. Theory*, vol. 20, no. 2, pp. 227–239, Mar. 1974.
- [47] A. Sahai and H. Palaiyanur, "A simple encoding and decoding strategy for stabilization over discrete memoryless channels," in *Proceedings of the Allerton Conference on Communication, Control, and Computing*, Monticello, IL, Sept. 2005. [Online]. Available: <http://www.eecs.berkeley.edu/~sahai/Papers/allerton05-sequential-stabilization.pdf>
- [48] M. Costa, "Writing on dirty paper," *IEEE Trans. Inform. Theory*, vol. 29, no. 3, pp. 439–441, May 1983.
- [49] A. Sahai, "Anytime coding on the infinite bandwidth AWGN channel: a sequential semi-orthogonal code," in *Proceedings of the 2005 Conference on Information Sciences and Systems*, Baltimore, MD, Mar. 2005, paper 196. [Online]. Available: <http://www.eecs.berkeley.edu/~sahai/Papers/ciss05wideband.pdf>
- [50] —, "Stabilization over discrete memoryless and wideband channels using nearly memoryless observations," submitted to the 44th IEEE Conference on Decision and Control, 2005. [Online]. Available: <http://www.eecs.berkeley.edu/~sahai/Papers/memorylesscontrol.pdf>
- [51] J. E. Hopcroft and J. D. Ullman, *Introduction to Automata Theory, Languages, and Computation*. Reading, MA: Addison-Wesley, 1979.