Control Over Gaussian Channels With and Without Source–Channel Separation

Anatoly Khina, *Member, IEEE*, Elias Riedel Gårding, Gustav M. Pettersson, Victoria Kostina, *Member, IEEE*, and Babak Hassibi, *Member, IEEE*

***Abstract*—We consider the problem of controlling an unstable linear plant with Gaussian disturbances over an additive white Gaussian noise channel with an average transmit power constraint, where the signaling rate of communication may be different from the sampling rate of the underlying plant. Such a situation is quite common since sampling is done at a rate that captures the dynamics of the plant and that is often lower than the signaling rate of the communication channel. This rate mismatch offers the opportunity of improving the system performance by using coding over multiple channel uses to convey a single control action. In a traditional, separation-based approach to source and channel coding, the analog message is first quantized down to a few bits and then mapped to a channel codeword whose length is commensurate with the number of channel uses per sampled message. Applying the separation-based approach to control meets its challenges: first, the quantizer needs to be capable of zooming in and out to be able to track unbounded system disturbances, and second, the channel code must be capable of improving its estimates of the past transmissions exponentially with time, a characteristic known as anytime reliability. We implement a separated scheme by leveraging recently developed techniques for control over quantized-feedback channels and for efficient decoding of anytime-reliable codes. We further propose an alternative, namely, to perform analog joint source–channel coding, by this avoiding the digital domain altogether. For the case where the communication signaling rate is twice the sampling rate, we employ analog linear repetition as well as Shannon–Kotel'nikov maps to show a significant improvement in stability margins and linear-quadratic costs over separation-based schemes. We conclude that such analog coding performs better than separation, and can stabilize all moments as well as guarantee almost-sure stability.**

***Index Terms*—Channel coding, combined source–channel coding, Gaussian channel, Lloyd–Max algorithm, networked control systems, quantization, tree codes.**

Manuscript received August 11, 2018; accepted November 3, 2018. Date of publication April 19, 2019; date of current version August 28, 2019. This work was supported by the European Union's Horizon 2020 research and innovation program under the Marie Skłodowska-Curie Grant 708932. The work of E. Riedel Gårding was supported by the National Science Foundation (NSF) under Grant CCF-1566567 through the SURF program. The work of G. M. Pettersson was supported by The Boeing Company under the SURF program. The work of V. Kostina was supported in part by the NSF under Grant CCF-1566567. The work of B. Hassibi was supported in part by the NSF under Grant CNS-0932428, Grant CCF-1018927, Grant CCF-1423663, and Grant CCF-1409204, by a grant from Qualcomm Inc., by NASA's Jet Propulsion Laboratory through the President and Director's Fund, and by King Abdullah University of Science and Technology. This paper was presented in part at the IEEE Conference on Decision and Control, Las Vegas, NV, USA, Dec. 2016. This work was done in part while A. Khina and V. Kostina were visiting the Simons Institute for the Theory of Computing. Recommended by Associate Editor K. Kashima. *(Corresponding author: Anatoly Khina.)*

A. Khina was with the Department of Electrical Engineering, California Institute of Technology, Pasadena, CA 91125 USA. He is now with the Department of Electrical Engineering—Systems, Tel Aviv University, Tel Aviv 6997801, Israel (e-mail: anatolyk@eng.tau.ac.il).

E. R. Gårding was with the Department of Applied Mathematics and Theoretical Physics, Centre for Mathematical Sciences, University of Cambridge, Cambridge CB3 0WA, U.K. He is now with the School of Engineering Sciences, KTH Royal Institute of Technology, SE-100 44 Stockholm, Sweden (e-mail: eliasrg@kth.se).

G. M. Pettersson is with the School of Engineering Sciences, KTH Royal Institute of Technology, SE-100 44 Stockholm, Sweden (e-mail: gupet@kth.se).

V. Kostina is with the Department of Electrical Engineering, California Institute of Technology, Pasadena, CA 91125 USA (e-mail: vkostina@caltech.edu).

B. Hassibi is with the Department of Electrical Engineering, California Institute of Technology, Pasadena, CA 91125 USA (e-mail: hassibi@caltech.edu).

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/TAC.2019.2912255

I. Introduction

THE current technological era of ubiquitous wireless connectivity and the *Internet of Things* exhibits an ever-growing demand for new and improved techniques for stabilizing *cyber-physical* and *networked control systems*, which as a result have been the subject of intense recent investigations [1]–[4]. Unlike traditional control with colocated plant, observer and controller, the components of such systems are separated by unreliable communication links. In many of these systems, the rate at which the output of the plant is sampled and observed, as well as the rate at which control inputs are applied to the plant, is different from the signaling rate with which communication occurs. The rate at which the plant is sampled and controlled is often governed by how fast the dynamics of the plant is, whereas the signaling rate of the communication depends on the bandwidth available, the noise levels, etc. As a result, there is no inherent reason why these two rates should be related and, in fact, the communication signaling rate is almost always higher than the control sampling rate.

This latest fact clearly gives us the opportunity to improve the performance of the system by conveying the information about each sampled output of the plant, and/or each control signal, through *multiple* uses of the communication channel.

The standard information-theoretic approach suggests quantizing the analog messages (the sampled output or state signal) and then protecting the quantized bits with a channel error-correcting code whose block length is commensurate with the

number of channel uses available per sample. This approach relies on the *source–channel separation principle*, which proffers that quantization of the messages and channel coding of the quantized bits can be done independently of one another.

Nonetheless, while source–channel separation-based schemes become optimal in communication systems where large blocks of the message and the channel code are processed together (necessitating noncausal knowledge of all the message signals and entailing large delays)—a celebrated result [5], [6, Ch. 3.9]—it is not true for control systems that require real-time (*low-delay*) communication of causally available messages. Furthermore, since any error made in the past is magnified in each subsequent time step due to the unstable nature of the plant, the source–channel separation principle requires a stronger notion of error protection, termed *anytime reliability* by Sahai and Mitter [7]. Anytime reliability guarantees that the error probability of causally encoded quantized ("information") bits decays faster than the inflation factor at each step. Sahai and Mitter [7] further observed that *anytime-reliable codes* have a natural *tree code* structure reminiscent of the codes developed by Schulman [8] for the related problem of interactive communication.

Sukhavasi and Hassibi [9] showed that anytime reliability can be guaranteed with high probability by concentrating on the family of linear time-invariant (LTI) codes and choosing their coefficients at random. Unfortunately, maximum-likelihood (ML) (optimum) decoding of tree codes is infeasible.[1] To overcome this problem, a sequential decoder [10], [11, Ch. 6.4], [12, Ch. 10], [13, Sec. 6.9], [14, Ch. 6], [15, Ch. 6] for tree codes was proposed in [16] and was shown to achieve anytime reliability with high probability while maintaining bounded expected decoding complexity, albeit with some loss of performance.

Tree codes transform the control task over a noisy channel to that over a noiseless channel with finite-capacity C, implying that the channel code needs to be supplemented with an adequate *fixed-rate quantizer* (a.k.a. *fixed-length lossy source coder*). Such a quantizer compresses the analog signal to packets of exactly C bits to be communicated from the observer to the controller at every time step. As unstable systems with disturbances that have distributions with unbounded support cannot be stabilized by a static quantizer [17, Sec. III-A], adaptive uniform and logarithmic quantizers that establish stabilizability guarantees were devised by Yüksel [18] and Minero *et al.* [19] (and others), respectively. A greedy algorithm that re-calculates, at every time step, the probability density function (PDF) of the source conditioned on the previously transmitted packets (*a la* sequential Bayesian filtering [20]) and applies Lloyd–Max quantization [21, Ch. 6] with respect to this PDF, was proposed by Bao *et al.* [22] (albeit without any optimality claims). For the scenario of Gaussian disturbances and scalar measurements, this algorithm has been recently shown to be greedily optimal in [23], i.e., it minimizes the linear-quadratic cost at each time step; it was further shown there to be nearly globally optimal.

An obvious alternative strategy to separated source/channel coding is to simply repeat the transmitted (analog) signal—this adds a linear factor to the signal-to-noise ratio (SNR) (3 dB for a single repetition). This strategy maps the analog control signals directly into analog communication signals, avoiding the digital domain altogether, and can therefore be viewed as a simple instance of joint source–channel coding (JSCC) [24].

Surprisingly, in the Gaussian rate-matched case, in which one additive white Gaussian noise (AWGN) channel use is available per one white Gaussian source sample, a simple amplifier achieves the Shannon limit with zero delay [25], [26]. The optimality of linear schemes extends further to the case where $K_C > 1$ uses of an AWGN channel with *perfect instantaneous feedback* are available per one white Gaussian source sample [26]–[30], the reason being that a Gaussian source is *probabilistically matched* to a Gaussian channel [31]—uncommon coincidence. Tatikonda and Mitter [32] exploited a special property of the erasure channel with feedback, in which a retransmission scheme attains its capacity without delay. A related example is control over a packet-drop channel, considered by Sinopoli *et al.* [33]. There, a simple retransmission scheme attains the optimum, as long as the packet drop probability is not too high. Coding of Gauss–Markov sources over a packet erasure channel with feedback is studied in [34].

JSCC in the absence of probabilistic matching is challenging. In the Gaussian rate-mismatched case with no communication feedback, in which $K_C > 1$ AWGN channel uses are available per one source sample, repetitive transmission of the source sample is suboptimal. Non-linear mappings are known to achieve better performance, as noted originally by Shannon [35] and Kotel'nikov [36], and in the context of control—in [37]–[39] (and is akin to Witsenhausen's celebrated counterexample [40]).

In this paper, we concentrate on the simple case of stabilizing a scalar discrete-time linear quadratic Gaussian (LQG) control system over an AWGN channel with $K_C = 2$ channel uses per control sample, with a fixed SNR. As we show in the sequel, this SNR imposes an upper limit on the size of the maximum unstable eigenvalue of the plant that can be stabilized.

We develop end-to-end separation-based and JSCC schemes and compare their LQG costs, as well as the minimum required SNR for stabilizing the system; to the best of our knowledge, this is the first time the latter is devised for the scenario of LQG control over an AWGN channel (cf., [7]), whereas the former (in the absence of instantaneous feedback) was offered in the conference version of this paper [41]. We show that JSCC schemes achieve far better performance while requiring far less computational and memory resources. We further observe that an inherent advantage of JSCC schemes is that they allow a graceful improvement in performance with the SNR, while the performance of separation-based schemes saturates due to their digital nature. Moreover, whereas separation-based schemes can guarantee only a finite number of bounded moments, certain JSCC schemes, e.g., linear ones (repetition-based included), can stabilize all moments (as well as guarantee almost-sure stability).

Another important product of this paper is the implementation and simulation of all the aforementioned schemes and algorithms along with a comparison of their performance; the Python 3 code is available online in [42].

An outline of the rest of the paper is as follows. We formulate the problem setup in Section II. Three different ingredients that are used to construct the schemes for control over noisy channels of this paper, namely, a quantizer for control over a noiseless finite-rate channel, an anytime-reliable tree code, and an Archimedean bi-spiral-based JSCC map, are described in Sections III, IV, and V, respectively. They are subsequently used

[1]Except over erasure channels, over which ML decoding amounts to solving linear equations [9].

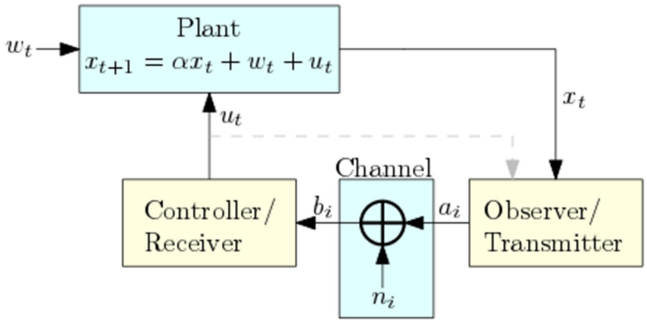

Fig. 1. Scalar control system with an AWGN driving disturbance and an AWGN communication channel. The dashed line represents the assumption that the past control signals are available at the transmitter.

in Sections VI and VII to develop source–channel separation-based and JSCC-based schemes for LQG control over an AWGN channel, and are compared in terms of their LQG cost in Section VIII. We conclude the paper with Section IX, by discussing the principal differences between the proposed schemes along with possible extensions.

II. Problem Setup

We now formulate the control–communications setting that will be treated in this paper, depicted in Fig. 1. We concentrate on the simple case of a scalar linear fully observable system. In contrast to traditional control settings, the observer and the controller are not colocated in this setting, and are connected instead via a scalar AWGN channel.

Remark 2.1: The model and solutions can be extended to more complex cases of vector states and channels; see Section IX-B.

The control and transmission duration spans the time interval $[T] \triangleq \{1, \dots, T\}$.

Plant: A scalar discrete-time linear system dynamics

$$x_{t+1} = \alpha x_t + w_t + u_t, \qquad t \in [T-1], \tag{1}$$

where x_t is the (scalar) state at time t, w_t is an AWGN of power W, α is a known scalar satisfying $|\alpha| > 1$, and u_t is the control signal. We further assume that $x_0 = 0$ for simplicity.[2]

Channel: We assume $K_C \in \mathbb{N}$ channel uses are available per each control sample. Hence, at each time instant t, we can use the channel

$$b_{t;i} = a_{t;i} + n_{t;i}, \qquad i \in [K_C], \qquad t \in [T-1] \tag{2}$$

K_C times, where $b_{t;i}$ is the ith channel output corresponding to control sample t, $a_{t;i}$ is the corresponding channel input subject to a unit power constraint

$$\mathbb{E}\left[a_{t;i}^2\right] \le 1, \tag{3}$$

and $n_{t;i}$ is an AWGN of power $1/\mathrm{SNR}$.[3] We collect all the K_C channel uses in a column vector and denote it, with a slight abuse of notation, by $\boldsymbol{a}_t \triangleq (a_{t;1}, \dots, a_{t;K_C})^{\mathrm{T}}$, where "T" denotes the transpose operation. The corresponding channel output vector is denoted by $\boldsymbol{b}_t \triangleq (b_{t;1}, \dots, b_{t;K_C})^{\mathrm{T}}$.

Causal Transmitter: At time t, generates K_C channel inputs by applying a causal function $\mathcal{E}_t : \mathbb{R}^t \times \mathbb{R}^{t-1} \to \mathbb{R}^{K_C}$ to the measured states $x^t \triangleq (x_1, \dots, x_t)$ and all past control signals $u^{t-1} \triangleq (u_1, \dots, u_{t-1})$:

$$\boldsymbol{a}_t = \mathcal{E}_t\left(x^t, u^{t-1}\right), \tag{4}$$

where the input is subject to an average power constraint (3).

Remark 2.2: In this paper, we assume that the observer/transmitter knows all past control signals u^{t-1}; for a discussion of the scenario when such information is not available at the observer, see Section IX-C.

Causal Receiver: At time t, observes K_C channel outputs and generates a control signal u_t by applying a causal function $\mathcal{D}_t : \mathbb{R}^{tK_C} \to \mathbb{R}$ to all the available channel outputs

$$u_t = \mathcal{D}_t\left(\boldsymbol{b}^t\right), \tag{5}$$

where $\boldsymbol{b}^t \triangleq (\boldsymbol{b}_1^{\mathrm{T}}, \dots, \boldsymbol{b}_t^{\mathrm{T}})^{\mathrm{T}}$.

Cost: Similarly to the classical LQG control setting (in which the controller and the observer are colocated), we wish to minimize the average-stage LQG cost at the time horizon T,

$$\bar{J}_T \triangleq \frac{1}{T}\mathbb{E}\left[Q_T x_T^2 + \sum_{t=1}^{T-1}\left(Q_t x_t^2 + R_t u_t^2\right)\right],$$

for known non-negative weights $\{Q_t\}$ and $\{R_t\}$, by designing appropriate operations at the observer, which also plays the role of the transmitter over the channel (2); and the controller, which also serves as the receiver over the channel (2).

For the important special case of fixed parameters,

$$Q_t \equiv Q, R_t \equiv R,$$

we further define the infinite-horizon cost

$$\bar{J}_\infty \triangleq \lim_{T\to\infty} \bar{J}_T, \tag{6}$$

assuming the limit exists.

We recall next recently developed schemes for quantization and channel coding for control as well as results from information theory for JSCC design with low delay.

III. Control With Noiseless Finite-Rate Feedback

In this section, we consider the model of Section II with the AWGN channel (2) replaced with a noiseless channel of finite capacity C, depicted in Fig. 2. That is, in this case, the channel, transmitter, and receiver are as follows.

Channel: At time t, a packet $\ell_t \in \{0, \dots, 2^C - 1\}$ is sent over a noiseless channel of capacity C, meaning that the receiver obtains ℓ_t at time t.

Transmitter: The function $\mathcal{E}_t$ (4), in this case, has a discrete codomain $\{0, \dots, 2^C - 1\}$ (with no power constraints).

Receiver: The domain of $\mathcal{D}_t$ (5) is $\{0, \dots, 2^C - 1\}$.

Thus, the transmitter–receiver design amounts, in this case, to fixed-length sequential quantization.

A recent result [23] shows that an adaptive quantizer that successively calculates the PDF of x_t given ℓ^t and applies Lloyd–Max quantization with respect to it is greedily optimal and close-to-globally optimal whenever f_w is log-concave.

[2]This assumption can be replaced with an initial state x_0 with a Gaussian PDF.

[3]This representation is without loss of generality since the case of an average power P_C and noise power N can always be transformed to an equivalent channel with average power 1 and noise power $N/P_C \triangleq 1/\mathrm{SNR}$ by multiplying both sides of (2) by $1/\sqrt{P_C}$.

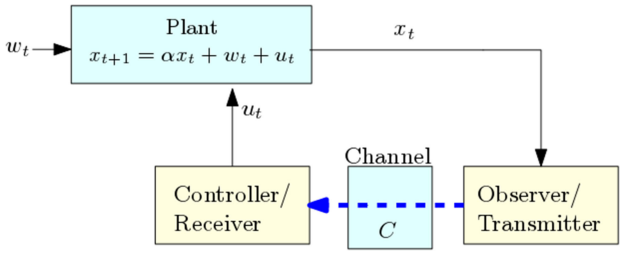

Fig. 2. Scalar control system with a driving white Gaussian disturbance and a noiseless finite-capacity channel ("bit pipe"). The dashed line represents a bit-pipe of capacity C.

Definition 3.1 (Log-concave function; see [43]): A function $f : \mathbb{R} \to \mathbb{R}_{\geq 0}$ is said to be log-concave if its logarithm $\log \circ f$ is concave; we use the extended definition that allows $f(x)$ to assign zero values, i.e., $\log f(x) \in \mathbb{R} \cup \{-\infty\}$ is an extended real-value function that can take the value $-\infty$.

We recall the Lloyd–Max algorithm and its optimality guarantees in Section III-A; the appropriate adaptive networked control system is described in Section III-B.

A. Quantizer Design

The exposition in this section is based on [21, Part II].

Definition 3.2 (Quantizer): A scalar quantizer $\mathcal{Q}$ of rate C is described by an encoder $\mathcal{E}_\mathcal{Q} : \mathbb{R} \to \{0, \ldots, 2^C - 1\}$ and a decoder $\mathcal{D}_\mathcal{Q} : \{0, \ldots, 2^C - 1\} \to \{c[0], \ldots, c[2^C - 1]\} \subset \mathbb{R}$. With a slight abuse of notation, we shall define the quantization operation $\mathcal{Q} : \mathbb{R} \to \{c[0], \ldots, c[2^C - 1]\}$ as the composition of the encoding and decoding operations: $\mathcal{Q} = \mathcal{D}_\mathcal{Q} \circ \mathcal{E}_\mathcal{Q}$.[4] The *reproduction points* are assumed, without loss of generality, to be ordered:[5]

$$c[0] < c[1] < \cdots < c[2^C - 1].$$

We denote by $\mathcal{I}[\ell]$ the collection of all points that are mapped to index ℓ (equivalently to the reproduction point $c[\ell]$):

$$\mathcal{I}[\ell] \triangleq \{x | x \in \mathbb{R}, \mathcal{E}_\mathcal{Q} = \ell\} = \{x | x \in \mathbb{R}, \mathcal{Q} = c[\ell]\}.$$

We shall concentrate on the following class of quantizers.

Definition 3.3 (Regular quantizer): A scalar quantizer is *regular* if every cell $\mathcal{I}[\ell]$ $(\ell = 0, \ldots, 2^C - 1)$ is a contiguous interval that contains its reproduction point $c[\ell]$:

$$c[\ell] \in \mathcal{I}[\ell] = [p[\ell], p[\ell + 1]), \qquad \ell = 0, \ldots, 2^C - 1,$$

where $p \triangleq \{p[0], \ldots, p[2^C]\}$ is the set of *partition levels*—the cells boundaries. Hence, a regular scalar quantizer can be represented by the input partition-level set and the reproduction-point set $c \triangleq \{c[0], \ldots, c[2^C - 1]\}$.

Cost: The cost we wish to minimize is the mean squared error distortion between a source w with a given log-concave PDF f_w and its quantization $\mathcal{Q}(w)$

$$D \triangleq \mathbb{E}\left[(w - \mathcal{Q}(w))^2\right] \tag{7a}$$

$$= \sum_{\ell=0}^{2^C - 1} \int_{p[\ell]}^{p[\ell+1]} (w - c[\ell])^2 f_w(w) dw. \tag{7b}$$

Denoted by D^*, the minimal achievable distortion D; the *optimal quantizer* is the one that achieves D^*.

Remark 3.1: Since f_w is log-concave, it is continuous [43]. Hence, the inclusion/exclusion of the boundary points of each cell does not affect the distortion of the quantizer, meaning that the boundary points can be broken systematically.

Remark 3.2: We concentrate, in this paper, on input PDFs with an infinite support. Consequently, $p[0] = -\infty$ and $p[2^C] = \infty$, and the leftmost and rightmost intervals are open.

The optimal quantizer satisfies the following necessary conditions [21, Ch. 6.2].

Proposition 3.1 (Nearest neighbor condition): For a fixed reproduction-point set c (fixed decoder), the partition-level set p (encoder) that minimize the distortion D (7) is

$$p[\ell] = \frac{c[\ell - 1] + c[\ell]}{2}, \qquad \ell = 1, 2, \ldots, 2^C - 1, \tag{8}$$

and $p[0] = -\infty$ and $p[2^C] = \infty$.

Proposition 3.2 (Centroid condition): For a fixed partition-level set p (fixed encoder), the reproduction-point set c (decoder) that minimizes the distortion D (7) is

$$c[\ell] = \mathbb{E}\left[w \,\middle|\, p[\ell] \leq w < p[\ell + 1]\right], \ \ell = 0, \ldots, 2^C - 1. \tag{9}$$

The optimal quantizer must simultaneously satisfy both (8) and (9); iterating between these two necessary conditions gives rise to the Lloyd–Max algorithm.

Algorithm 3.1 (Lloyd–Max): Initialization. Pick an initial reproduction-point set c.

Iteration. Repeat the two steps

1) fix c and set p as in (8),
2) fix p and set c as in (9),

interchangeably, until the decrease in the distortion D per iteration goes below a desired threshold.

Propositions 3.2 and 3.1 suggest that the distortion at every iteration decreases; since the distortion is bounded from below by zero, the Lloyd–Max algorithm is guaranteed to converge to a local optimum.

Unfortunately, multiple local optima may exist in general (e.g., Gaussian mixtures with well-separated components), rendering the algorithm sensitive to the initial choice c.

Nonetheless, sufficient conditions for the existence of a unique global optimum were established in [44]–[46]. These guarantee the convergence of the algorithm to the global optimum for any initial choice of c. An important class of PDFs that satisfy these conditions is that of the log-concave PDFs.

Theorem 3.1 (see [44]–[46]): Let the source PDF f_w be log-concave. Then, the Lloyd–Max algorithm converges to a unique solution that minimizes the quadratic distortion (7).

B. Controller Design

We now describe the optimal greedy control policy, the implementation of which is available in [42, `tree/master/code/separate/control`]. To that end,

[4]The encoder and decoder that give rise to the same quantizer are unique up to a permutation of the labeling of the index ℓ.

[5]If some inequalities are not strict, then the quantizer can be reduced to another quantizer of lower rate.

we make use of the following lemma that extends the separation principle of estimation and control to networked control systems.

Lemma 3.1 (see [47], [29]): The optimal control law is given by

$$u_t = -L_t \hat{x}_t,$$

where

$$L_t = \frac{S_{t+1}}{R_t + S_{t+1}} \alpha \tag{10}$$

is the optimal linear quadratic regulator control gain, $\hat{x}_t \triangleq \mathbb{E}\left[x_t \middle| \ell^t\right]$, and S_t satisfies the dynamic Riccati backward recursion [48] with $S_T = Q_T$ and $S_{T+1} = L_T = 0$:

$$S_t = Q_t + \frac{S_{t+1} R_t}{S_{t+1} + R_t} \alpha^2.$$

Moreover, this controller achieves the cost[6]

$$\bar{J}_T = \frac{1}{T} \sum_{t=1}^{T} \left(S_t W + G_t \mathbb{E}\left[(x_t - \hat{x}_t)^2 \right] \right)$$

with $G_t = S_{t+1}\alpha^2 - S_t + Q_t$.

The optimal greedy algorithm minimizes the estimation distortion $\mathbb{E}\left[(x_t - \hat{x}_t)^2\right]$ at time t, without regard to its effect on future distortions. To that end, at time t, the transmitter and the receiver calculate the PDF of x_t conditioned on ℓ^{t-1} and u^{t-1}, $f_{x_t|\ell^{t-1},u^{t-1}}$, and apply the Lloyd–Max quantizer to this PDF.[7] We refer to $f_{x_t|\ell^{t-1},u^{t-1}}$ and to $f_{x_t|\ell^t,u^{t-1}}$ as the *prior* and *posterior PDFs* , respectively.

Although the optimal greedy algorithm does not achieve global optimality [49], its loss is negligible [23].

Algorithm 3.2 (Optimal greedy control):

Initialization. Both the transmitter and the receiver set

1) $\ell_0 = x_0 = u_0 = 0$.
2) Prior PDF: $f_{x_1|\ell_0,u_0}(x_1|0,0) \equiv f_{x_1}(x)$.

Observer/Transmitter. At time $t \in [T-1]$:

1) Observes x_t.
2) Runs the Lloyd–Max algorithm (Algorithm 3.1) with respect to the prior PDF $f_{x_t|\ell^{t-1},u^{t-1}}$ to obtain the quantizer $\mathcal{Q}_t(x_t)$ of rate C; we denote its partition-level and reproduction-point sets by p_t and c_t, respectively.
3) Quantizes the system state x_t [recall Definition 3.2]:

$$\ell_t = \mathcal{E}_{\mathcal{Q}_t}(x_t),$$
$$\hat{x}_t = \mathcal{Q}_t(x_t) = \mathcal{D}_{\mathcal{Q}_t}(\ell_t).$$

4) Transmits the quantization index ℓ_t.
5) Calculates the posterior PDF:

$$f_{x_t|\ell^t,u^{t-1}}(x_t|\ell^t, u^{t-1}) = \begin{cases} f_{x_t|\ell^{t-1},u^{t-1}}(x_t|\ell^{t-1}, u^{t-1})/\gamma, & x_t \in \mathcal{I}[\ell_t] \\ 0, & \text{otherwise} \end{cases}$$

where $\mathcal{I}[\ell_t] \triangleq [p_t[\ell_t], p_t[\ell_t+1])$ as in Definition 3.3, and $\gamma \triangleq \int_{p_t[\ell_t]}^{p_t[\ell_t+1]} f_{x_t|\ell^{t-1},u^{t-1}}(\alpha|\ell^{t-1}, u^{t-1}) d\alpha$.

6) Calculates the next prior PDF using (1) and $u_t = -L_t \hat{x}_t$

$$f_{x_{t+1}|\ell^t,u^t}(x_{t+1}|\ell^t, u^t) = \frac{1}{|\alpha|} f_{x_t|\ell^t,u^{t-1}}\left(\frac{x_{t+1} - u_t}{\alpha} \middle| \ell^t, u^{t-1} \right) * f_w(x_{t+1})$$

where "$*$" denotes the convolution operation, and the two convolved terms correspond to the PDFs of the quantization error $\alpha(x_t - \hat{x}_t)$ and the disturbance w_t.

Controller/Receiver. At time $t \in [T-1]$:

1) Runs the Lloyd–Max algorithm (see Algorithm 3.1) with respect to the prior PDF $f_{x_t|\ell^{t-1},u^{t-1}}$ as in Step 2 of the observer/transmitter protocol.
2) Receives the index ℓ_t.
3) Reconstructs the quantized value: $\hat{x}_t = \mathcal{D}_{\mathcal{Q}_t}(\ell_t)$.
4) Applies the control actuation $u_t = -L_t \hat{x}_t$ to the system.
5) Calculates the posterior PDF $f_{x_t|\ell^t,u^{t-1}}$ and the next prior PDF $f_{x_{t+1}|\ell^t,u^t}$ as in Steps 5 and 6 of the observer/transmitter protocol.

Theorem 3.2 (see [23]): Let f_w be a Gaussian PDF. Then, Algorithm 3.2 constitutes the optimal greedy control policy.

IV. Anytime-Reliable Codes

We now describe causal error-correcting codes that allow us to meet the assumption of a noiseless finite-capacity channel of Section III over the AWGN channel (2).

Since any decoding mistake is multiplied by α at every time step, and the corresponding second moment (power)—by α^2, the code should have an error probability that decays exponentially with time with an exponent that is greater than α^2;[8] see [7], [9], and [16] for further details and discussion.

We construct such codes, termed anytime-reliable codes [7], for memoryless binary-input output-symmetric (MBIOS) channels and then apply these results for the AWGN channel by employing appropriate digital constellations.

Definition 4.1 (MBIOS channel): A *binary-input channel* is a system with binary input alphabet $\{0,1\}$, output alphabet $\mathcal{Z}$, and two probability transition functions: $q(z|0)$ for input $c=0$ and $q(z|1)$ for input $c=1$. The channel is said to be *memoryless* if the probability distribution of the output depends only on the input at that time and is conditionally independent of previous and future channel inputs and outputs. It is further said to be *output-symmetric* if there exists an involution $\pi : \mathcal{Z} \to \mathcal{Z}$, i.e., a permutation that satisfies $\pi^{-1} = \pi$, such that $q(\pi(z)|0) = q(z|1)$ for all $z \in \mathcal{Z}$.[9]

The encoder and resulting code need to be causal, in our case, due to the sequential nature of the information stream. That is, at time instant t, k new information bits $\boldsymbol{\imath}_t$ are fed to the encoder; the encoder, then, produces n coded bits $\boldsymbol{c}_t$ by encoding all of the available information bits $\boldsymbol{\imath}^t$,

$$\boldsymbol{c}_t = \mathcal{E}_t^{\mathbb{A}}\left(\boldsymbol{\imath}^t\right), \tag{11}$$

[6] We set $R_T = 0$ and $\ell_T = 0$ for the definition of $\hat{x}_T$, as no transmission or control action are performed at time T.

[7] Since u^{t-1} is a deterministic function of ℓ^{t-1}, it suffices to condition on ℓ^{t-1}. However, when incorporating Algorithm 3.2 into a separation-based scheme, this distinction becomes useful since u^{t-1} becomes a function of possibly corrupted channel outputs in this case.

[8] To stabilize higher moments, one needs higher exponents. See the discussion in Section IX-A below.

[9] This also extends to additive noise channels, such as the binary-input AWGN channel.

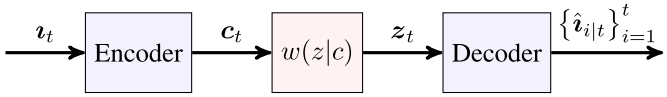

Fig. 3. MBIOS channel with reconstructions of all past information bits.

using an encoding function $\mathcal{E}_t^{\text{A}} : \{0,1\}^{kt} \to \{0,1\}^n$, agreed upon by the encoder and the decoder before transmission.

The sequential encoding operation can be conveniently viewed as advancing over a prefix tree (trie) and the corresponding codes are therefore referred to as tree codes.

At time t, the decoder recovers estimates $\{\boldsymbol{\imath}_{i|t}\}_{i=1}^{t}$ of all the past information bits $\boldsymbol{\imath}^t$ by applying a causal function $\mathcal{D}^{\text{A}} : \mathcal{Z}^{nt} \to \{0,1\}^{kt}$ to all the received channel outputs $\boldsymbol{c}^t$, to produce (see also Fig. 3)

$$\left(\boldsymbol{\imath}_{1|t}, \boldsymbol{\imath}_{2|t}, \ldots, \boldsymbol{\imath}_{t|t}\right) = \mathcal{D}_t^{\text{A}}\left(\boldsymbol{c}^t\right). \tag{12}$$

One is then assigned the task of choosing a sequence of function pairs $\{(\mathcal{E}_t^{\text{A}}, \mathcal{D}_t^{\text{A}}) | t \in \mathbb{N}\}$ that provides anytime reliability. We recall this definition as stated in [9].

Definition 4.2 (Anytime reliability): Define the probability of the first error event at time t happening d steps back (delay) as

$$P_e(t,d) \triangleq P\left(\boldsymbol{\imath}_{t-d} \neq \hat{\boldsymbol{\imath}}_{t-d|t}, \forall \delta > d, \boldsymbol{\imath}_{t-\delta} = \hat{\boldsymbol{\imath}}_{t-\delta|t}\right),$$

where the probability is over the randomness of the information bits $\{\boldsymbol{\imath}_t\}$ and the channel noise. Suppose we are assigned a budget of n channel uses per time step of the evolution of the plant. Then, an encoder–decoder pair is called (R, β) anytime reliable if there exist $A \in \mathbb{R}$ and $d_0 \in \mathbb{N}$, such that

$$P_e(t,d) \leq A 2^{-\beta n d} \quad \forall t, d \geq d_0, \tag{13}$$

where β is called the anytime exponent.

Remark 4.1: The requirement of $d \geq d_0$ in (13) can always be dropped by replacing A by a larger constant. Conversely, A can be replaced with 1 by reducing β by $\epsilon > 0$, however small, and taking a large enough d_0. Nonetheless, we use both A and d_0 in the definition for convenience.

A. LTI Anytime-Reliable Codes Under ML Decoding

Following Sukhavasi and Hassibi [9], we now present an LTI anytime-reliable code ensemble under ML decoding.

When restricted to an LTI ("tree") code, each function ε_t^{A} can be characterized by a set of matrices $\{\mathbf{G}_1, \ldots, \mathbf{G}_t\}$, where $\mathbf{G}_t \in \mathbb{Z}_2^{n \times k}$. The sequence of quantized measurements at time t, $\{\boldsymbol{b}_i\}_{i=1}^{t}$, is encoded as

$$\boldsymbol{c}_t = \mathbf{G}_1 \boldsymbol{\imath}_1 + \mathbf{G}_2 \boldsymbol{\imath}_2 + \cdots + \mathbf{G}_t \boldsymbol{\imath}_t \tag{14}$$

or equivalently in a matrix form

$$\boldsymbol{c} = \mathcal{G}_{n;R} \boldsymbol{\imath},$$

with

$$\mathcal{G}_{n,R} = \begin{bmatrix} \mathbf{G}_1 & \mathbf{0} & \mathbf{0} & \cdots & \cdots \\ \mathbf{G}_2 & \mathbf{G}_1 & \mathbf{0} & \cdots & \cdots \\ \vdots & \vdots & \ddots & \ddots & \cdots \\ \mathbf{G}_t & \mathbf{G}_{t-1} & \cdots & \mathbf{G}_1 & \mathbf{0} \\ \vdots & \vdots & \vdots & \vdots & \ddots \end{bmatrix} \tag{15a}$$

$$\boldsymbol{\imath}^{\text{T}} = \begin{bmatrix} \boldsymbol{\imath}_1^{\text{T}} & \boldsymbol{\imath}_2^{\text{T}} & \cdots & \boldsymbol{\imath}_t^{\text{T}} & \ldots \end{bmatrix} \tag{15b}$$

$$\boldsymbol{c}^{\text{T}} = \begin{bmatrix} \boldsymbol{c}_1^{\text{T}} & \boldsymbol{c}_2^{\text{T}} & \cdots & \boldsymbol{c}_t^{\text{T}} & \ldots \end{bmatrix}. \tag{15c}$$

We now define the random LTI tree code ensemble.

Definition 4.3 (LTI tree code ensemble): An ensemble of LTI tree codes of rate $R = k/n$ that map kt information bits into n bits at every time step t, where the entries in all $\{\mathbf{G}_i\}$ of $\mathcal{G}_{n;R}$ of (15a) are i.i.d. and uniform.

Theorem 4.1 (Error exponent under ML decoding): Let q be an MBIOS channel. Let further $\epsilon > 0$ and $d_0 \in \mathbb{N}$. Then, the probability that a particular code from the random LTI tree code ensemble of Definition IV.3 has an anytime exponent (13) of $E_G(R) - \epsilon$, for all $t \in \mathbb{N}$ and $d > d_0$, under optimal (ML) decoding, is bounded from below by

$$\Pr\left(\bigcap_{t=1}^{\infty} \bigcap_{d=d_0}^{t} \left\{P_e(t,d) \leq 2^{-[E_G(R)-\epsilon]nd}\right\}\right) \geq 1 - \frac{2^{-\epsilon n d_0}}{1 - 2^{-\epsilon n}}$$

where E_G is the block random-coding error exponent [50, Ch. 9], [13, Sec. 5.6], [12, Ch. 7]

$$E_G(R) \triangleq \max_{0 \leq \rho \leq 1} \left[E_0(\rho) - \rho R\right] \tag{16a}$$

$$E_0(\rho) \triangleq -\log \sum_{z \in \mathcal{Z}} \left[\frac{1}{2} q^{\frac{1}{1+\rho}}(z|0) + \frac{1}{2} q^{\frac{1}{1+\rho}}(z|1)\right]^{1+\rho}. \tag{16b}$$

Thus, for any $\epsilon > 0$, however small, this probability can be made arbitrarily close to 1 by taking d_0 to be large enough.

Unfortunately, ML decoding requires searching over all possible codewords—the number of which grows exponentially fast with time—rendering it infeasible except over erasure channels [9]. We therefore turn to sequential decoding, which trades some performance for feasible expected complexity.

B. LTI Anytime-Reliable Codes Under Sequential Decoding

Instead of an exhaustive search over all possible codewords—the complexity of which grows as $\mathcal{O}(2^{kt})$—as is done in ML decoding, one may restrict the search to only the most likely codeword paths, such that their per letter complexity does not grow significantly with time. Such algorithms are known collectively as sequential decoding algorithms.

In this paper, we shall concentrate on one of the two popular variants of this algorithm—the *Stack Algorithm* , the other being the *Fano Algorithm* . The former achieves better time complexity and smaller error probability but is more expensive in terms of memory, compared to the latter. Nonetheless, the anytime exponent of both algorithms is the same; for a treatment of the Fano algorithm, which is similar to the one presented next, the reader is referred to [16].

Algorithm 1: Sequential Decoding Stack Algorithm.

```
Q ← MaxPriorityQueue(node ↦ node.metric)        ▷ Leaf nodes, ordered by Fano metric
Q.push_with_priority(root)
while Q.top.depth < t do                        ▷ Stop at the first sequence to reach full length
    node ← Q.pop()                              ▷ Take the sequence with the largest metric
    for child ∈ node.create_children() do       ▷ Replace it with its 2^k extensions
        Q.push_with_priority(child)
return Q.top.input_sequence()                   ▷ Reconstruct the input sequence by backtracking
```

We next summarize the relevant properties of the stack decoding algorithm when using the generalized Fano metric (see, e.g., [12, Ch. 10]) to compare possible codeword paths

$$M(\boldsymbol{c}_1,\dots,\boldsymbol{c}_N)=\sum_{t=1}^{T} M(\boldsymbol{c}_t) \tag{17a}$$

$$M(\boldsymbol{c}_t)\triangleq \log\frac{q(\boldsymbol{z}_t|\boldsymbol{c}_t)}{\sum_{\boldsymbol{c}'\in\{0,1\}^n}\left(\frac{1}{2}\right)^n q(\boldsymbol{z}_t|\boldsymbol{c}')}-nB \tag{17b}$$

where B is referred to as the metric bias. It penalizes longer paths when the metrics of different-length paths are compared.

In contrast to ML decoding, where at time t, all possible paths (of length kt) are explored to determine the path with the total maximal metric,[10] when using the stack sequential decoding algorithm, a list of partially explored paths is stored in a priority queue, where at each step the path with the highest metric is further explored and replaced with its immediate descendants and their metrics. The stack algorithm is outlined in Algorithm 1 and implemented in [42, `tree/master/code/separate/coding`], [51]; for a detailed description of the stack algorithm (as well as the Fano algorithm and variants thereof), see [11, Ch. 6.4], [12, Ch. 10], [13, Sec. 6.9], [14, Ch. 6], and [15, Ch. 6].

Theorem 4.2 (Error exponent under sequential decoding): Let q be an MBIOS channel. Let further $\epsilon>0$ and $d_0\in\mathbb{N}$. Then, the probability that a particular code from the random LTI tree code ensemble of Definition 4.3 has an anytime exponent (13) of $E_J(R)-\epsilon$, for all $t\in\mathbb{N}$ and $d>d_0$, under sequential stack decoding, is bounded from below by

$$\Pr\left(\bigcap_{t=1}^{\infty}\bigcap_{d=d_0}^{t}\left\{P_e(t,d)\le A2^{-[E_J(R)-\epsilon]nd}\right\}\right)\ge 1-\frac{2^{-\epsilon n d_0}}{1-2^{-\epsilon n}}$$

where E_J is Jelinek's sequential decoding exponent

$$E_J(B,R)\triangleq\max_{0\le\rho\le 1}\frac{\rho}{1+\rho}\Big\{E_0(\rho)+B-(1+\rho)R\Big\},$$

E_0 is given in (16b), and A is finite for $B<E_0(1)$ and is bounded from above by[11]

$$A\le\frac{1-\mathrm{e}^{-t[E_0(\rho)-\rho B]}}{1-\mathrm{e}^{-[E_0(\rho)-\rho B]}}\le\frac{1}{1-\mathrm{e}^{-[E_0(\rho)-\rho B]}}<\infty.$$

Thus, for any $\epsilon>0$, however small, this probability can be made arbitrarily close to 1 by taking d_0 to be large enough.

[10]Note that optimizing (17a) in this case is equivalent to ML decoding.

[11]Note that $E_0(\rho)/\rho$ is a monotonically decreasing function of ρ, therefore $B<E_0(1)$ guarantees that $E_0(\rho)-\rho B>0$.

Since $E_J(B,R)$ is a monotonically increasing function of B, choosing $B=E_0(1)$ maximizes the exponential decay of $\bar{P}_e(d)$ in d.[12] Interestingly, for this choice of bias, we have $E_J(E_0(1),R)=E_G(R)$ whenever $E_G(R)$ is achieved by $\rho=1$ in (16a), i.e., for rates below the critical rate. For other values of ρ, $E_J(E_0(1),R)$ is strictly smaller than $E_G(R)$.

The choice $B=R$, on the other hand, is known to minimize the expected computational complexity (which has a Pareto distribution; see [16], for details), and is therefore a popular choice in practice. Moreover, for rates below the cutoff rate $R<E_0(1)$, the expected number of metric evaluations (17b) at each time instant is finite and does not depend on d, for any $B\le E_0(1)$ [13, Sec. 6.9], [12, Ch. 10]. Thus, the only increase in expected complexity of this algorithm with d comes from an increase in the complexity of evaluating the metric of a single symbol (17b). Since the latter increases (at most) linearly with d, the total complexity of the algorithm grows polynomially in d. Furthermore, for rates above the cutoff rate, $R>E_0(1)$, the expected complexity is known to grow rapidly with the code length for *any metric* [52], implying that the algorithm is applicable only for rates below the cutoff rate $E_0(1)$.

C. Modulation

In order to support the transmission of more than one coded bit per channel use, we modulate the bits using pulse-amplitude modulation (PAM). Specifically, we map every n/K_C consecutive coded bits of $\boldsymbol{c}_t=(c_{t;1},\dots,c_{t;n})^{\mathrm{T}}$ into a constellation point of size 2^k, where k is the number of information bits than need to be conveyed at each time step t. We normalize the constellation to have an average unit power:

$$a_{t;i}=\sqrt{\frac{3}{2^{k-1}-1}}\sum_{j=1}^{k}2^{j-1}(-1)^{c_{t;j+ki}},\quad i\in[K_C]. \tag{18}$$

V. Low-Delay JSCC

In this section, we review known results from information theory for transmitting an i.i.d. zero-mean Gaussian source s_t of power P_S over an AWGN channel (2). Following the problem setup of Section II, we consider the case where $K_C\in\mathbb{N}$ channel uses of (2) are available per each source sample s_t. We suppress the time index t throughout this section.

The goal of the transmitter is to convey the source s to the receiver with a minimal possible average distortion, where the appropriate distortion measure for our case of interest is the mean square error (MSE) distortion.

[12]For finite values of d a lower choice of B may be better, since the constant A might be smaller in this case.

Transmitter: Similarly to (4), it generates K_C channel inputs by applying a function $\mathcal{E}:\mathbb{R}\to\mathbb{R}^{K_C}$ to the source sample s,

$$\boldsymbol{a}=\mathcal{E}(s),$$

where $\boldsymbol{a}$ is subject to an average power constraint (3).

Receiver: Observes the K_C channel outputs $\boldsymbol{b}$ [defined as in Section II], and constructs an estimate $\hat{s}$ of s, by applying a function $\mathcal{D}:\mathbb{R}^{K_C}\to\mathbb{R}$:

$$\hat{s}=\mathcal{D}(\boldsymbol{b}).$$

Cost: The cost, commonly referred to as average distortion in the context of JSCC, is defined by

$$D=\mathbb{E}\left[(s-\hat{s})^2\right],$$

and the corresponding (source) signal-to-distortion ratio (SDR) is defined as

$$\mathrm{SDR}\triangleq\frac{P_S}{D}.$$

Our results here are more easily presented in terms of unbiased errors, as these can be regarded as uncorrelated additive noise in the sequel (when used as part of the developed control scheme). Therefore, we consider the use of (sample-wise) correlation-sense unbiased estimators (CUBE), namely, estimators that satisfy

$$\mathbb{E}\left[s\,(s-\hat{s})\right]=0.$$

We note that any estimator $\hat{s}^B$ can be transformed into a CUBE $\hat{s}$ by multiplying by a suitable constant:

$$\hat{s}=\frac{\mathbb{E}\left[s^2\right]}{\mathbb{E}\left[s\hat{s}^B\right]}\hat{s}^B. \tag{19}$$

For a further discussion of such estimators and their use in communications, the reader is referred to [24].

Shannon's celebrated result [5] states that the minimal achievable distortion, using *any* transmitter–receiver scheme, is dictated, in the case of a Gaussian source, by[13]

$$\frac{1}{2}\log(1+\mathrm{SDR})=R(D)\le K_C C=\tfrac{K_C}{2}\log(1+\mathrm{SNR}), \tag{20}$$

where $R(D)$ is the rate–distortion function of the source and C is the channel capacity [5]; this result remains true even in the presence of feedback—when the channel outputs are available at the transmitter [31, Ch. 1.5]. Thus, the optimal SDR, commonly referred to as *optimum performance theoretically achievable* (OPTA) SDR, is given by

$$\mathrm{SDR}_{\mathrm{OPTA}}=(1+\mathrm{SNR})^{K_C}-1. \tag{21}$$

Although (21) is attainable via separation for K_S source samples and $K_S K_C$ channel uses in the limit of large K_S, it is unknown, in general, how closely (21) can be approached at finite delay. Here, we focus on the scenario of interest to control, namely, the zero-delay case, in which a single Gaussian sample is instantaneously mapped to K_C channel uses.

We next concentrate on the case of $K_C>1$ and perfect instantaneous feedback, in Section V-A. We further treat the case of $K_C=2$ when no feedback is available, in Section V-B.

[13]The rate–distortion function here is written in terms of the unbiased SDR, in contrast to the more common-biased SDR expression $\log(\mathrm{SDR})$.

A. With Feedback

When perfect instantaneous feedback is available, the following simple scheme, due to Elias [26] (see also [31, Ch. 3.5] and the references therein), is known to achieve $\mathrm{SDR}_{\mathrm{OPTA}}$ for $K_C\in\mathbb{N}$.

Scheme V.1 (JSCC with feedback):

Transmitter. At channel use $i\in[K_C]$

1) Calculates the MMSE estimation error of the source s given all past outputs $(b_1,\ldots,b_{i-1})$ (available via the instantaneous feedback):

$$\tilde{s}_{i-1}^{\mathrm{MMSE}}=s-\hat{s}_{i-1}^{\mathrm{MMSE}},$$

where the MMSE estimate $\hat{s}_i$ of s_t given $(b_1,\ldots,b_i)$ is equal to (clearly $\hat{s}_0^{\mathrm{MMSE}}=0$)

$$\hat{s}_i^{\mathrm{MMSE}}=\hat{s}_{i-1}^{\mathrm{MMSE}}+\mathrm{SNR}\sqrt{\frac{P_S}{(1+\mathrm{SNR})^{i+1}}}b_i. \tag{22}$$

2) Transmits the estimation error $\tilde{s}_{i-1}$ after a suitable power adjustment:

$$a_i=\frac{(1+\mathrm{SNR})^{i-1}}{P_S}\tilde{s}_{i-1}^{\mathrm{MMSE}}. \tag{23}$$

Receiver. At channel use $i\in[K_C]$

1) Calculates the MMSE estimate $\hat{s}_i^{\mathrm{MMSE}}$ of s from $(b_1,\ldots,b_i)$ as in (22);
2) Calculates the CUBE estimate $\hat{s}_i$ of s from $(b_1,\ldots,b_i)$ using (19):

$$\hat{s}_i=\frac{(1+\mathrm{SNR})^{i}}{(1+\mathrm{SNR})^{i}-1}\hat{s}_i^{\mathrm{MMSE}}.$$

Theorem 5.1 (see [26]): Scheme V.1 achieves the OPTA SDR (21).

We provide a short proof, for completeness.

Proof: The transmitter calculates the MMSE estimate from the channel outputs, which are available to it via the feedback and transmits the estimation error with a proper power adjustment.

Clearly, $\hat{s}_{t;0}^{\mathrm{MMSE}}=\mathbb{E}[s_t]=0$.

At channel use i, the MMSE estimate is given by

$$\hat{s}_i^{\mathrm{MMSE}}\triangleq\mathbb{E}\left[s\middle|b_1,\ldots,b_i\right] \tag{24a}$$

$$=\mathbb{E}\left[\hat{s}_{i-1}^{\mathrm{MMSE}}+\tilde{s}_{i-1}^{\mathrm{MMSE}}\middle|b_1,\ldots,b_i\right] \tag{24b}$$

$$=\hat{s}_{i-1}^{\mathrm{MMSE}}+\mathbb{E}\left[\tilde{s}_{i-1}^{\mathrm{MMSE}}\middle|b_1,\ldots,b_i\right] \tag{24c}$$

$$=\hat{s}_{i-1}^{\mathrm{MMSE}}+\frac{\mathrm{SNR}}{1+\mathrm{SNR}}\sqrt{\frac{P_S}{(1+\mathrm{SNR})^{i-1}}}b_i \tag{24d}$$

where (24c) holds since $(\tilde{s}_{i-1}^{\mathrm{MMSE}},b_i)$ are independent of $b_1,\ldots,b_{i-1}$ due to the structure of a_i (23), the fact that the MMSE estimation error is orthogonal to all the measurements, and hence also independent by Gaussianity, and (24d) holds since the MMSE estimator is linear in the Gaussian case.

The MMSE is equal to the conditional MMSE in the Gaussian case, and is given by

$$\mathbb{E}\left[\tilde{s}_i^2 \middle| b_1, \ldots, b_i\right] = \mathbb{E}\left[\tilde{s}_i^2\right] = \frac{1}{(1+\mathrm{SNR})^i} P_S \,.$$

This concludes the proof. ■

Remark 5.1 (Non-Gaussian noise): For the case of an additive non-Gaussian noise channel with a given SNR, Scheme V.1 achieves an SDR of $(1+\mathrm{SNR})^{K_C} - 1$. Since linear optimization is generally suboptimal in the non-Gaussian case, better performance can be attained using an appropriate scheme instead of Scheme V.1 (which is the optimal *linear* scheme); a notable attempt in this direction was made by Shayevitz and Feder [53]. In fact, for most noises, OPTA performance can be attained only in the limit of large information and code blocks, even in the presence of feedback [31, Ch. 3.5].

B. Without Feedback

We now turn to the more involved case of low-delay JSCC without feedback. We concentrate on the case of $K_C = 2$. That is, the case in which one source sample is conveyed over two channel uses.

A naïve approach is to send the source as is over both channel uses, up to a power adjustment. The corresponding unbiased SDR in this case is

$$\mathrm{SDR}_{\mathrm{lin}} = 2\mathrm{SNR}$$

i.e., a linear improvement rather than an exponential one as in (21). This scheme approaches (21) for very low SNRs, but suffers great losses at high SNRs. Note that the linear factor 2 comes from the fact that the total power available over both channel uses has doubled, and the same performance can be attained by allocating all of the available power to the first channel use while remaining silent during the second channel use.

This suggests that better mappings that truly exploit the extra channel use can be constructed. The first to propose an improvement for the 1:2 case were Shannon [35] and Kotel'nikov [36], in the late 1940s. In their works, the source sample is viewed as a point on a single-dimensional line, whereas the two channel uses correspond to a two-dimensional (2-D) space. In these terms, the linear scheme corresponds to mapping the one-dimensional (1-D) source line to a straight line in the 2-D channel space (represented by a dashed line in Fig. 4), and hence clearly cannot provide any improvement, as AWGN is invariant to rotations. However, by mapping the 1-D source line into a 2-D curve that fills the space better, a great boost in performance can be attained, as was demonstrated in [35], [36], [11, Ch. 8.2], [54]–[58], and references therein, for different families of mappings.

In this paper, we concentrate on such a family that is based on the Archimedean spiral, which was considered in several works [54], [58]–[61] (represented by the solid line in Fig. 4):

$$\begin{cases} a_1^{\mathrm{reg}}(s) = c^{\mathrm{reg}}\ s\ \cos(\omega s) = c^{\mathrm{reg}}\,|s|\ \cos(\omega|s|)\,\mathrm{sign}(s) \\ a_2^{\mathrm{reg}}(s) = c^{\mathrm{reg}}\,|s|\ \sin(\omega s) = c^{\mathrm{reg}}\,|s|\ \sin(\omega|s|)\,\mathrm{sign}(s) \end{cases} \tag{25}$$

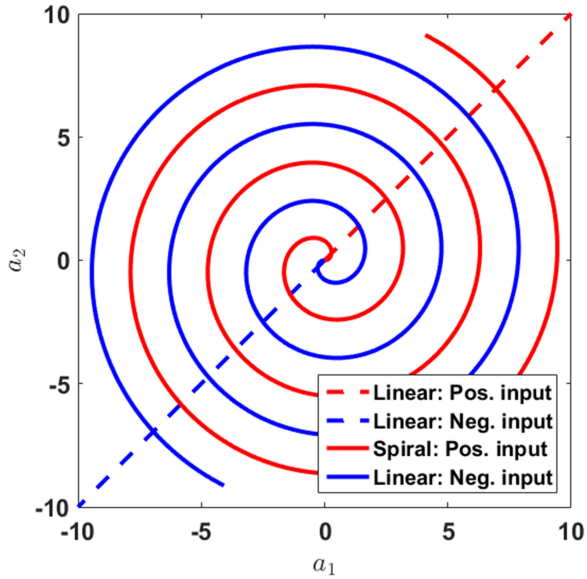

Fig. 4. Linear and Archimedean bispiral curves.

where ω determines the rotation frequency, the factor c^{reg} is chosen to satisfy the power constraint, and the $\mathrm{sign}(s)$ term is needed to avoid overlap of the curve for positive and negative values of s (for each of which now corresponds a distinct spiral, and the two meet only at the origin). This (bi)spiral allows to effectively improve the resolution with respect to small noise values, since the 1-D source line is effectively stretched compared to the noise, and hence the noise magnitude shrinks when the source curve is mapped (contracted) back. However, for large noise values, a jump to a different branch—referred to as a *threshold effect* —may occur, incurring a large distortion. Thus, the value ω needs to be chosen to be as large as possible to allow maximal stretching of the curve for the same given power, while maintaining a low threshold event probability. The SDRs for different values of ω are depicted in Fig. 5(a).

Another ingredient that is used in conjunction with (25) is stretching s prior to mapping it to a bispiral using $\phi_\lambda(s) \triangleq \mathrm{sign}(s)|s|^\lambda$:

$$\begin{cases} a_1^{\mathrm{stretch}}(s) = a_1^{\mathrm{reg}}(\phi_\lambda(s)) = c^{\mathrm{stretch}}|s|^\lambda \cos\left(\omega|s|^\lambda\right)\mathrm{sign}(s) \\ a_2^{\mathrm{stretch}}(s) = a_2^{\mathrm{reg}}(\phi_\lambda(s)) = c^{\mathrm{stretch}}|s|^\lambda \sin\left(\omega|s|^\lambda\right)\mathrm{sign}(s) \end{cases} \tag{26}$$

The choice $\lambda = 0.5$ promises a great boost in performance in the high-SNR regime, as is seen in Fig. 5(b). We further note that although the optimal decoder $\mathbb{E}\left[s|b_1, b_2\right]$ is an MMSE estimator, in this case, the ML decoder $p(b_1, b_2|s)$ achieves similar performance for moderate and high SNRs. A joint optimization of λ and ω for each SNR, for both MMSE and ML decoding, was carried out in [59], with the latter depicted in Fig. 5.

A desired property of the linear JSCC schemes is that their SDR improves with the channel SNR ("SNR universality"). Such an improvement is not allowed by the separation-based technique, as it fails when the actual SNR is lower than the design SNR, and does not promise (almost) any improvement for SNRs above it. This motivated much work in designing JSCC schemes whose performance improves with the SNR, even for the case of large block lengths [62]–[64]. The schemes

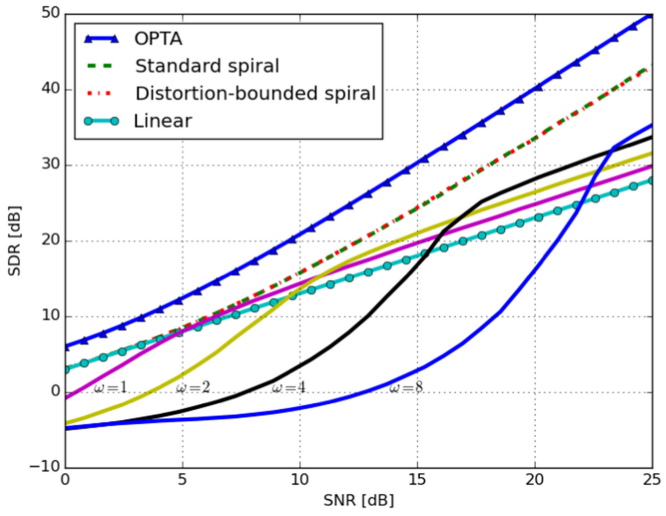

(a) $\lambda = 1$.

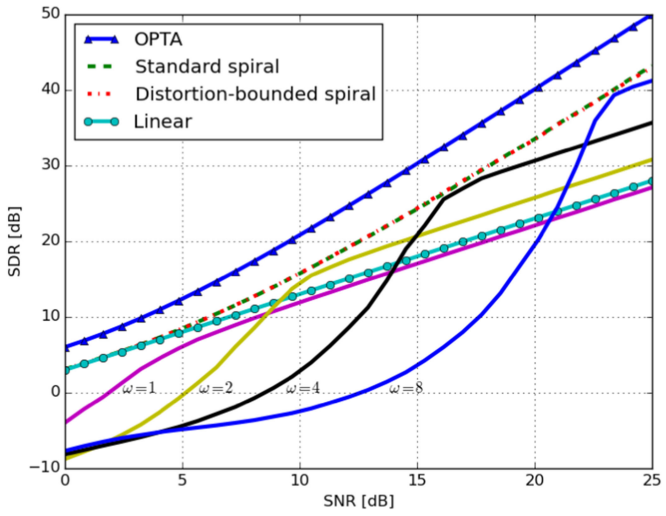

(b) $\lambda = 0.5$.

Fig. 5. Performances of the JSCC linear repetition scheme, OPTA bound, and the JSCC SK spiral scheme for optimized λ and ω, for the standard case ($\beta = 1$) and distortion-bounded case. The solid lines depict the performance of the standard spiral for various values of ω for two stretch parameters $\lambda = 0.5$ and 1, which perform better at high and low SNRs, respectively.

in these works achieve optimal performance (21) for a specific design SNR (21), and improve linearly for higher SNRs. Similar behavior is observed also in Fig. 5 where the optimal ω value varies with the (design) SNR, and mimics closely the quadratic growth in the SDR. Above the design SNR, linear growth is achieved for a particular choice of ω.

We further note that the distortion component due to the threshold event grows with $|s|$. To avoid this behavior, instead of increasing the magnitude $\left\|a^{\text{stretch}}\right\|$ proportionally to the phase $\angle\left(a^{\text{stretch}}\right)$ as in (26), we propose to increase it slightly faster at a pace that guarantees that the incurred distortion does not grow with $|s|$:

$$\begin{cases} a_1^{\text{bounded}}(s) = c^{\text{bounded}}|s|^{\lambda\beta}\cos\left(\omega|s|^{\lambda}\right)\operatorname{sign}(s) \\ a_2^{\text{bounded}}(s) = c^{\text{bounded}}|s|^{\lambda\beta}\sin\left(\omega|s|^{\lambda}\right)\operatorname{sign}(s) \end{cases} \tag{27}$$

for some $\beta > 1$. This has only a slight effect on the resulting SDRs, as is illustrated in Fig. 5.

Finally, note that in no way do we claim that the spiral-based Shannon–Kotel'nikov (SK) scheme is optimal. Various other techniques exist, most using a hybrid of digital and analog components [56], [57], [65], which outperform the spiral-based scheme for various parameters. Nevertheless, this scheme gives good performance boosts, which suffice for our demonstration.

VI. Control via Source–Channel Separation

The separation-based control scheme, outlined next, applies Algorithm 3.2 and sends the resulting quantization indices after encoding with a tree code generated as in Section IV. The observer/transmitter knowingly ignores any decoding errors made by the controller/receiver by internally simulating the system without any decoding errors. On the other hand, the controller/receiver, upon detecting an error in the past, recalculates the steps of Algorithm 3.2 starting from this error and corrects for it in the following steps. This scheme is implemented in [42, `tree/master/code/separate`].

Scheme VI.1 (Separation-based):

Initialization.

1) Selects the number of information bits k that are encoded at every time step t [recall (11), (14)].
2) Sets the size M of the PAM constellation to be 2^k and the number of coded bits, n, to be $K_C k$.
3) Generates $\mathbf{G}_1, \ldots, \mathbf{G}_T$ as in Definition IV.3.
4) Assigns the noiseless-channel capacity C of Algorithm 3.2 to equal k.
5) Initializes Algorithm 3.2.

Observer/Transmitter. At time $t \in [T]$:

1) Runs the observer/transmitter steps of Algorithm 3.2 with the control signal u_t in Step 6 replaced by the signal generated by the controller in (28) below.
2) Maps the resulting quantization index ℓ_t into the k-bit input $\boldsymbol{\imath}_t$ of the tree encoder.
3) Encodes $\boldsymbol{\imath}^t$ into n coded bits $\boldsymbol{c}_t$ according to (14).
4) Maps $\boldsymbol{c}_t$ into K_C constellation points $\boldsymbol{a}_t$ as in (18).
5) Transmits the K_C constellation points $\boldsymbol{a}_t$ over the K_C channel uses.

Controller/Receiver. At time $t \in [T]$:

1) Receives the K_C channel outputs $\boldsymbol{b}_t$.
2) Recovers estimates of all information bits until time t, $\left(\hat{\boldsymbol{\imath}}_{1|t}, \hat{\boldsymbol{\imath}}_{2|t}, \ldots, \hat{\boldsymbol{\imath}}_{t|t}\right)$ as in (12) using Algorithm 1.
3) Maps each $\hat{\boldsymbol{\imath}}_{\tau|t}$ (for $\tau \in [t]$) into a quantization index estimate $\hat{\ell}^{\mathrm{r}}_{\tau|t}$, where the superscript "r" stands for "receiver."
4) Finds the earliest time $\tau \in [t-1]$ for which $\hat{\ell}^{\mathrm{r}}_{\tau|t} \neq \hat{\ell}^{\mathrm{r}}_{\tau|t-1}$. We denote this time instant by t_0. If no such time instant exists, set $t_0 = t$.
5) Runs the controller/receiver steps of Algorithm 3.2 for time instants $\tau = t_0, \ldots, t-1$, with ℓ^t replaced with $(\hat{\ell}^{\mathrm{r}}_{1|t}, \ldots, \hat{\ell}^{\mathrm{r}}_{t|t})$ and the used control signals u^{t-1}.
6) Runs Steps 1 and 3 of the controller/receiver of Algorithm 3.2 for time instant t with ℓ_t replaced with $\hat{\ell}^{\mathrm{r}}_{t|t}$.

7) Applies the control signal

$$u_t = -L_t \hat{x}^{\mathrm{r}}_{t|t} + \sum_{\tau=0}^{t-1} \alpha^{t-1-\tau} L_\tau \left(\hat{x}^{\mathrm{r}}_{\tau|t} - \hat{x}^{\mathrm{r}}_{\tau|t-1} \right) \quad (28a)$$

$$= -L_t \hat{x}^{\mathrm{r}}_{t|t} + \sum_{\tau=t_0}^{t-1} \alpha^{t-1-\tau} L_\tau \left(\hat{x}^{\mathrm{r}}_{\tau|t} - \hat{x}^{\mathrm{r}}_{\tau|t-1} \right) \quad (28b)$$

to the system, where $\hat{x}^{\mathrm{r}}_{\tau|t}$ denotes the estimate of the source x_τ given $\left(\hat{\ell}^{\mathrm{r}}_{1|t}, \hat{\ell}^{\mathrm{r}}_{2|t}, \ldots, \hat{\ell}^{\mathrm{r}}_{\tau|t} \right)$ at the receiver.

VII. Control via Low-Delay JSCC

In this section, we construct a Kalman-filter-like solution [48] by employing JSCC schemes. We note that the additional complication here is due to the communication channel (2) and its inherent input power constraint.

Denote by $\hat{x}^{\mathrm{r}}_{t_1|t_2}$, the estimate of x_{t_1} at the receiver given $\boldsymbol{b}^{t_2}$. Denote further its MSE by

$$P^{\mathrm{r}}_{t_1|t_2} \triangleq \mathbb{E}\left[\left(\tilde{x}^{\mathrm{r}}_{t_1|t_2} \right)^2 \right],$$

where $\tilde{x}^{\mathrm{r}}_{t_1|t_2} \triangleq x_{t_1} - \hat{x}^{\mathrm{r}}_{t_1|t_2}$.

Then, the scheme works as follows. At time instant t, the controller constructs an estimate $\hat{x}^{\mathrm{r}}_{t|t}$ of x_t. It then applies the control signal $u_t = -L_t \hat{x}^{\mathrm{r}}_{t|t}$ to the plant, with L_t given in (10). Note that, since both the controller and the observer know the previously applied control signals u^t, they also know $\hat{x}^{\mathrm{r}}_{t|t}$ and $\hat{x}^{\mathrm{r}}_{t+1|t}$.

Hence, in order to describe x_t, the observer can save transmit power by transmitting the error signal $(x_t - \hat{x}^{\mathrm{r}}_{t|t-1})$, instead of x_t. The controller can then add back $\hat{x}^{\mathrm{r}}_{t|t-1}$ to the received signal to construct $\hat{x}^{\mathrm{r}}_{t|t}$.

Scheme VII.1:

Observer/transmitter: At time t

1) Generates the desired error signal

$$s_t = \tilde{x}^{\mathrm{r}}_{t|t-1} \quad (29a)$$

$$= x_t - \hat{x}^{\mathrm{r}}_{t|t-1} \quad (29b)$$

of average power $P^{\mathrm{r}}_{t|t-1}$ (determined in the sequel).

2) Since the channel input is subject to a unit power constraint (3), s_t is normalized:

$$\bar{s}_t = \frac{1}{\sqrt{P^{\mathrm{r}}_{t|t-1}}} s_t \,. \quad (30)$$

3) Maps $\bar{s}_t$ into K_C channel inputs, constituting the entries of $\boldsymbol{a}_t$, using a bounded-distortion JSCC scheme of choice with (maximum given any input) average distortion $1/\mathrm{SDR}_0$ for the given channel SNR.

4) Sends the K_C channel inputs $\boldsymbol{a}_t$ over the channel (2).

Controller/Receiver: At time $t \in [T]$:

1) Receives the K_C channel outputs $\boldsymbol{b}_t$.

2) Recovers a CUBE of the source signal $\bar{s}_t$: $\hat{\bar{s}}_t = \bar{s}_t + n^{\mathrm{eff}}_t$, where $n^{\mathrm{eff}}_t \perp \bar{s}_t$ is an additive noise of power of (at most) $1/\mathrm{SDR}_0$.

3) Unnormalizes $\hat{\bar{s}}_t$ to construct an estimate of s_t:

$$\hat{s}_t = \sqrt{P^{\mathrm{r}}_{t|t-1}}\, \hat{\bar{s}}_t \quad (31a)$$

$$= \sqrt{P^{\mathrm{r}}_{t|t-1}} \left(\bar{s}_t + n^{\mathrm{eff}}_t \right) \quad (31b)$$

$$= \tilde{x}^{\mathrm{r}}_{t|t-1} + \sqrt{P^{\mathrm{r}}_{t|t-1}}\, n^{\mathrm{eff}}_t . \quad (31c)$$

4) Constructs an estimate $\hat{x}^{\mathrm{r}}_{t|t}$ of x_t given $\boldsymbol{b}^t$. Since $\hat{s}_t \perp \hat{x}^{\mathrm{r}}_{t|t-1}$, the linear MMSE estimate amounts to[14]

$$\hat{x}^{\mathrm{r}}_{t|t} = \hat{x}^{\mathrm{r}}_{t|t-1} + \frac{\mathrm{SDR}_0}{1+\mathrm{SDR}_0} \hat{s}_t \quad (32)$$

with an MSE of

$$P^{\mathrm{r}}_{t|t} = \frac{P^{\mathrm{r}}_{t|t-1}}{1+\mathrm{SDR}_0}. \quad (33)$$

5) Generates the control signal $u_t = -L_t \hat{x}^{\mathrm{r}}_{t|t}$, and the receiver prediction of the next system state

$$\hat{x}^{\mathrm{r}}_{t|t-1} = \alpha \hat{x}^{\mathrm{r}}_{t-1|t-1} + u_{t-1},$$

where L_t is given as in Lemma 3.1.

Using (33) and (1), the prediction error at the receiver is given by the following recursion:

$$P^{\mathrm{r}}_{t+1|t} = \frac{\alpha^2 P^{\mathrm{r}}_{t|t-1}}{1+\mathrm{SDR}_0} + W. \quad (34)$$

The recursive relation (34) leads to the following condition for the stabilizability of the control system and bound on the achievable cost.

Theorem 7.1 (Achievable): The scalar control system of Section II is stabilizable using Scheme VII.1 if $\alpha^2 < 1 + \mathrm{SDR}_0$, and its infinite-horizon average-stage LQG cost $\bar{J}_\infty$ (6) is bounded from above by

$$\bar{J} \le \frac{Q + \left(\alpha^2 - 1\right) S}{1 + \mathrm{SDR}_0 - \alpha^2} W. \quad (35)$$

The following theorem is an adaptation of the lower bound in [66] to our setting of interest.

Theorem 7.2 (Lower bound): The scalar control system of Section II is stabilizable only if $\alpha^2 < 1 + \mathrm{SDR}_{\mathrm{OPTA}}$, and the optimal achievable infinite-horizon average-stage LQG cost is bounded from below by

$$\bar{J} \ge \frac{Q + \left(\alpha^2 - 1\right) S}{1 + \mathrm{SDR}_{\mathrm{OPTA}} - \alpha^2} W. \quad (36)$$

By comparing (35) and (36), we see that the potential gap between the two bounds stems only from the gap between the bounds on the achievable SDR over the AWGN channel (2).

Remark 7.1 (Stabilizability): The stabilizability condition of Theorem. 7.1 is distant from that of Theorem 7.2 in this case since $\mathrm{SDR}_0 < \mathrm{SDR}_{\mathrm{OPTA}}$. By improving SDR_0, one can improve the achievable stabilizability of the system.

[14] If the resulting effective noise n^{eff}_t is not an AWGN with power that does not depend on the channel input, then a better estimator than that in (32) may be constructed.

It is interesting to note that in this case, in stark contrast to the classical LQG setting in which the system is stabilizable for *any values* of α and W, low values of the SDR render the system unstable. Hence, it provides, among others, the minimal required transmit power for the system to remain stable. The difference from the classical LQG case stems from the additional input power constraint, which effectively couples the power of the effective observation noise with that of the estimation error, and was previously observed in, e.g., [28], [29],[66], and [67]. The existence of a threshold SDR below which the system cannot be stabilized parallels the result of Sinopoli *et al.* [33] for control over packet drop channels, which shows that the system cannot be stabilized if the packet drop probability exceeds a certain threshold.

We next discuss the special cases of $K_C = 1$ in Section VII-A, $K_C \in \mathbb{N}$ and instantaneous perfect output feedback—in Section VII-B, and $K_C = 2$ in Section VII-C.

A. Source–Channel Rate Match

In this section, we treat the case of $K_C = 1$, namely, the case where the sample rate of the control system and the signaling rate of the communication channel match.

As we have seen in Section V, analog linear transmission of a Gaussian source over an AWGN channel achieves optimal performance (even when infinite delay is allowed), namely, the OPTA SDR (21), for any given input value. Thus, the JSCC scheme that we use in this case is linear transmission—the source is transmitted as is, up to a power adjustment [recall (29) and (30)]

$$a_t = \bar{s}_t = \frac{1}{\sqrt{P^{\mathrm{r}}_{t|t-1}}} s_t \, .$$

Since in this case $\mathrm{SDR}_0 = \mathrm{SDR}_{\mathrm{OPTA}}$, the upper and lower bounds of Theorems 7.1 and 7.2 coincide, establishing the optimum performance in this case.

Corollary 7.1: The scalar control system of Section II with $K_C = 1$ is stabilizable if only if $\alpha^2 < 1 + \mathrm{SNR}$, and the optimal achievable infinite-horizon average stage LQG cost satisfies (35) with equality with $\mathrm{SDR}_0 = \mathrm{SNR}$.

Remark 7.2: The stabilizability condition and optimum MMSE performance were previously established in [28] and [67] and extend also to the noisy-observation case [23].

B. Source–Channel Rate Mismatch With Feedback

When the AWGN channel outputs b^t (2) are available to the transmitter via an instantaneous feedback, we can incorporate Scheme V.1 in Scheme VII.1 to attain the OPTA SDR and again establish the optimal LQG cost of this setting.

Corollary 7.2: The scalar control system of Section II $K_C \in \mathbb{N}$ is stabilizable if only if $\alpha^2 < (1 + \mathrm{SNR})^{K_C}$, and the optimal achievable infinite-horizon average stage LQG cost satisfies (35) with equality with $\mathrm{SDR}_0 = (1 + \mathrm{SNR})^{K_C} - 1$.

Remark 7.3 (Non-Gaussian noise): Following Remark V.1, the achievability of Theorem 7.1 is attainable with $\mathrm{SDR}_0 = (1 + \mathrm{SNR})^{K_c} - 1$ even when the noise is non-Gaussian. In this case, however, the variance W in the lower bound of Theorem 7.2 should be replaced with its entropy power (which is strictly lower than the variance for non-Gaussian processes) [66], and therefore better performance might be achievable. We note that this lower bound is not achievable in general with or without feedback, as is implied by [31, Ch. 3.5].

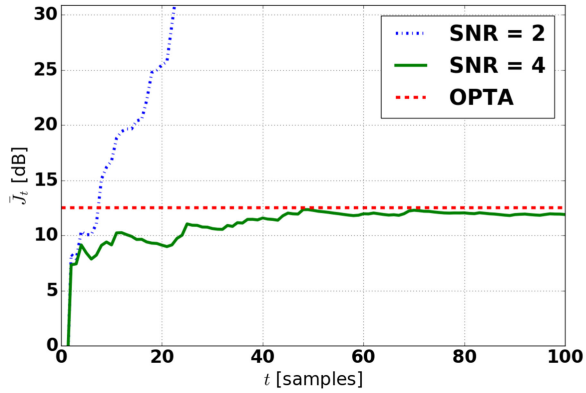

Fig. 6. Optimal average stage LQG cost $\bar{J}$ of a single representative run for $K_C = K_S = 1$, $\alpha = 2$, and SNRs 2 and 4, which correspond to a stabilizable and an unstabilizable systems. The driving noise and observation noise powers and the LQG penalty coefficients are $Q_t \equiv R_t \equiv 1, W = 1$.

C. Source–Channel Rate Mismatch Without Feedback

We now consider the case of $K_C = 2$ channel uses per sample. As we saw in Section V, linear schemes are suboptimal outside the low-SNR regime. Instead, by using nonlinear maps, e.g., the (modified) Archimedean spiral-based SK maps (27), better performance can be achieved. This scheme is implemented in [42, `tree/master/code/joint`].

We note that the improvement in the SDR of the JSCC scheme is substantial when the SDR of the linear scheme is close to $\alpha^2 - 1$, using an improved scheme with better SDR improves substantially the LQG cost. Unfortunately, the spiral-based SK schemes do not promise any improvement for SNRs below 5 dB under ML decoding.

Remark 7.4: By replacing the ML decoder with an MMSE one, strictly better performance can be achieved over the linear scheme for all SNR values.

Remark 7.5: The resulting effective noise at the output of the JSCC receiver is not necessarily Gaussian, and hence the resulting system state, x_t, is not necessarily Gaussian either. Nevertheless, for the bounded-distortion scheme (27), this has no effect on the resulting performance, as is demonstrated next.

VIII. Simulations

A. Rate-Matched Case

The optimal average-stage LQG cost is illustrated in Fig. 6, for a system with $\alpha = 2$ and two SNRs—2 and 4. $\mathrm{SNR} = 4$ satisfies the stabilizability condition $\alpha^2 < 1 + \mathrm{SNR}$, whereas $\mathrm{SNR} = 2$ fails to do so. Unit LQG penalty coefficients $Q_t \equiv R_t \equiv 1$ and unit driving noise power $W = 1$ are used.

B. JSCC for the Rate-Mismatched Case

The effect of the SDR improvement is illustrated in Fig. 7 for a system with $\alpha = 3$ and $W = 1$, for $Q_t \equiv 1$ and $R_t \equiv 0$, by comparing the achievable costs and lower bound of Theorems 7.1 and 7.2. We note that the JSCC scheme with (intermediate) feedback always achieves OPTA.

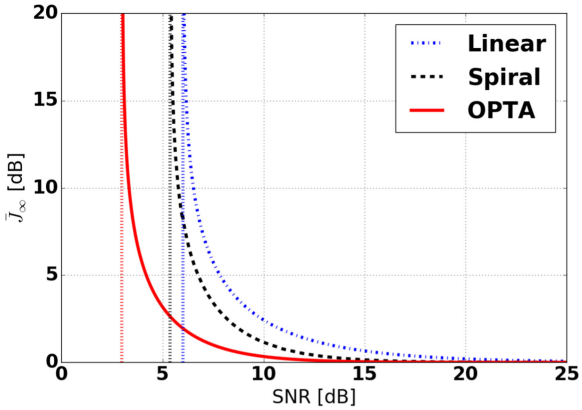

Fig. 7. Average-stage LQG costs when using the (distortion-bounded) SK Archimedean bispiral, repetition, and the lower bound of Theorem 7.2 for $\alpha = 3, W = 1, Q_t \equiv 1, R_t \equiv 0$. The vertical dotted lines represent the minimum SNR below which the cost diverges to infinity.

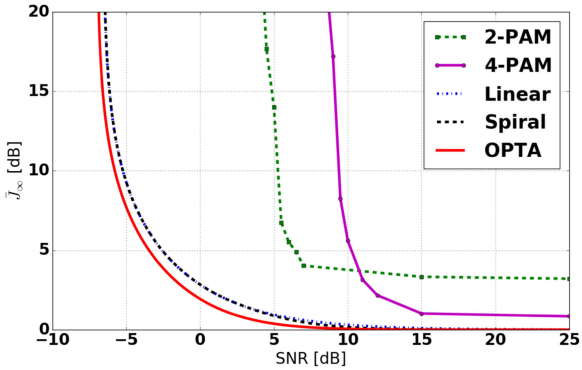

Fig. 8. Average-stage LQG costs averaged over 256 runs when using the (distortion-bounded) JSCC SK Archimedean bispiral scheme, linear (repetition) scheme, separation-based scheme for 2-PAM and 4-PAM, and the (OPTA) lower bound of Theorem 7.2 for $\alpha = 1.2, W = 1, Q_t \equiv 1, R_t = 0$. The SNRs for which the separation based schemes were simulated are marked with squares and circles for 2-PAM and 4-PAM, respectively.

C. Comparison of Separation-Based and JSCC Schemes for the Rate-Mismatched Case

We now compare the performance of the separation-based scheme of Section VI with the JSCC schemes of Section VII. The implementations of these schemes are available in [42].

Fig. 8 shows a comparison of the control costs $\bar{J}_t$ achieved by these schemes for a fully observable scalar plant with $\alpha = 1.2$, $W = 1$, $Q_t \equiv 1$, and $R_t \equiv 0$, over 256 runs.

Clearly, the JSCC-based schemes outperform the separation-based schemes by a large margin.

Note that while larger constellations perform better for high SNRs, the situation is reversed when the SNR is low.

We further include a simulation of a single run of each of the schemes in Fig. 9. For the separation-based scheme, decoding errors (times when $\hat{\ell}^{\mathrm{r}}_{t|t} \neq \ell_t$) are highlighted. Their impact is clear: while the decoder is in error, it applies the wrong control signal, causing the cost function to rapidly deteriorate. In the instance shown, these decoding errors are clearly the major factor degrading performance.

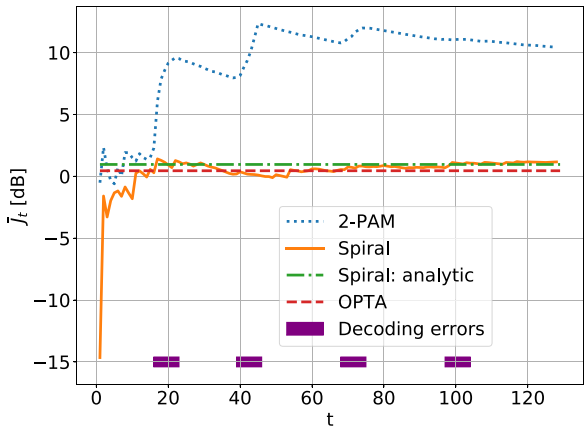

Fig. 9. LQG cost comparison of the separation-based scheme with $k = 2$ (2-PAM constellation), the JSCC SK bispiral scheme, and the OPTA lower bound for $\alpha = 1.2$, $\mathrm{SNR} = 4.5$ dB, $W = 1, Q_t \equiv 1, R_t \equiv 0$. The schemes were simulated for the same disturbance and noise sequences and their results are compared to the analytically derived cost of the JSCC scheme and the OPTA lower bound. Times when decoding errors in the separation-based scheme occur are marked by a thick line.

IX. Discussion and Future Research

A. Separation-Based and JSCC Schemes Comparison

As is evident from the simulation results in Section VIII-C, in addition to demanding far less computation time and memory, and being considerably simpler to implement than separation-based schemes, the JSCC-based schemes also perform much better in terms of control cost.

A key component behind this improvement is the fact that the JSCC schemes of Section V allow the (rare) utilization of large (unbounded) excess transmission power. The separation-based schemes, on the other hand, are limited by the transmission power of the maximum constellation point, which increases as the square root of the average power.[15] Namely, these schemes have a peak power constraint, which is known to have a detrimental effect on the performance [68].

Another unfortunate shortcoming of using separation-based schemes is their incompetence to stabilize higher moments. As was noted already in the seminal work of Sahai and Mitter [7], in order to stabilize higher moments, increased error exponents are required, that need to grow linearly with the moment's order—this behavior is manifested by the abrupt jumps in the cost of this scheme in Fig. 9. In contrast, JSCC schemes (which in our case enjoy an implicit feedback via the control-system loop) can attain a super exponential decay of the error probability, when used to send bits [69] (cf., [68]). Thus, such schemes can stabilize more and even all moments, and guarantee almost-sure stability (by using, e.g., the simple linear/repetition based scheme), as was noted already by Sahai and Mitter [7, Sec. III-C] in the context of anytime reliability.

On the other hand, the JSCC schemes developed in this paper assumed knowledge of the control objectives *at the transmitter* in order to recover the estimates of the controller at the transmitter. The separation-based scheme on the other hand did not

[15] In the limit of infinite-size constellations, the distribution of the constellation tends to a continuous uniform PDF over $[-M/2, M/2]$ of power $P = M^2/12$.

require such knowledge and only assumed knowledge of the control signal u_t.

B. Partially Observable Vector Systems

In this paper, we focused on the simplest case of scalar systems, and $K_C = 2$. As implied by the JSCC theorem (20), an exponentially large (in K_C) gain in the cost can be achieved. Furthermore, the results of Theorems 7.1 and 7.2 readily extend to systems with noisy observations [41] as well as *vector* states $\boldsymbol{x}_t$ and *vector* control signals $\boldsymbol{u}_t$ but *scalar* observed outputs y_t.

Interestingly, for the case of vector (observation, state and control) signals, even if the signaling rate of the channel and the sample rate of the observer are equal (rate-matched case), conveying several analog observations over a single channel input may be of the essence. This is achieved by a compression JSCC scheme, e.g., by reversing the roles of the source and the channel inputs in the SK spiral-based scheme. Similarly to their expansion counterparts, such compression JSCC schemes provide exponentially growing gains with the SNR and dimension [35], [36], [54], [57]–[59], [61], and promise better LQG costs than their linear counterparts [28].

The vector case is much more challenging, in general. Beyond the difficulty of designing good higher-dimensional curves/surfaces (and more general geometric structures) [55], [57] and lower bounds [66], there is a challenge of even determining and prioritizing the subspaces of the state space of $\boldsymbol{x}_t$ that needs to be conveyed. To make this point more clear, consider a linear plant

$$\boldsymbol{x}_{t+1} = A\boldsymbol{x}_t + Bu_t + \boldsymbol{w}_t$$

with *2-D* state ($\boldsymbol{x}_t$), output ($\boldsymbol{y}_t$), and i.i.d. Gaussian noise ($\boldsymbol{w}_t$) signals, a *scalar* control signal (u_t), and known 2×2 and 2×1 matrices A and B, respectively. Since u_t is 1-D, optimizing the cost of the next step mandates transmitting the information that is needed to reconstruct/approximate the optimal u_t (say a 1-D linear combination of $\boldsymbol{x}_t$: $u_t = -\mathbf{K}_t^T \boldsymbol{x}_t$, where $\mathbf{K}_t$ is a predetermined column vector). However, since A mixes both the subspace dictated by $\mathbf{K}^T$ and its perpendicular, better performance might be attained in future steps ("cost-to-go") by transmitting both subspaces.

C. Oblivious Transmitter

In this paper, we assumed that the observer knows all past control signals. We note that such information is not needed for the JSCC schemes (without feedback), for the special case of variance control, i.e., $R_t \equiv 0$.

For the more general LQG-cost setting ($R_t \not\equiv 0$), this assumption can be viewed as a two-sided side-information scenario. Nevertheless, although this is a common situation in practice, there are scenarios in which the observer is oblivious of the control signal applied or has only a noisy measurement of the actuation signal generated by the controller. Such settings can be regarded as a JSCC problem with side information at the receiver (only), and can be treated using JSCC techniques designed for this case, some of which combine naturally with the JSCC schemes for rate mismatch [57], [64], [65]. In fact, this idea was recently applied for the related problem of communication over an AWGN channel with AWGN feedback channel in [70].

We further note that for bounded noise (even worst-case/arbitrary), parallel results can be achieved.

D. Packeted Transmission With Erasures

In the separation-based schemes, following the work of Sahai and Mitter [7], we used a decoder that ought to make a decision on all information bits transmitted until that time, even if its "belief" of a particular bit—quantified by an appropriate metric, say Fano's metric—is low.

An alternative to this approach is to allow declaring an *erasure*, for bits of "low belief." This idea was advocated and explored in the celebrated work of Fano [71] for block codes, where a tradeoff between the achievable error erasure exponents was established.

Furthermore, by substantially increasing the error exponent (lowering the error probability), at the expense of decreasing the erasure exponent (increasing the erasure probability), one can drive the separation-based scheme toward that of a noiseless channel with occasional packet erasures. Interestingly, Algorithm 3.2 and the lower bound of Theorem 7.2 readily extend to this case due to their "greedy nature" [34].

ACKNOWLEDGMENT

The authors thank H. Yıldız for valuable discussions and help with parts of the simulation.

REFERENCES

[1] J. P. Hespanha, P. Naghshtabrizi, and U. Xu, "A survey of recent results in networked control systems," *Proc. IEEE*, vol. 95, no. 1, pp. 138–162, Jan. 2007.
[2] V. Gupta, A. F. Dana, J. P. Hespanha, R. M. Murray, and B. Hassibi, "Data transmission over networks for estimation and control," *IEEE Trans. Autom. Control*, vol. 54, no. 8, pp. 1807–1819, Aug. 2009.
[3] L. Schenato, B. Sinopoli, M. Franceschetti, K. Poola, and S. S. Sastry, "Foundations of control and estimation over lossy networks," *Proc. IEEE*, vol. 95, no. 1, pp. 163–187, Jan. 2007.
[4] S. Yüksel and T. Başar, *Stochastic Networked Control Systems: Stabilization and Optimization Under Information Constraints*. Boston, MA, USA: Birkhäuser, 2013.
[5] C. E. Shannon, "A mathematical theory of communication," *Bell Syst. Tech. J.*, vol. 27, pp. 379–423, Jul. 1948.
[6] A. El Gamal and Y.-H. Kim, *Network Information Theory*. Cambridge, U.K.: Cambridge Univ. Press, 2011.
[7] A. Sahai and S. K. Mitter, "The necessity and sufficiency of anytime capacity for stabilization of a linear system over a noisy communication link—part I: Scalar systems," *IEEE Trans. Inf. Theory*, vol. 52, no. 8, pp. 3369–3395, Aug. 2006.
[8] L. J. Schulman, "Coding for interactive communication," *IEEE Trans. Inf. Theory*, vol. 42, no. 6, pp. 1745–1756, Nov. 1996.
[9] R. T. Sukhavasi and B. Hassibi, "Linear time-invariant anytime codes for control over noisy channels," *IEEE Trans. Autom. Control*, vol. 61, no. 12, pp. 3826–3841, Dec. 2016.
[10] J. M. Wozencraft, "Sequential decoding for reliable communications," Tech. Rep. 325, Res. Lab. Electron., Massachusetts Inst. Technol., Cambridge, MA, USA, Aug. 1957.
[11] J. M. Wozencraft and I. M. Jacobs, *Principles of Communication Engineering*. New York, NY, USA: Wiley, 1965.
[12] F. Jelinek, *Probabilistic Information Theory: Discrete and Memoryless Models*. New York, NY, USA: McGraw-Hill, 1968.
[13] R. G. Gallager, *Information Theory and Reliable Communication*. New York, NY, USA: Wiley, 1968.
[14] A. J. Viterbi and J. K. Omura, *Principles of Digital Communication and Coding*. New York, NY, USA: McGraw-Hill, 1979.
[15] R. Johannesson and K. S. Zigangirov, *Fundamentals of Convolutional Coding*. New York, NY, USA: Wiley-IEEE Press, 1999.

[16] A. Khina, W. Halbawi and B. Hassibi, "(almost) practical tree codes," in *Proc. IEEE Int. Symp. Inf. Theory*, Barcelona, Spain, Jul. 2016, pp. 2404–2408.

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

[18] S. Yüksel, "Stochastic stabilization of noisy linear systems with fixed-rate limited feedback," *IEEE Trans. Autom. Control*, vol. 55, no. 12, pp. 2847–2853, Dec. 2010.

[19] P. Minero, M. Franceschetti, S. Dey, and G. N. Nair, "Data rate theorem for stabilization over time-varying feedback channels," *IEEE Trans. Autom. Control*, vol. 54, no. 2, pp. 243–255, Feb. 2009.

[20] S. Särkkä, *Bayesian Filtering and Smoothing*, vol. 3. Cambridge, U.K.: Cambridge Univ. Press, 2013.

[21] A. Gersho and R. M. Gray, *Vector Quantization and Signal Compression*. Boston, MA, USA: Kluwer, 1992.

[22] L. Bao, M. Skoglund, and K. H. Johansson, "Iterative encoder–controller design for feedback control over noisy channels," *IEEE Trans. Autom. Control*, vol. 56, no. 2, pp. 265–278, Feb. 2011.

[23] A. Khina, Y. Nakahira, Y. Su, and B. Hassibi, "Algorithms for optimal control with fixed-rate feedback," in *Proc. IEEE Conf. Decision Control*, Melbourne, VIC, Australia, Dec. 2017, pp. 6015–6020.

[24] Y. Kochman, A. Khina, U. Erez, and R. Zamir, "Rematch-and-forward: Joint source/channel coding for parallel relaying with spectral mismatch," *IEEE Trans. Inf. Theory*, vol. 60, no. 1, pp. 605–622, Jan. 2014.

[25] T. Goblick, "Theoretical limitations on the transmission of data from analog sources," *IEEE Trans. Inf. Theory*, vol. 11, no. 4, pp. 558–567, Oct. 1965.

[26] P. Elias, "Channel capacity without coding," *Proc. IRE*, vol. 45, no. 3, pp. 381–381, Jan. 1957.

[27] R. Bansal and T. Başar, "Simultaneous design of measurement and control strategies for stochastic systems with feedback," *Automatica*, vol. 25, no. 5, pp. 679–694, Sep. 1989.

[28] J. S. Freudenberg, R. H. Middleton, and V. Solo, "Stabilization and disturbance attenuation over a Gaussian communication channel," *IEEE Trans. Autom. Control*, vol. 55, no. 3, pp. 795–799, Mar. 2010.

[29] S. Tatikonda, A. Sahai, and S. K. Mitter, "Stochastic linear control over a communication channel," *IEEE Trans. Autom. Control*, vol. 49, no. 9, pp. 1549–1561, Sep. 2004.

[30] N. Elia, "When Bode meets Shannon: Control-oriented feedback communication schemes," *IEEE Trans. Autom. Control*, vol. 49, no. 9, pp. 1477–1488, Sep. 2004.

[31] M. Gastpar, "To code or not to code," Ph.D. dissertation, School Comput. Commun. Sci., École Polytechnique Fédérale de Lausanne, Lausanne, Switzerland, Dec. 2002.

[32] S. Tatikonda and S. K. Mitter, "Control over noisy channels," *IEEE Trans. Autom. Control*, vol. 49, no. 7, pp. 1196–1201, Jul. 2004.

[33] B. Sinopoli, L. Schenato, M. Franceschetti, K. Poola, M. I. Jordan, and S. S. Sastry, "Kalman filtering with intermittent observations," *IEEE Trans. Autom. Control*, vol. 49, no. 9, pp. 1453–1464, Sep. 2004.

[34] A. Khina, V. Kostina, A. Khisti, and B. Hassibi, "Tracking and control of Gauss–Markov processes over packet-drop channels with acknowledgments," *IEEE Trans. Control Netw. Syst.*.

[35] C. E. Shannon, "Communication in the presence of noise," *Proc. IRE*, vol. 37, no. 1, pp. 10–21, Jan. 1949.

[36] V. A. Kotel'nikov, *The Theory of Optimum Noise Immunity*. New York, NY, USA: McGraw-Hill, 1959.

[37] S. Yüksel and S. Tatikonda, "A counterexample in distributed optimal sensing and control," *IEEE Trans. Autom. Control*, vol. 54, no. 4, pp. 841–844, Apr. 2009.

[38] U. Kumar, V. Gupta, and J. N. Laneman, "On stability across a Gaussian product channel," in *Proc. IEEE Conf. Decision Control Eur. Control Conf.*, Orlando, FL, USA, Dec. 2011, pp. 3142–3147.

[39] A. A. Zaidi, T. J. Oechtering, S. Yüksel, and M. Skoglund, "Stabilization and control over Gaussian networks," in *Information and Control in Networks*, B. Bernhardsson, G. Como, and A. Rantzer, Eds. New York, NY, USA: Springer, 2014, pp. 39–85.

[40] H. S. Witsenhausen, "A counterexample in stochastic optimum control," *SIAM J. Control*, vol. 6, no. 1, pp. 131–147, Feb. 1968.

[41] A. Khina, G. M. Pettersson, V. Kostina, and B. Hassibi, "Multi-rate control over AWGN channels via analog joint source–channel coding," in *Proc. IEEE Conf. Decision Control*, Las Vegas, NV, USA, Dec. 2016, pp. 5968–5973.

[42] E. Riedel Gårding, "Control over AWGN channels with and without source–channel separation: Python implementation," Aug. 2017. [Online]. Available: https://github.com/eliasrg/SURF2017/

[43] S. Dharmadhikari and K. Joag-Dev, *Unimodality, Convexity, and Applications* Probability and Mathematical Statistics. Boston, MA, USA: Academic, 1988.

[44] P. E. Fleischer, "Sufficient conditions for achieving minimum distortion in a quantizer," in *Proc. IEEE Int. Convention Rec., Part 1*, Jan. 1964, pp. 104–111.

[45] J. C. Kieffer, "Uniqueness of locally optimal quantizer for log-concave density and convex error weighting function," *IEEE Trans. Inf. Theory*, vol. 29, no. 1, pp. 42–47, Jan. 1983.

[46] A. V. Trushkin, "Sufficient conditions for uniqueness of a locally optimal quantizer for a class of convex error weighting functions," *IEEE Trans. Inf. Theory*, vol. 28, no. 2, pp. 187–198, Mar. 1982.

[47] T. Fischer, "Optimal quantized control," *IEEE Trans. Autom. Control*, vol. 27, no. 4, pp. 996–998, Aug. 1982.

[48] D. P. Bertsekas, *Dynamic Programming and Optimal Control*, vol. I, 2nd ed. Belmont, MA, USA: Athena Sci., 2000.

[49] M. Fu, "Lack of separation principle for quantized linear quadratic Gaussian control," *IEEE Trans. Autom. Control*, vol. 57, no. 9, pp. 2385–2390, Sep. 2012.

[50] R. M. Fano, *Transmission of Information*. Cambridge, MA, USA: MIT Press, 1961.

[51] A. Khina, W. Halbawi, and B. Hassibi, "Sequential stack decoder: Python implementation," Jan. 2016. [Online]. Available: http://www.eng.tau.ac.il/anatolyk/code/tree_codes.py

[52] E. Arıkan, "An upper bound on the cutoff rate of sequential decoding," *IEEE Trans. Inf. Theory*, vol. 34, no. 1, pp. 55–63, Jan. 1988.

[53] O. Shayevitz and M. Feder, "The posterior matching feedback scheme for joint source–channel coding with bandwidth expansion," in *Proc. Data Comput. Conf.*, Snowbird, UT, USA, 2009, pp. 83–92.

[54] S. Y. Chung, "On the construction of some capacity-approaching coding schemes," Ph.D. dissertation, , Dept. EECS, Massachusetts Inst. Technol., Cambridge, MA, USA, 2000.

[55] P. A. Floor, "On the theory of Shannon–Kotelnikov mappings in joint source–channel coding," Ph.D. dissertation, Dept. Electron. Telecommun., Norwegian Univ. Sci. Technol., Trondheim, Norway, May 2008.

[56] M. Kleiner and B. Rimoldi, "Asymptotically optimal joint source–channel with minimal delay," in *Proc. IEEE Global Telecommun. Conf.*, Honolulu, HI, USA, Nov./Dec. 2009, pp. 1–5.

[57] E. Akyol, K. B. Viswanatha, K. Rose, and T. A. Ramstad, "On zero-delay source–channel coding," *IEEE Trans. Inf. Theory*, vol. 60, no. 12, pp. 7473–7489, Dec. 2014.

[58] F. Hekland, P. A. Floor, and T. A. Ramstad, "Shannon–Kotel'nikov mappings in joint source–channel coding," *IEEE Trans. Commun.*, vol. 57, no. 1, pp. 94–105, Jan. 2009.

[59] Y. Hu, J. Garcia-Frias, and M. Lamarca, "Analog joint source–channel coding using non-linear curves and MMSE decoding," *IEEE Trans. Commun.*, vol. 59, no. 11, pp. 3016–3026, Nov. 2011.

[60] I. Kvecher and D. Rephaeli, "An analog modulation using a spiral mapping," in *Proc. IEEE Convection Electr. Electron. Eng. Israel*, Eilat, Israel, Nov. 2006, pp. 194–198.

[61] A. Ingber and M. Feder, "Power preserving 2:1 bandwidth reduction mappings," in *Proc. Data Compression Conf.*, Snowbird, UT, USA, Mar. 1997, p. 385.

[62] U. Mittal and N. Phamdo, "Hybrid digital-analog (HDA) joint source–channel codes for broadcasting and robust communications," *IEEE Trans. Inf. Theory*, vol. 48, no. 5, pp. 1082–1102, May 2002.

[63] Z. Reznic, M. Feder, and R. Zamir, "Distortion bounds for broadcasting with bandwidth expansion," *IEEE Trans. Inf. Theory*, vol. 52, no. 8, pp. 3778–3788, Aug. 2006.

[64] Y. Kochman and R. Zamir, "Analog matching of colored sources to colored channels," *IEEE Trans. Inf. Theory*, vol. 57, no. 6, pp. 3180–3195, Jun. 2011.

[65] X. Chen and E. Tuncel, "Zero-delay joint source–channel coding using hybrid digital–analog schemes in the Wyner–Ziv setting," *IEEE Trans. Commun.*, vol. 62, no. 2, pp. 726–735, Feb. 2014.

[66] V. Kostina and B. Hassibi, "Rate–cost tradeoffs in control," in *Proc. Allerton Conf. Commun., Control, Comput.*, Monticello, IL, USA, Sep. 2016, pp. 1157–1164.

[67] J. S. Freudenberg, R. H. Middleton, and J. H. Braslavsky, "Stabilization and disturbance attenuation over a Gaussian communication channel," in *Proc. IEEE Conf. Decision Control*, New Orleans, LA, USA, Dec. 2007, pp. 3958–3963.

[68] A. D. Wyner, "On the Schalkwijk–Kailath coding scheme with a peak energy constraint," *IEEE Trans. Inf. Theory*, vol. 14, no. 1, pp. 129–134, Jan. 1968.

[69] J. P. M. Schalkwijk and T. Kailath, "A coding scheme for additive noise channels with feedback–I: No bandwidth constraint," *IEEE Trans. Inf. Theory*, vol. IT-12, no. 2, pp. 172–182, Apr. 1966.
[70] A. Ben-Yishai and O. Shayevitz, "Interactive schemes for the AWGN channel with noisy feedback," *IEEE Trans. Inf. Theory*, vol. 63, no. 4, pp. 2409–2427, Apr. 2017.
[71] G. D. Forney Jr., "Exponential error bounds for erasure, list, and decision feedback schemes," *IEEE Trans. Inf. Theory*, vol. IT-14, no. 2, pp. 206–220, Mar. 1968.

Anatoly Khina (S'08–M'17) was born in Moscow, USSR, in 1984, and moved to Israel in 1991. He received the B.Sc. (*summa cum laude*), the M.Sc. (*summa cum laude*), and the Ph.D. degrees in 2006, 2010, and 2016, respectively, all in electrical engineering from Tel Aviv University, Tel Aviv, Israel.

He is currently a Senior Lecturer with the Department of Electrical Engineering-Systems, Tel Aviv University. He was a Postdoctoral Scholar with the Department of Electrical Engineering, California Institute of Technology, Pasadena, CA, USA, from 2015 to 2018, and a Research Fellow with the Simons Institute for the Theory of Computing, University of California, Berkeley, Berkeley, CA, USA, during the Spring of 2018. His research interests include information theory, control theory, signal processing, and matrix analysis. In parallel to his studies, he has worked as an Engineer in various roles focusing on algorithms, software, and hardware R&D.

Dr. Khina is a recipient of the Simons-Berkeley and Qualcomm Research Fellowships; Fulbright, Rothschild, and Marie Skłodowska-Curie Postdoctoral Fellowships; Clore Scholarship; Trotsky Award; Weinstein Prize in signal processing; Intel award for Ph.D. research; and the first prize for outstanding research work in the field of communication technologies from the Advanced Communication Center's Feder Family Award program.

Elias Riedel Gårding received the B.Sc. degree in engineering physics from the Royal Institute of Technology (KTH), Stockholm, Sweden, in 2017 and the Master's degree in applied mathematics, focusing on theoretical physics, from the University of Cambridge, Cambridge, U.K., in 2018. He is currently working toward the M.Sc. degree in subatomic and astrophysics with KTH.

He contributed to this work as a Summer Undergraduate Research Fellow with the California Institute of Technology (2017). He was a summer student with the Joint Research Centre of Photonics, Zhejiang University, Hangzhou, China, in 2016 and contributed to NMR spectroscopy research with the Department of Biochemistry, University of Cambridge in the summer of 2018.

Gustav M. Pettersson received the B.Sc. degree in engineering physics in 2017 and the M.Sc. degree in aerospace engineering in 2018, from the KTH Royal Institute of Technology, Stockholm, Sweden.

He is currently with the European Space Agency, Noordwijk, the Netherlands. He has previously been with the California Institute of Technology, Pasadena, CA, USA, in 2016, the Science for Life Laboratory, Solna, Sweden, in 2017, and at the NASA Ames Research Center, Moffett Field, CA, USA, in 2018.

Victoria Kostina (S'12–M'14) received the Bachelor's degree in applied mathematics and physics from Moscow institute of Physics and Technology, Dolgoprudny, Russia, in 2004, where she was affiliated with the Institute for Information Transmission Problems, Russian Academy of Sciences, the Master's degree in electrical engineering from the University of Ottawa, Ottawa, ON, Canada, in 2006, and the Ph.D. degree in electrical engineering from Princeton University, Princeton, NJ, USA, 2013.

She joined Caltech as an Assistant Professor of electrical engineering in the fall of 2014. Her research interests include information theory, coding, control, and communications.

Dr. Kostina received the Natural Sciences and Engineering Research Council of Canada postgraduate scholarship (2009–2012), the Princeton Electrical Engineering Best Dissertation Award (2013), the Simons-Berkeley Research Fellowship (2015), and the NSF CAREER award (2017).

Babak Hassibi (M'08) was born in Tehran, Iran, in 1967. He received the B.S. degree from the University of Tehran, Tehran, Iran, in 1989, and the M.S. and Ph.D. degrees from Stanford University, Stanford, CA, USA, in 1993 and 1996, respectively, all in electrical engineering.

He has been with the California Institute of Technology, Pasadena, CA, USA, since January 2001, where he is currently the Mose and Lilian S. Bohn Professor of electrical engineering. From 20013 to 2016, he was the Gordon M. Binder/Amgen Professor of electrical engineering and from 2008 to 2015, he was an Executive Officer of electrical engineering, as well as an Associate Director of information science and technology. From October 1996 to October 1998, he was a Research Associate with the Information Systems Laboratory, Stanford University, and from November 1998 to December 2000, he was a Member of the Technical Staff with the Mathematical Sciences Research Center, Bell Laboratories, Murray Hill, NJ, USA. He has also held short-term appointments at Ricoh California Research Center, the Indian Institute of Science, and Linköping University, Sweden. He is the coauthor of the books (both with A. H. Sayed and T. Kailath) *Indefinite Quadratic Estimation and Control: A Unified Approach to* H^2 *and* H^∞ *Theories* (SIAM, 1999) and *Linear Estimation* (Prentice Hall, 2000). His research interests include communications and information theory, control and network science, and signal processing and machine learning.

Dr. Hassibi is a recipient of an Alborz Foundation Fellowship, the 1999 O. Hugo Schuck best paper award of the American Automatic Control Council (with H. Hindi and S. P. Boyd), the 2002 National Science Foundation Career Award, the 2002 Okawa Foundation Research Grant for Information and Telecommunications, the 2003 David and Lucille Packard Fellowship for Science and Engineering, the 2003 Presidential Early Career Award for Scientists and Engineers, and the 2009 Al-Marai Award for Innovative Research in Communications, and was a participant in the 2004 National Academy of Engineering "Frontiers in Engineering" program. He has been a Guest Editor for the IEEE TRANSACTIONS ON INFORMATION THEORY special issue on "space-time transmission, reception, coding and signal processing," was an Associate Editor for Communications of the IEEE TRANSACTIONS ON INFORMATION THEORY during 2004–2006, and is currently an Editor for the Journal "*Foundations and Trends in Information and Communication*" and for the IEEE TRANSACTIONS ON NETWORK SCIENCE AND ENGINEERING. He was an IEEE Information Theory Society Distinguished Lecturer for 2016–2017.