

Joint Optimization of Wireless Communication and Networked Control Systems

Lin Xiao¹, Mikael Johansson², Haitham Hindi³,
Stephen Boyd⁴, and Andrea Goldsmith⁴

¹ Dept. of Aeronautics & Astronautics, Stanford University, Stanford CA 94305, USA

² Department of Signals, Sensors and Systems, KTH, SE 100 44 Stockholm, Sweden

³ Systems and Practices Laboratory, PARC, Palo Alto, CA 94304, USA

⁴ Electrical Engineering Department, Stanford University, Stanford, CA 94305, USA

Abstract. We consider a linear system, such as an estimator or a controller, in which several signals are transmitted over wireless communication channels. With the coding and medium access schemes of the communication system fixed, the achievable bit rates are determined by the allocation of communications resources such as transmit powers and bandwidths, to different channels. Assuming conventional uniform quantization and a standard white-noise model for quantization errors, we consider two specific problems. In the first, we assume that the linear system is fixed and address the problem of allocating communication resources to optimize system performance. We observe that this problem is often convex (at least, when we ignore the constraint that individual quantizers have an integral number of bits), hence readily solved. We describe a dual decomposition method for solving these problems that exploits the problem structure. We briefly describe how the integer bit constraints can be handled, and give a bound on how suboptimal these heuristics can be. The second problem we consider is that of jointly allocating communication resources and designing the linear system in order to optimize system performance. This problem is in general not convex. We present an iterative heuristic method based on alternating convex optimization over subsets of variables, which appears to work well in practice.

1 Introduction

We consider a linear system in which several signals are transmitted over wireless communication links, as shown in figure 1. All signals are vector-valued: w is a vector of exogenous signals (such as disturbances or noises acting on the system); z is a vector of performance signals (including error signals and actuator signals); and y and y_r are the signals transmitted over the communication network, and received, respectively. This general arrangement can represent a variety of systems, for example a controller or estimator in which sensor, actuator, or command signals are sent over wireless links. It can also represent a distributed controller or estimator, in which some signals (*i.e.*, inter-process communication) are communicated across a network. In this paper, we address the

problem of optimizing the stationary performance of the linear system by jointly allocating resources in the communication network and tuning the parameters of the linear system.

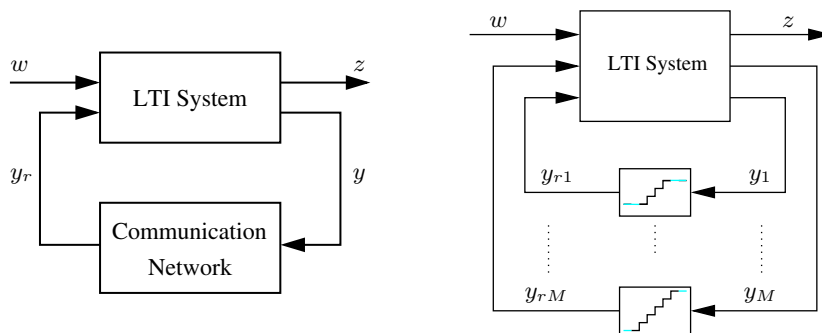


Fig. 1: System set-up (left) and uniform quantization model (right).

Many issues arise in the design of networked controllers and the associated communication systems, including bit rate limitations [WB99, NE00, TSM98], communication delays [NBW98], data packet loss [XHH00], transmission errors [SSK99], and asynchronicity [Özg89]. In this paper we consider only the first issue, *i.e.*, bit rate limitations. In other words, we assume that each communication link has a fixed and known delay (which we model as part of the LTI system), does not drop packets, transfers bits without error, and operates (at least for purposes of analysis) synchronously with the discrete-time linear system.

The problem of control with bit-rate limitations has achieved a lot of attention recently. Much of the research has concentrated on joint design of control and coding to find the minimum bit rate required to stabilize a linear system. For example, [WB99] and [NE98] established various closed-loop stability conditions involving the feedback data rate and eigenvalues of the open-loop system, while [BM97, TSM98] studied control with communication constraints within the classical linear quadratic Gaussian framework. Closely related is also the research on control with quantized feedback information, see [Cur70, Del90, KH94, BL00, EM01].

Our focus is different. We assume that the source coding, channel coding and medium access scheme of the communication system are fixed and concentrate on finding the allocation of communications resources (such as transmit powers and bandwidths) and linear system parameters that yields the optimal closed-loop performance. For a fixed sampling frequency of the linear system, the limit on communication resources translates into a constraint on the number of bits that can be transmitted over each communication channel during one sampling period. We assume that the individual signals y_i are coded using conventional memoryless uniform quantizers, as shown in figure 1. This coding scheme is certainly not optimal (see, *e.g.*, [WB97, NE98]), but it is conventional,

easily implemented, and leads to a simple model of how the system performance depends on the bit-rates. In particular, by imposing lower bounds on the number of quantization bits, we ensure that data rates are high enough for stabilization and that the white-noise model for quantization errors introduced by Widrow (see [WKL96] and the references therein) is valid. This approach has clear links to the research in the signal processing literature on allocation of bits in linear systems with quantizers. The main effort of that research has been to derive analysis and design methods for fixed-point filter and controller implementations, (see [Wil85, WK89, SW90]). However, joint optimization of communications resource allocation and linear system design, interacting through bit rate limitations and quantization, has not been addressed in the literature before. Even in the simplified setting under our assumptions, the joint optimization problem is quite nontrivial and its solution requires concepts and techniques from communication, control, and optimization.

We address to specific problems in this paper. First, we assume the linear system is fixed and consider the problem of allocating communication resources to optimize the overall system performance. We observe that this problem is often convex, provided we ignore the constraint that the number of bits for each quantizer is an integer. This means that these communication resource allocation problems can be solved efficiently, using a variety of convex optimization techniques. We describe a general approach for solving these problems based on dual decomposition. The method results in very efficient procedures for solving for many communication resource allocation problems, and reduces to well known water-filling in simple cases. We also show several methods that can be used to handle the integrality constraint. The simplest is to round down the number of bits for each channel to the nearest integer. We show that this results in an allocation of communication resources that is feasible, and at most a factor of two suboptimal in terms of the RMS (root-mean-square) value of critical variable z . We also describe a simple and effective heuristic that often achieves performance close to the bound obtained by solving the convex problem, ignoring the integrality constraints.

The second problem we consider is the problem of jointly allocating communication resources and designing the linear system in order to optimize performance. Here we have two sets of design variables: the communication variables (which indirectly determine the number of bits assigned to each quantizer), and the controller variables (such as estimator or controller gains in the linear system). Clearly the two are strongly coupled, since the effect of quantization errors depends on the linear system, and similarly, the choice of linear system will affect the choice of communication resource allocation. We show that this joint problem is in general not convex. We propose an alternating optimization method that exploits problem structure and appears to work well in practice.

The paper is organized as follows. In §2, we describe the linear system and our model for the effect of uniform quantization error on overall system performance. In §3, we describe a generic convex model for the bit rate limitations imposed by communication systems, and describe several examples. In §4, we formulate

the communication resource allocation problem for fixed linear systems, describe the dual decomposition method which exploits the separable structure, and give a heuristic rounding method to deal with the integrality of bit allocations. In §5, we demonstrate the nonconvexity of the joint design problem, and give a iterative heuristic to solve such problems. Two examples, a networked linear estimator and a LQG control system over communication networks, are used to illustrate the optimization algorithms in §4 and §5. We conclude the paper in §6.

2 Linear System and Quantizer Model

2.1 Linear System Model

To simplify the presentation we assume a synchronous, single-rate discrete-time system. The linear time-invariant (LTI) system can be described as

$$z = G_{11}(\varphi)w + G_{12}(\varphi)y_r, \quad y = G_{21}(\varphi)w + G_{22}(\varphi)y_r, \quad (1)$$

where G_{ij} are LTI operators (*i.e.*, convolution systems described by transfer or impulse matrices). Here, $\varphi \in \mathbf{R}^q$ is the vector of design parameters in the linear system that can be tuned or changed to optimize performance. To give lighter notation, we suppress the dependence of G_{ij} on φ except when necessary. We assume that $y(t)$, $y_r(t) \in \mathbf{R}^M$, *i.e.*, the M scalar signals y_1, \dots, y_M are transmitted over the network during each sampling period.

We assume that the signals sent (*i.e.*, y) and received (*i.e.*, y_r) over the communication links are related by memoryless scalar quantization, which we describe in detail in the next subsections. This means that all communication delays are assumed constant and known, and included in the LTI system model.

2.2 Quantization Model

Unit Uniform Quantizer A unit range uniform b_i -bit quantizer partitions the range $[-1, 1]$ into 2^{b_i} intervals of uniform width 2^{1-b_i} . To each quantization interval a codeword of b_i bits is assigned. Given a received codeword, the input signal y_i is approximated by (or reconstructed as) y_r , the midpoint of the interval. As long as the quantizer does not overflow (*i.e.*, as long as $|y_i| \leq 1$), the relationship between original and reconstructed values can be expressed as

$$Q_{b_i}(y_i) = \frac{\text{round}(2^{b_i-1}y_i)}{2^{b_i-1}}$$

and the quantization error $y_{ri} - y_i$ lies in the interval $\pm 2^{-b_i}$.

The behavior of the quantizer when y_i overflows (*i.e.*, $|y_i| > 1$) is not specified. One approach is to introduce two more codewords, corresponding to negative and positive overflow, respectively, and to extend Q_{b_i} to saturate for $|y_i| \geq 1$. The details of the overflow behavior will not affect our analysis or design, since we assume by appropriate scaling (described below) that overflow does not occur, or occurs rarely enough to not affect overall system performance.

Scaling To avoid overflow, each signal $y_i(t)$ is scaled by the factor $s_i^{-1} > 0$ prior to encoding with a unit uniform b_i -bit quantizer, and re-scaled by the factor s_i after decoding (figure 2), so that

$$y_{ri}(t) = s_i Q_{b_i}(y_i(t)/s_i).$$

The associated quantization error is given by

$$q_i(t) = y_{ri}(t) - y_i(t) = s_i E_{b_i}(y_i(t)/s_i),$$

which lies in the interval $\pm s_i 2^{-b_i}$, provided $|y_i(t)| < s_i$.




Fig. 2: Scaling before and after the quantizer.

To minimize quantization error while ensuring no overflow (or ensuring that overflow is rare) the scale factors s_i should be chosen as the maximum possible value of $|y_i(t)|$, or as a value that with very high probability is larger than $|y_i(t)|$. For example, we can use the so-called *3 σ -rule*,

$$s_i = 3 \mathbf{rms}(y_i),$$

where $\mathbf{rms}(y_i)$ denotes the root-mean-square value of y_i ,

$$\mathbf{rms}(y_i) = \left(\lim_{t \rightarrow \infty} \mathbf{E} y_i(t)^2 \right)^{1/2}.$$

If y_i has a Gaussian amplitude distribution, this choice of scaling ensures that overflow occurs only about 0.3% of the time.

White-Noise Quantization Error Model We adopt the standard stochastic quantization noise model introduced by Widrow (see, *e.g.*, [FPW90, Chapter 10]). Assuming that overflow is rare, we model the quantization errors $q_i(t)$ as independent random variables, uniformly distributed on the interval

$$s_i[-2^{-b_i}, 2^{-b_i}].$$

In other words, we model the effect of quantizing $y_i(t)$ as an additive white noise source $q_i(t)$ with zero mean and variance $\mathbf{E} q_i(t)^2 = (1/3)s_i^2 2^{-2b_i}$, see figure 3. When allocating bits to quantizers, we will impose a lower bound on each b_i . This value should be high enough for stabilizing the closed-loop system (*cf.* [WB99, NE00]) and make the white noise model a reasonable assumption in a feedback control context (*cf.* [WKL96, FPW90]).




Fig. 3: LTI system with white noise quantization noise model.

2.3 Performance of the Closed-Loop System

We can express z and y in terms of the inputs w and q as

$$z = G_{zw}w + G_{zq}q, \quad y = G_{yw}w + G_{yq}q,$$

where G_{zw} , G_{zq} , G_{yw} and G_{yq} are the closed-loop transfer matrices from w and q to z and y , respectively. From the expression for z , we see that it consists of two terms: $G_{zw}w$, which is what z would be if the quantization were absent, and $G_{zq}q$, which is the component of z due to the quantization. The variance of z induced by the quantization is given by

$$V_q = \mathbf{E} \|G_{zq}q\|^2 = \sum_{i=1}^M \|G_{zqi}\|^2 \left(\frac{1}{3} s_i^2 2^{-2b_i} \right) \quad (2)$$

where G_{zqi} is the i th column of the transfer matrix G_{zq} , and $\|\cdot\|$ denotes the \mathbf{L}^2 norm (see [BB91, §5.2.3]). We can use V_q as a measure of the effect of quantization on the overall system performance. If w is also modeled as a stationary stochastic process, the overall variance of z is given by

$$V = \mathbf{E} \|z\|^2 = V_q + \mathbf{E} \|G_{zw}w\|^2. \quad (3)$$

The above expression shows how V_q depends on the allocation of quantizer bits b_1, \dots, b_M , as well as the scalings s_1, \dots, s_M and LTI system (which affect the a_i 's). Note that while the formula (2) was derived assuming that b_i are integers, it makes sense for $b_i \in \mathbf{R}$.

3 Communications Model and Assumptions

3.1 A Generic Model for Bit Rate Constraints

The capacity of communication channels depend on the media access scheme and the selection of certain critical parameters, such as transmission powers and bandwidths or time-slot fractions allocated to individual channels (or groups of channels). We refer to these critical communication parameters collectively as *communication variables*, and denote the vector of communication variables by θ .

The communication variables are themselves limited by various resource constraints, such as limits on the total power or total bandwidth available. We will assume that the medium access methods and coding and modulation schemes are fixed, but that we can optimize over the underlying communication variables θ .

We let $b \in \mathbf{R}^M$ denote the vector of bits allocated to each quantized signal. The associated communication rate r_i (in bits per second) can be expressed as $b_i = \alpha r_i$, where the constant α has the form $\alpha = c_s/f_s$. Here f_s is the sample frequency, and c_s is the channel coding efficiency in source bits per transmission bit. This relationship will allow us to express capacity constraints in terms of bit allocations rather than communication rates.

We will use the following general model to relate the vector of bit allocations b , and the vector of communication variables θ :

$$\begin{aligned} f_i(b, \theta) &\leq 0, \quad i = 1, \dots, m_f \\ h_i^T \theta &\leq d_i, \quad i = 1, \dots, m_h \\ \theta_i &\geq 0, \quad i = 1, \dots, m_\theta \\ \underline{b}_i &\leq b_i \leq \bar{b}_i, \quad i = 1, \dots, M \end{aligned} \tag{4}$$

We make the following assumptions about this generic model.

- The functions f_i are convex functions of (b, θ) , monotone increasing in b and monotone decreasing in θ . These inequalities describe capacity constraints on the communication channels. We will show below that many classical capacity formula satisfy these assumptions.
- The second set of constraints describes resource limitations, such as a total available power or bandwidth for a group of channels. We assume the vectors h_i have nonnegative entries. We assume that d_i , which represent resource limits, are positive.
- The third constraint specifies that the communication resource variables (which represent powers, bandwidths, time-slot fractions) are nonnegative.
- The last group of inequalities specify lower and upper bounds for each bit allocation. We assume that \underline{b}_i and \bar{b}_i are (nonnegative) integers. The lower bounds are imposed to ensure that the white noise model for quantization errors is reasonable. The upper bounds can arise from hardware limitations.

This generic model will allow us to formulate the communication resource allocation problem, *i.e.*, the problem of choosing θ to optimize overall system performance, as a convex optimization problem.

There is also one more important constraint on b not included above:

$$b_i \text{ is an integer, } \quad i = 1 \dots, M. \tag{5}$$

For the moment, we ignore this constraint. We will return to it in §4.2.

3.2 Capacity Constraints

In this section, we describe some simple channel models and show how they fit the generic model (4) given above. More detailed descriptions of these channel models, as well as derivations, can be found in, *e.g.*, [CT91, Gol99].

Gaussian Channel We start by considering a single Gaussian channel. The communication variables are the bandwidth $W > 0$ and transmission power $P > 0$. Let N be the power spectral density of the additive white Gaussian noise at the front-end of the receiver. The channel capacity is given by ([CT91])

$$R = W \log_2 \left(1 + \frac{P}{NW} \right)$$

(in bits per second). The achievable communication rate r is bounded by this channel capacity, *i.e.*, we must have $r \leq R$. Expressed in terms of b , we have

$$b \leq \alpha W \log_2 \left(1 + \frac{P}{NW} \right). \quad (6)$$

We can express this in the form

$$f(b, W, P) = b - \alpha W \log_2 \left(1 + \frac{P}{NW} \right) \leq 0,$$

which fits the generic form (4). To see that the function f is jointly convex in the variables (b, W, P) , we note that the function $g(P) = -\alpha \log_2(1 + P/N)$ is a convex function of P and, therefore its *perspective function* (see [BV04])

$$Wg(P/W) = -\alpha W \log_2 \left(1 + \frac{P}{NW} \right)$$

is a convex function of (P, W) . Adding the linear (hence convex) function b establishes convexity of f . It is easily verified that f is monotone increasing in b , and monotone decreasing in W and P .

Gaussian Broadcast Channel with FdMA In the Gaussian broadcast channel with frequency-domain multiple access (FDMA), a transmitter sends information to n receivers over disjoint frequency bands with bandwidths $W_i > 0$. The communication parameters are the bandwidths W_i and the transmit powers $P_i > 0$ for each individual channel. The communication variables are constrained by a total power limit

$$P_1 + \cdots + P_n \leq P_{\text{tot}}$$

and a total available bandwidth limit

$$W_1 + \cdots + W_n \leq W_{\text{tot}},$$

which have the generic form for communication resource limits.

The receivers are subject to independent white Gaussian noises with power spectral densities N_i . The transmitter assigns power P_i and bandwidth W_i to the i th receiver. The achievable bit rates b are constrained by

$$b_i \leq \alpha W_i \log_2 \left(1 + \frac{P_i}{N_i W_i} \right), \quad i = 1, \dots, n. \quad (7)$$

Again, the constraints relating b and $\theta = (P, W)$ have the generic form (4).

Gaussian Multiple Access Channel with FDMA In a Gaussian multiple access channel with FDMA, n transmitters send information to a common receiver, each using a transmit power P_i over a bandwidth W_i . It has the same set of constraints as for the broadcast channel, except that $N_i = N$, $i = 1, \dots, n$ (since they have a common receiver).

Variations and Extensions The capacity formulas for many other channel models, including the Parallel Gaussian channel, Gaussian broadcast channel with TDMA and the Gaussian broadcast channel with CDMA, are also concave in communications variables and can be included in our framework. It is also possible to combine the channel models above to model more complex communication systems. Finally, channels with time-varying gain variations (fading) as well as rate constraints based on bit error rates (with or without coding) can be formulated in a similar manner; see, *e.g.*, [LG01, CG01].

4 Optimal Resource Allocation for Fixed Linear System

In this section, we assume that the linear system is fixed and consider the problem of choosing the communication variables to optimize the system performance. We take as the objective (to be minimized) the variance of the performance signal z , given by (3). Since this variance consists of a fixed term (related to w) and the variance induced by the quantization, we can just as well minimize the variance of z induced by the quantization error, *i.e.*, the quantity V_q defined in (2). This leads to the optimization problem

$$\begin{aligned} & \text{minimize} && \sum_{i=1}^M a_i 2^{-2b_i} \\ & \text{subject to} && f_i(b, \theta) \leq 0, \quad i = 1, \dots, m_f \\ & && h_i^T \theta \leq d_i, \quad i = 1, \dots, m_h \\ & && \theta_i \geq 0, \quad i = 1, \dots, m_\theta \\ & && \underline{b}_i \leq b_i \leq \bar{b}_i, \quad i = 1, \dots, M \end{aligned} \quad (8)$$

where $a_i = (1/3)\|G_{zqi}\|^2 s_i^2$, and the optimization variables are θ and b . For the moment we ignore the constraint that b_i must be integers.

Since the objective function, and each constraint function in the problem (8) is a convex function, this is a convex optimization problem. This means that it can be solved globally and efficiently using a variety of methods, *e.g.*, interior-point methods (see, *e.g.*, [BV04]). In many cases, we can solve the problem (8) more efficiently than by applying general convex optimization methods by exploiting its special structure. This is explained in the next subsection.

4.1 The Dual Decomposition Method

The objective function in the communication resource allocation problem (8) is separable, *i.e.*, a sum of functions of each b_i . In addition, the constraint functions $f_k(b, \theta)$ usually involve only one b_i , and a few components of θ , since the

channel capacity is determined by the bandwidth, power, or time-slot fraction, for example, allocated to that channel. In other words, the resource allocation problem (8) is almost separable; the small groups of variables (that relate to a given link or channel) are coupled mostly through the resource limit constraints $h_i^T \theta \leq d_i$. These are the constraints that limit the total power, total bandwidth, or total time-slot fractions.

This almost separable structure can be efficiently exploited using a technique called dual decomposition (see, *e.g.*, [BV04, Ber99]). We will explain the method for a simple FDMA system to keep the notation simple, but the method applies to any communication resource allocation problem with almost separable structure. We consider an FDMA system with M channels, and variables $P \in \mathbf{R}^M$ and $W \in \mathbf{R}^M$, with a total power and a total bandwidth constraint. We will also impose lower and upper bounds on the bits. This leads to

$$\begin{aligned}
 & \text{minimize} && \sum_{i=1}^M a_i 2^{-2b_i} \\
 & \text{subject to} && b_i \leq \alpha W_i \log_2(1 + P_i/N_i W_i), \quad i = 1, \dots, M \\
 & && P_i \geq 0, \quad i = 1, \dots, M \\
 & && \sum_{i=1}^M P_i \leq P_{\text{tot}} \\
 & && W_i \geq 0, \quad i = 1, \dots, M \\
 & && \sum_{i=1}^M W_i \leq W_{\text{tot}} \\
 & && \underline{b}_i \leq b_i \leq \bar{b}_i, \quad i = 1, \dots, M.
 \end{aligned} \tag{9}$$

Here N_i is the receiver noise spectral density of the i th channel, and \underline{b}_i and \bar{b}_i are the lower and upper bounds on the number of bits allocated to each channel. Except for the total power and total bandwidth constraint, the constraints are all local, *i.e.*, involve only b_i , P_i , and W_i .

We first form the Lagrange dual problem, by introducing Lagrange multipliers but *only* for the two coupling constraints. The Lagrangian has the form

$$L(b, P, W, \lambda, \mu) = \sum_{i=1}^M a_i 2^{-2b_i} + \lambda \left(\sum_{i=1}^M P_i - P_{\text{tot}} \right) + \mu \left(\sum_{i=1}^M W_i - W_{\text{tot}} \right).$$

The dual function is defined as

$$\begin{aligned}
 g(\lambda, \mu) &= \inf \left\{ L \mid P_i \geq 0, W_i \geq 0, \underline{b}_i \leq b_i \leq \bar{b}_i, b_i \leq \alpha W_i \log_2(1 + P_i/N_i W_i) \right\} \\
 &= \sum_{i=1}^M g_i(\lambda, \mu) - \lambda P_{\text{tot}} - \mu W_{\text{tot}}
 \end{aligned}$$

where

$$\begin{aligned}
 g_i(\lambda, \mu) &= \inf \left\{ a_i 2^{-2b_i} + \lambda P_i + \mu W_i \mid \right. \\
 & \quad \left. P_i \geq 0, W_i \geq 0, \underline{b}_i \leq b_i \leq \bar{b}_i, b_i \leq \alpha W_i \log_2(1 + P_i/N_i W_i) \right\}.
 \end{aligned}$$

Finally, the Lagrange dual problem associated with the communication resource allocation problem (9) is given by

$$\begin{aligned}
 & \text{maximize} && g(\lambda, \mu) \\
 & \text{subject to} && \lambda \geq 0, \quad \mu \geq 0.
 \end{aligned} \tag{10}$$

This problem has only two variables, namely the variables λ and μ associated with the total power and bandwidth limits, respectively. It is a convex optimization problem, since g is a concave function (see [BV04]). Assuming that Slater's condition holds, the optimal value of the dual problem (10) and the primal problem (9) are equal. Moreover, from the optimal solution of the dual problem, we can recover the optimal solution of the primal. Suppose (λ^*, μ^*) is the solution to the dual problem (10), then the primal optimal solution is the minimizer (b^*, P^*, W^*) when evaluating the dual function $g(\lambda^*, \mu^*)$. In other words, we can solve the original problem (9) by solving the dual problem (10).

The dual problem can be solved using a variety of methods, for example, cutting-plane methods. To use these methods we need to be able to evaluate the dual objective function, and also obtain a subgradient for it (see [BV04]), for any given $\mu \geq 0$ and $\lambda \geq 0$. To evaluate $g(\lambda, \mu)$, we simply solve the M separate problems,

$$\begin{aligned} & \text{minimize} && a_i 2^{-2b_i} + \lambda P_i + \mu W_i \\ & \text{subject to} && P_i \geq 0, W_i \geq 0, \\ & && \underline{b}_i \leq b_i \leq \bar{b}_i, \\ & && b_i \leq \alpha W_i \log_2(1 + P_i/N_i W_i), \end{aligned}$$

each with three variables, which can be carried out separately or in parallel. Many methods can be used to very quickly solve these small problems.

A subgradient of the concave function g at (λ, μ) is a vector $h \in \mathbf{R}^2$ such that

$$g(\tilde{\lambda}, \tilde{\mu}) \leq g(\lambda, \mu) + h^T \begin{bmatrix} \tilde{\lambda} - \lambda \\ \tilde{\mu} - \mu \end{bmatrix}$$

for all $\tilde{\lambda}$ and $\tilde{\mu}$. To find such a vector, let the optimal solution to the subproblems be denoted

$$b_i^*(\lambda, \mu), \quad P_i^*(\lambda, \mu), \quad W_i^*(\lambda, \mu).$$

Then, a subgradient of the dual function g is readily given by

$$\begin{bmatrix} \sum_{i=1}^M P_i^*(\lambda, \mu) - P_{\text{tot}} \\ \sum_{i=1}^M W_i^*(\lambda, \mu) - W_{\text{tot}} \end{bmatrix}.$$

This can be verified from the definition of the dual function.

Putting it all together, we find that we can solve the dual problem in time linear in M , which is far better than the standard convex optimization methods applied to the primal problem, which require time proportional to M^3 .

The same method can be applied whenever there are relatively few coupling constraints, and each link capacity is dependent on only a few communication resource parameters. In fact, when there is only one coupling constraint, the subproblems that we must solve can be solved analytically, and the master problem becomes an explicit convex optimization problem with only one variable. It is easily solved by bisection, or any other one-parameter search method. This is the famous water-filling algorithm (see, *e.g.*, [CT91]).

4.2 Integrality of Bit Allocations

We now come back to the requirement that the bit allocations must be integers. The first thing we observe is that we can always round down the bit allocations found by solving the convex problem to the nearest integers. Let b_i denote the optimal solution of the convex resource allocation problem (8), and define $\tilde{b}_i = \lfloor b_i \rfloor$. Here, $\lfloor b_i \rfloor$ denotes the *floor* of b_i , *i.e.*, the largest integer smaller than or equal to b_i . First we claim that \tilde{b} is feasible. To see this, recall that f_k and h_k are monotone decreasing in b , so since b is feasible and $\tilde{b} \leq b$, we have \tilde{b} feasible.

We can also obtain a crude performance bound for \tilde{b} . Clearly the objective value obtained by ignoring the integer constraint, *i.e.*,

$$J_{\text{cvx}} = \sum_{i=1}^M a_i 2^{-2b_i},$$

is a lower bound on the optimal objective value J_{opt} of the problem with integer constraints. The objective value of the rounded-down feasible bit allocation \tilde{b} is

$$J_{\text{rnd}} = \sum_{i=1}^M a_i 2^{-2\tilde{b}_i} \leq \sum_{i=1}^M a_i 2^{-2(b_i-1)} = 4J_{\text{cvx}} \leq 4J_{\text{opt}},$$

using the fact that $\tilde{b}_i \geq b_i - 1$. Putting this together we have

$$J_{\text{opt}} \leq J_{\text{rnd}} \leq 4J_{\text{opt}},$$

i.e., the performance of the suboptimal integer allocation obtained by rounding down is never more than a factor of four worse than the optimal solution. In terms of RMS, the rounded-down allocation is never more than a factor of two suboptimal.

Variable Threshold Rounding Of course, far better heuristics can be used to obtain better integer solutions. Here we give a simple method based on a variable rounding threshold.

Let $0 < t \leq 1$ be a threshold parameter, and round b_i as follows:

$$\tilde{b}_i = \begin{cases} \lfloor b_i \rfloor, & \text{if } b_i - \lfloor b_i \rfloor \leq t, \\ \lceil b_i \rceil, & \text{otherwise.} \end{cases} \quad (11)$$

Here, $\lceil b_i \rceil$ denotes the *ceiling* of b_i , *i.e.*, the smallest integer larger than or equal to b_i . In other words, we round b_i down if its remainder is smaller than or equal to the threshold t , and round up otherwise. When $t = 1/2$, we have standard rounding, with ties broken down. When $t = 1$, all bits are rounded down, as in the scheme described before. This gives a feasible integer solution, which we showed above has a performance within a factor of four of optimal. For $t < 1$ feasibility of the rounded bits \tilde{b} is not guaranteed, since bits can be rounded up.

For a given fixed threshold t , we can round the b_i 's as in (11), and then solve a convex feasibility problem over the remaining continuous variables θ :

$$\begin{aligned} f_i(\tilde{b}, \theta) &\leq 0 \\ h_i^T \theta &\leq d_i \\ \theta_i &\geq 0 \end{aligned} \quad (12)$$

The upper and lower bound constraints $\underline{b}_i \leq \tilde{b}_i \leq \bar{b}_i$ are automatically satisfied because \underline{b}_i and \bar{b}_i are integers. If this problem is feasible, then the rounded \tilde{b}_i 's and the corresponding θ are suboptimal solutions to the integer constrained bit allocation problem.

Since f_i is monotone increasing in b , hence in t , and monotone decreasing in θ , there exists a t^* such that (12) is feasible if $t \geq t^*$ and infeasible if $t < t^*$. In the variable threshold rounding method, we find t^* , the smallest t which makes (12) feasible. This can be done by bisection over t : first try $t = 1/2$. If the resulting rounded bit allocation is feasible, we try $t = 1/4$; if not, we try $t = 3/4$, etc.

Roughly speaking, the threshold t gives us a way to vary the conservativeness of the rounding procedure. When t is near one, almost all bits are rounded down, and the allocation is likely to be feasible. When t is small, we round many bits up, and the bit allocation is unlikely to be feasible. But if it is, the performance (judged by the objective) will be better than the bit allocation found using more conservative rounding (*i.e.*, with a larger t). A simple bisection procedure can be used to find a rounding threshold close to the aggressive one that yields a feasible allocation.

4.3 Example: Networked Linear Estimator

To illustrate the ideas of this section, we consider the problem of designing a networked linear estimator with the structure shown in figure 4. We want to estimate an unknown point $x \in \mathbf{R}^{20}$ using $M = 200$ linear sensors,

$$y_i = c_i^T x + v_i, \quad i = 1, \dots, M.$$

Each sensor uses b_i bits to code its measurements and transmits the coded signal to a central estimator over a Gaussian multiple access channel with FDMA. The performance of the estimator is judged by the estimation error variance $J_K = \mathbf{E} \|\hat{x} - x\|^2$. We assume that $\|x\| \leq 1$ and that the sensor noises v_i are




Fig. 4: Networked linear estimator over a multiple access channel

IID with $\mathbf{E} v_i = 0$, $\mathbf{E} v_i^2 = 10^{-6}$. In this example, the sensor coefficients c_i are uniformly distributed on $[0, 5]$. Since $\|x\| \leq 1$, we choose scaling factors $s_i = \|c_i\|$.

The noise power density of the Gaussian multiple access channel is $N = 0.1$, the coding constant is $\alpha = 2$, and the upper and lower bounds for bit allocations are $\underline{b} = 5$ and $\bar{b} = 12$. The total available power is $P = 300$ and the total available bandwidth is $W = 200$.

The estimator is a linear unbiased estimator

$$\hat{x} = Ky_r,$$

where $KC = I$, with $C = [c_1, \dots, c_M]^T$. In particular, the minimum variance estimator is given by

$$K = (C^T(R_v + R_q)^{-1}C)^{-1}C^T(R_v + R_q)^{-1} \quad (13)$$

where R_v and R_q are the covariance matrices for the sensor noises and quantization noises, respectively. (Note that the estimator gain depends on the bit allocations.) The associated estimation error variance is

$$J_K(b) = \frac{1}{3} \sum_{i=1}^M s_i^2 \|k_i\|^2 2^{-2b_i} + \text{Tr}(KR_vK^T)$$

where k_i is the i th column of the matrix K . Clearly, $J_K(b)$ is on the form (3) and will serve as the objective function for the resource allocation problem (8).

First we allocate power and bandwidth evenly to all sensors, which results in $b_i = 8$ for each sensor. Based on this allocation, we compute the quantization noise variances $\mathbf{E}q_i^2 = (1/3)s_i^2 2^{-2b_i}$ and design a least-squares estimator as in (13). The resulting RMS estimation error is 3.676×10^{-3} . Then we fix the estimator gain K , and solve the relaxed optimization problem (8) to find the resource allocation that minimizes the estimation error variance. The resulting RMS value is 3.1438×10^{-3} . Finally, we perform a variable threshold rounding with $t^* = 0.4211$. Figure 5 shows the distribution of rounded bit allocation. The resulting RMS estimation error is 3.2916×10^{-3} . Thus, the allocation obtained from optimization and variable threshold rounding gives a 10% improved performance compared to the uniform resource allocation, which is not very far from the performance bound given by the relaxed convex optimization problem.

We can see that the allocation obtained from optimization and variable threshold rounding give a 10% improved performance compared to the uniform resource allocation, and is not very far from the performance bound given by the relaxed convex optimization problem.

Note that with the new bit allocations, the quantization covariance changes — it is not the one that was used to design K . We will address this issue of the coupling between the choice of the communication variables and the estimator.

5 Joint Design of Communication and Linear Systems

We have seen that when the linear system is fixed, the problem of optimally allocating communication resources is convex (when we ignore integrality of bit

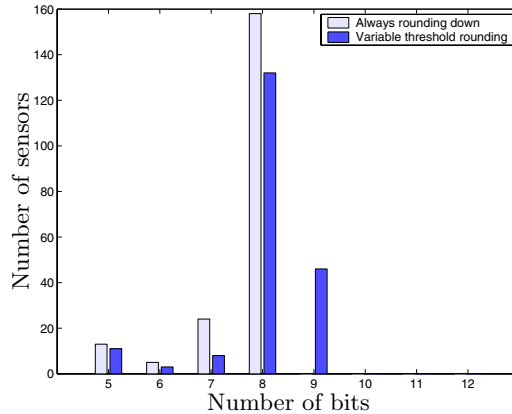


Fig. 5: Bit allocation for networked least-squares estimator.

allocations), and can be efficiently solved. In order to achieve the optimal system performance, however, one should optimize the parameters of the linear system *and* the communication system *jointly*. Unfortunately, this joint design problem is in general not convex. In some cases, however, the joint design problem is bi-convex: for fixed resource allocation the controller design problem is convex, and for fixed controller design and scalings the resource allocation problem is convex. This special structure can be exploited to develop a heuristic method for the joint design problem, that appears to work well in practice.

5.1 Nonconvexity of the Joint Design Problem

To illustrate that the joint design problem is nonconvex, we consider the problem of designing a simple networked least-squares estimator for an example small enough that we can solve the joint problem globally.

An unknown scalar parameter $x \in \mathbf{R}$ is measured using two sensors that are subject to measurement noises:

$$y_1 = x + v_1, \quad y_2 = x + v_2.$$

We assume that v_1 and v_2 are independent zero-mean Gaussian random variables with variances $\mathbf{E}v_1^2 = \mathbf{E}v_2^2 = 0.001$. The sensor measurements are coded and sent over a communication channel with a constraint on the total bit rate. With a total of b_{tot} bits available we allocate b_1 bits to the first sensor and the $b_2 = b_{\text{tot}} - b_1$ remaining bits to the second sensor. For a given bit allocation, the minimum-variance unbiased estimate can be found by solving a weighted least-squares problem. Figure 6 shows the optimal performance as function of b_1 when $b_{\text{tot}} = 8$ and $b_{\text{tot}} = 12$. The relationship is clearly not convex.

These figures, and the optimal solutions, make perfect sense. When $b_{\text{tot}} = 8$, the quantization noise is the dominant noise source, so one should allocate all 8 bits to one sensor and disregard the other. When $b_{\text{tot}} = 12$, the quantization




Fig. 6: Estimator performance for $b_1 + b_2 = 8$ (top) and $b_1 + b_2 = 12$ (bottom).

noises are negligible in comparison with the sensor noise. It is then advantageous to use both sensors (*i.e.*, assign each one 6 bits), since it allows us to average out the effect of the measurement noises.

5.2 Alternating Optimization for Joint Design

The fact that the joint problem is convex in certain subsets of the variables while others are fixed can be exploited. For example (and ignoring the integrality constraints) the globally optimal communication variables can be computed very efficiently, sometimes even semi-analytically, when the linear system is fixed. Similarly, when the communication variables are fixed, we can (sometimes) compute the globally optimal variables for the linear system. Finally, when the linear system variables and the communication variables are fixed, it is straightforward to compute the quantizer scalings using the 3σ -rule. This makes it natural to apply an approach where we sequentially fix one set of variables and optimize over the others:

given initial linear system variables $\phi^{(0)}$, communication variables $\theta^{(0)}$, and scaling factors $s^{(0)}$.

$k := 0$

repeat

1. Fix $\phi^{(k)}$, $s^{(k)}$, and optimize over θ . Let $\theta^{(k+1)}$ be the optimal value.
2. Fix $\theta^{(k+1)}$, $s^{(k)}$, and optimize over ϕ . Let $\phi^{(k+1)}$ be the optimal value.
3. Fix $\phi^{(k+1)}$, $\theta^{(k+1)}$. Let $s^{(k+1)}$ be appropriate scaling factors.

$k := k + 1$

until convergence

Many variations on this basic heuristic method are possible. We can, for example, add trust region constraints to each of the optimization steps, to limit the variables changes in each step. Another variation is to convexify (by, for example, linearizing) the jointly nonconvex problem, and solve in each step using linearized versions for the constraints and objective terms in the remaining variables; see, *e.g.*, [HHB99] and the references therein. We have already seen how the optimization over θ can be carried out efficiently. In many cases, the optimization over ϕ can also be carried efficiently, using, *e.g.*, LQG or some other controller or estimator design technique.

Since the joint problem is not convex, there is no guarantee that this heuristic converges to the global optimum. On the other hand the heuristic method appears to work well in practice.

5.3 Example: Networked Linear Estimator

To demonstrate the heuristic method for joint optimization described above, we apply it to the networked linear estimator described in §4.3. The design of the linear system and the communication system couple through the weighting matrix Q in (13). The alternating procedure for this problem becomes

```

given initial estimator gain  $K^{(0)}$  and resource allocations  $(P^{(0)}, W^{(0)}, b^{(0)})$ .
k:=0
repeat
  1. Fix estimator gain  $K^{(k)}$  and solve the problem (9) to obtain resource
     allocation  $(P^{(k+1)}, W^{(k+1)}, b^{(k+1)})$ .
  2. Update the covariance matrix  $R_q^{(k+1)}$  and compute new estimator gain
      $K^{(k+1)}$  as in (13) using weight matrix  $Q^{(k+1)} = (R_v + R_q^{(k+1)})^{-1}$ .
     k:=k+1
until bit allocation converges.

```

Note that the scaling factors are fixed in this example, since neither the bit allocations nor the estimator gain affect the signals that are quantized, hence the scaling factors.

When we apply the alternating optimization procedure to the example given in §4.3, the algorithm converges in six iterations, and we obtain very different resource allocation results from before. Figure 7 shows the distribution of rounded bit allocation. This result is intuitive: try to assign as much resources as possible to the best sensors, and the bad sensors only get minimum number of bits. The RMS estimation error of the joint design is reduced significantly, 80%, as shown in Table 1. In this table, $\mathbf{rms}(e)$ is the total RMS error, $\mathbf{rms}(e_q)$ is the RMS error induced by quantization noise, and $\mathbf{rms}(e_v)$ is the RMS error induced by sensor noise.

We can see that joint optimization reduces the estimation errors due to both quantization and sensor noise. In the case of equal resource allocation, the RMS error due to quantization is much larger than that due to sensor noise. After

RMS values	equal allocation	joint optimization	variable threshold rounding
$\text{rms}(e_q)$	3.5193×10^{-3}	0.3471×10^{-3}	0.3494×10^{-3}
$\text{rms}(e_v)$	1.0617×10^{-3}	0.6319×10^{-3}	0.6319×10^{-3}
$\text{rms}(e)$	3.6760×10^{-3}	0.7210×10^{-3}	0.7221×10^{-3}

Table 1: RMS estimation errors of the networked LS estimator.

the final iteration of the alternating convex optimization, the RMS error due to quantization is at the same level as that due to sensor noise. Also, because the in the relaxed problem, most bits are integers (either $\underline{b} = 5$ or $\bar{b} = 12$; see Figure 7), variable threshold rounding (which gives $t^* = 0.6797$) does not change the solution, or the performance, much.

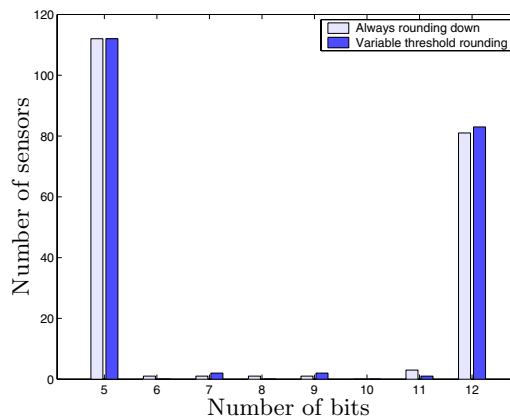


Fig. 7: Joint optimization of bit allocation and least-squares estimator

5.4 Example: LQG Control over Communication Networks

We now give a more complex example than the simple static, open-loop estimator described above. The situation is more complicated when the linear system is dynamic and involves feedback loops closed over the communication links. In this case, the RMS values of both control signals and output signals change when we re-allocate communication resources or adjust the controller. Hence, the alternating optimization procedure needs to include the step that modifies the scalings.

Basic System Setup First we consider the system setup in figure 8, where no communication links are included. The linear system has a state-space model

$$\begin{aligned} x(t+1) &= Ax(t) + B(u(t) + w(t)) \\ y(t) &= Cx(t) + v(t) \end{aligned}$$

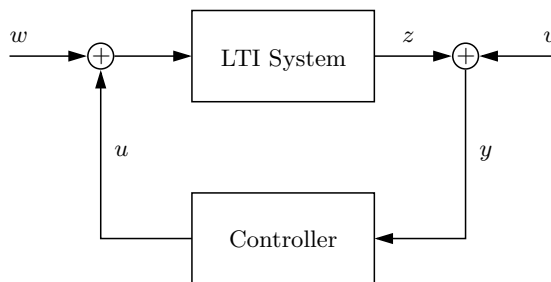


Fig. 8: Closed-loop control system without communication links.

where $u(t) \in \mathbf{R}^{M_u}$ and $y(t) \in \mathbf{R}^{M_y}$. Here $w(t)$ is the process noise and $v(t)$ is the sensor noise. Assume that $w(t)$ and $v(t)$ are independent zero-mean white noises with covariance matrices R_w and R_v respectively.

Our goal is to design the controller that minimizes the RMS value of $z = Cx$, subject to some upper bound constraints on the RMS values of the control signals:

$$\begin{aligned} & \text{minimize } \mathbf{rms}(z) \\ & \text{subject to } \mathbf{rms}(u_i) \leq \beta_i, \quad i = 1, \dots, M_u \end{aligned} \quad (14)$$

The limitations on the RMS values of the control signals are added to avoid actuator saturation.

It can be shown that the optimal controller for this problem has the standard estimated state feedback form,

$$\begin{aligned} \hat{x}(t+1|t) &= A\hat{x}(t|t-1) + Bu(t) + L(y(t) - C\hat{x}(t|t-1)) \\ u(t) &= -K\hat{x}(t|t-1) \end{aligned}$$

where K is the state feedback control gain and L is the estimator gain, found by solving the algebraic Riccati equations associated with an appropriately weighted LQG problem. Finding the appropriate weights, for which the LQG controller solves the problem (14), can be done via the dual problem; see, *e.g.*, [TM89, BB91].

Communications Setup We now describe the communications setup for the example. The sensors send their measurements to a central controller through a Gaussian multiple access channel, and the controller sends control signals to the actuators through a Gaussian broadcast channel, as shown in figure 9.

The linear system can be described as

$$\begin{aligned} x(t+1) &= Ax(t) + B(u(t) + w(t) + p(t)) \\ y_r(t) &= Cx(t) + v(t) + q(t), \end{aligned}$$

where p and q are quantization noises due to the bit rate limitations of the communication channels. Since these are modeled as white noises, we can include

the quantization noises in the process and measurement noises, by introducing the equivalent process noise and measurement noise

$$\tilde{w}(t) = w(t) + p(t), \quad \tilde{v}(t) = v(t) + q(t),$$

with covariance matrices

$$\begin{aligned} R_{\tilde{w}} &= R_w + \mathbf{diag} \left(\frac{s_{a,1}^2}{3} 2^{-2b_{a,1}}, \dots, \frac{s_{a,M_u}^2}{3} 2^{-2b_{a,M_u}} \right), \\ R_{\tilde{v}} &= R_v + \mathbf{diag} \left(\frac{s_{s,1}^2}{3} 2^{-2b_{s,1}}, \dots, \frac{s_{s,M_y}^2}{3} 2^{-2b_{s,M_y}} \right). \end{aligned} \quad (15)$$

Here b_a and b_s are number of bits allocated to the actuators and sensors.

The scaling factors can be found from the 3σ -rule, by computing the variance of the sensor and actuator signals. Hence, given the signal ranges and numbers of quantization bits, we can calculate $R_{\tilde{w}}$ and $R_{\tilde{v}}$, and then design a controller by solving (14). Notice that the signal ranges are determined by the RMS values, which in turn depend on the controller design. This intertwined relationship will show up in the iterative design procedures.

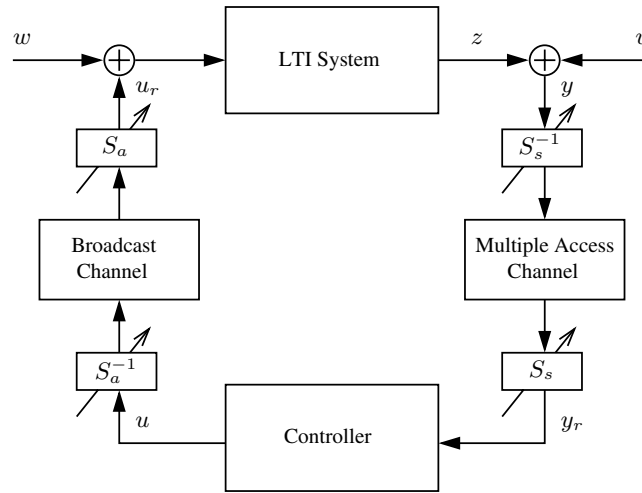


Fig. 9: Closed-loop control system over communication networks.

Iterative Procedure to Design a Controller with Uniform Bit Allocation First we allocate an equal number of bits to each actuator and sensor. This means that we assign power and bandwidth (in the case of FDMA) uniformly across all channels. We design a controller for such uniform resource allocation via the following iterative procedure (iterate on the scaling factors and the controller):

given $\beta_i = \mathbf{rms}(u_i)$ and estimated $\mathbf{rms}(z_j)$.
repeat
 1. Let $s_{a,i} = 3 \mathbf{rms}(u_i)$ and $s_{s,j} = 3 \mathbf{rms}(z_j)$, and compute R_w^\sim and R_v^\sim as in (15).
 2. Solve problem (14) and compute $\mathbf{rms}(u_i)$ and $\mathbf{rms}(z_j)$ of the closed-loop system.
until stopping criterion is satisfied.

If the procedure converges, the resulting controller variables K and L of this iterative design procedure will satisfy the constraints on the control signals.

The Alternating Optimization Procedure Our goal here is to do joint optimization of bit allocation and controller design. This involves an iteration procedure over controller design, scaling matrices update and bit allocation. The controller and scaling matrices designed for uniform bit allocation by the above iteration procedure can serve as a good starting point. Here is the alternating optimization procedure:

given $R_w, R_v, \beta_i = \mathbf{rms}(u_i)$ and $\mathbf{rms}(z_j)$ from the above iteration design procedure.
repeat
 1. Allocate bit rates $b_{a,i}, b_{s,j}$ and communication resources by solving a convex optimization problem of the form (8).
 2. Compute R_w^\sim and R_v^\sim as in (15), and find controller variables K and L by solving (14).
 3. Compute closed-loop system RMS values $\mathbf{rms}(u_i)$ and $\mathbf{rms}(z_j)$, then determine the signal ranges $s_{a,i}$ and $s_{s,j}$ by the 3σ rule.
until the RMS values $\mathbf{rms}(z_j)$ and bit allocation converges.

The convex optimization problem to be solved in step 1 depends on the communication system setup and resource constraints.

Numerical Example: Control of a Mass-Spring System Now we consider the specific example shown in figure 10. The position sensors on each mass send measurements $y_i = x_i + v_i$, where v_i is the sensor noise, to the controller through a Gaussian multiple access channel using FDMA. The controller receives data $y_{ri} = x_i + v_i + q_i$, where q_i is the quantization error due to bit rate limitation of the multiple access channel. The controller sends control signals u_j to actuators on each mass through a Gaussian broadcast channel using FDMA. The actual force acting on each mass is $u_{rj} = u_j + w_j + p_j$, where w_j is the exogenous disturbance force, and p_j is the quantization disturbance due to bit rate limitation of the broadcast channel. The mechanical system parameters are

$$m_1 = 10, \quad m_2 = 5, \quad m_3 = 20, \quad m_4 = 2, \quad m_5 = 15, \quad k = 1$$

The discrete-time system dynamics is obtained using a sampling frequency which is 5 times faster than the fastest mode of the continuous-time dynamics. The

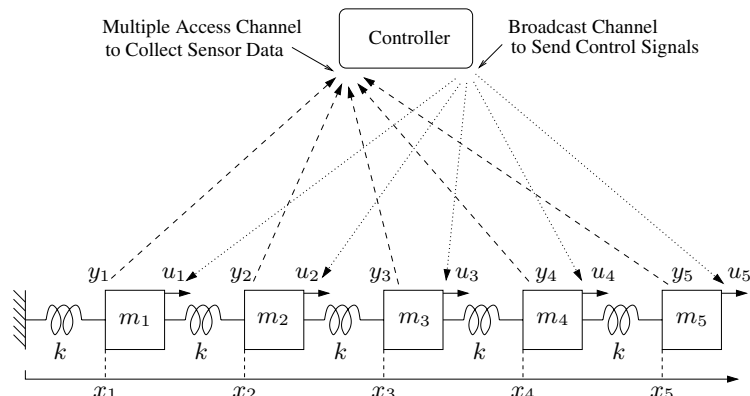


Fig. 10: Series-connected mass-spring system controlled over network.

independent zero mean noises w and v have covariance matrices $R_w = 10^{-6}I$ and $R_v = 10^{-6}I$ respectively. The actuators impose RMS constraints on the control signals:

$$\mathbf{rms}(u_i) \leq 1, \quad i = 1, \dots, 5.$$

For the Gaussian multiple access channel, the noise power density is $N = 0.1$, and the total power available is $P_{\text{mac,tot}} = 7.5$. For the Gaussian broadcast channel, the noise power density at each user is $N_i = 0.1$ for all i 's, and the total power available for all users is $P_{\text{bc,tot}} = 7.5$. All users of the multiple access channel and the broadcast channel share a total bandwidth of $W = 10$. The proportionality coefficient α in the capacity formula is set to 2. Finally, we impose a lower bound $\underline{b} = 5$ and an upper bound $\bar{b} = 12$ on the number of bits allocated to each quantizer.¹

First we allocate power and bandwidth evenly to all sensors and actuators, which results in a uniform allocation of 8 bits for each channel. We then designed a controller using the first iteration procedure based on this uniform resource allocation. This controller yields $\mathbf{rms}(u_i) = 1$ for all i 's, and the RMS-values of the output signal z are listed in Table 2.

Finally, we used the second iteration procedure to do joint optimization of bit allocation and controller design. The resulting resource allocation after four iterations is shown in figure 11. It can be seen that more bandwidth, and hence more bits are allocated to the broadcast channel than to the multiple access

¹ To motivate our choice of lower bound on the bit allocations, note that our system is critically stable and that the lower bound for stabilization given in [WB99, NE00, TSM98] is zero. In general, if we discretize an open-loop unstable continuous-time linear system using a sampling rate which is at least twice the largest magnitude of the eigenvalues (a traditional rule-of-thumb in the design of digital control systems [FPW90]), then the lower bound given in [WB99, NE00, TSM98] is less than one bit. The analysis in [WKL96] shows that $b_i \geq 3$ or 5 is usually high enough for assuming the white noise model for quantization errors.

RMS values	equal allocation	joint optimization	variable threshold rounding
$\mathbf{rms}(z_1)$	0.1487	0.0424	0.0438
$\mathbf{rms}(z_2)$	0.2602	0.0538	0.0535
$\mathbf{rms}(z_3)$	0.0824	0.0367	0.0447
$\mathbf{rms}(z_4)$	0.4396	0.0761	0.0880
$\mathbf{rms}(z_5)$	0.1089	0.0389	0.0346
$\mathbf{rms}(z)$	0.5493	0.1155	0.1258

Table 2: RMS-values of the output signal.

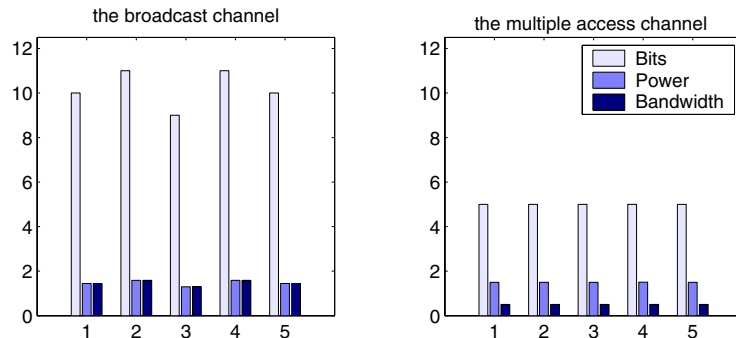


Fig. 11: Joint optimization of bit rates and linear control system.

channel. This means that the closed-loop performance is more sensitive to the equivalent process noises than to the equivalent sensor noises. The joint optimization resulted in $\mathbf{rms}(u_i) = 1$ for all i 's, and the RMS-values of the output signal z are listed in Table 2. At each step of the variable threshold rounding, we check the feasibility of the resource allocation problem. The optimal threshold found is $t^* = 0.6150$. Then we fix the integer bit allocation obtained with this threshold, and used the first iteration procedure to design the controller. We see a 77% reduction in RMS value over the result for uniform bit allocation, and the performance obtained by variable threshold rounding is quite close to that of the relaxed non-integer joint optimization.

6 Conclusions

We have addressed the problem of jointly optimizing the parameters of a linear system and allocating resources in the communication system that is used for transmitting sensor and actuator information. We considered a scenario where the coding and medium access scheme of the communication system is fixed, but the available communications resources, such as transmit powers and bandwidths, can be allocated to different channels in order to influence the achievable communication rates. To model the effect of limited communication rates on the performance of the linear system, we assumed conventional uniform quantization and used a simple white-noise model for quantization errors. We demonstrated that the problem of allocating communications resources to optimize the sta-

tionary performance of a fixed linear system (ignoring the integrality constraint) is often convex, hence readily solved. Moreover, for many important channel models, the communication resource allocation problem is separable except for a small number of constraints on the total communication resources. We illustrated how dual decomposition can be used to solve this class of problems efficiently, and suggested a variable threshold rounding method to deal with the integrality of bit allocations. The problem of jointly allocating communication resources and designing the linear system is in general not convex, but is often convex in subsets of variables while others are fixed. We suggested an iterative heuristic for the joint design problem that exploits this special structure, and demonstrated its effectiveness on the two examples: the design of a networked linear estimator, and the design of a multivariable networked LQG controller.

Acknowledgments

The authors are grateful to Wei Yu and Xiangheng Liu for helpful discussions.

References

- [BB91] S. Boyd and C. Barratt. *Linear Controller Design: Limits of Performance*. Prentice-Hall, 1991.
- [Ber99] D. P. Bertsekas. *Nonlinear Programming*. Athena Scientific, second edition, 1999.
- [BL00] R. Brockett and D. Liberzon. Quantized feedback stabilization of linear systems. *IEEE Transactions on Automatic Control*, 45(7):1279–1289, July 2000.
- [BM97] V. Borkar and S. Mitter. LQG control with communication constraints. In A. Paulraj, V. Roychowdhury, and C. D. Schaper, editors, *Communications, Computations, Control and Signal Processing, a Tribute to Thomas Kailath*, pages 365–373. Kluwer, 1997.
- [BV04] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004. Available at <http://www.stanford.edu/~boyd/cvxbook.html>.
- [CG01] S. T. Chung and A. J. Goldsmith. Degrees of freedom in adaptive modulation: A unified view. *IEEE Transactions on Communications*, 49(9):1561–1571, 2001.
- [CT91] T. Cover and J. Thomas. *Elements of Information Theory*. John Wiley & Sons, 1991.
- [Cur70] R. E. Curry. *Estimation and Control with Quantized Measurements*. The MIT Press, 1970.
- [Del90] D. F. Delchamps. Stabilizing a linear system with quantized state feedback. *IEEE Transactions on Automatic Control*, 35(8):916–924, August 1990.
- [EM01] N. Elia and S. K. Mitter. Stabilization of linear systems with limited information. *IEEE Transactions on Automatic Control*, 46(9):1384–1400, September 2001.
- [FPW90] G. F. Franklin, J. D. Powell, and M. L. Workman. *Digital Control of Dynamic Systems*. Addison-Wesley, 3rd edition, 1990.
- [Gol99] A. Goldsmith. *Course reader for EE359: Wireless Communications*. Stanford University, 1999.

- [HHB99] A. Hassibi, J. P. How, and S. P. Boyd. A path-following method for solving BMI problems in control. In *Proceedings of American Control Conference*, volume 2, pages 1385–9, June 1999.
- [KH94] P. T. Kabamba and S. Hara. Worst-case analysis and design of sampled-data control systems. *IEEE Transactions on Automatic Control*, 38(9):1337–1357, September 1994.
- [LG01] L. Li and A. J. Goldsmith. Capacity and optimal resource allocation for fading broadcast channels: Part I: Ergodic capacity. *IEEE Transactions on Information Theory*, 47(3):1103–1127, March 2001.
- [NBW98] J. Nilsson, B. Bernhardsson, and B. Wittenmark. Stochastic analysis and control of real-time systems with random time delays. *Automatica*, 34(1):57–64, 1998.
- [NE98] G. N. Nair and R. J. Evans. State estimation under bit-rate constraints. In *Proc. IEEE Conference on Decision and Control*, pages 251–256, Tampa, Florida, 1998.
- [NE00] G. N. Nair and R. J. Evans. Stabilization with data-rate-limited feedback: tightest attainable bounds. *Systems & Control Letters*, 41:49–56, 2000.
- [Özg89] Ü. Özgüner. Decentralized and distributed control approaches and algorithms. In *Proceedings of the 28th Conference on Decision and Control*, pages 1289–1294, Tampa, Florida, December 1989.
- [SSK99] K. Shoarinejad, J. L. Speyer, and I. Kanellakopoulos. An asymptotic optimal design for a decentralized system with noisy communication. In *Proc. of the 38th Conference on Decision and Control*, Phoenix, Arizona, December 1999.
- [SW90] R. E. Skelton and D. Williamson. Guaranteed state estimation accuracies with roundoff error. In *Proceedings of the 29th Conference on Decision and Control*, pages 297–298, Honolulu, Hawaii, December 1990.
- [TM89] H. T. Toivonen and P. M. Mäkilä. Computer-aided design procedure for multiobjective LQG control problems. *Int. J. Control*, 49(2):655–666, February 1989.
- [TSM98] S. Tatikonda, A. Sahai, and S. Mitter. Control of LQG systems under communication constraints. In *Proc. IEEE Conference on Decision and Control*, pages 1165–1170, December 1998.
- [WB97] W. S. Wong and R. W. Brockett. Systems with finite communication bandwidth constraints I: state estimation problems. *IEEE Transactions on Automatic Control*, 42:1294–1299, 1997.
- [WB99] W. S. Wong and R. W. Brockett. Systems with finite communication bandwidth constraints – II: Stabilization with limited information feedback. *IEEE Transactions on Automatic Control*, 44:1049–1053, May 1999.
- [Wil85] D. Williamson. Finite wordlength design of digital Kalman filters for state estimation. *IEEE Transactions on Automatic Control*, 30(10):930–939, October 1985.
- [WK89] D. Williamson and K. Kadiman. Optimal finite wordlength linear quadratic regulation. *IEEE Transactions on Automatic Control*, 34(12):1218–1228, December 1989.
- [WKL96] B. Widrow, I. Kóllar, and M.-L. Liu. Statistical theory of quantization. *IEEE Trans. Instrumentation and Measurements*, 45:353–361, April 1996.
- [XHH00] L. Xiao, A. Hassibi, and J. P. How. Control with random communication delays via a discrete-time jump system approach. In *Proc. American Control Conf.*, Chicago, IL, June 2000.