

GLOBALLY OPTIMAL PERFORMANCE OF FEEDBACK CONTROL SYSTEMS WITH LIMITED COMMUNICATION OVER NOISY CHANNELS

ADITYA MAHAJAN AND DEMOSTHENIS TENEKETZIS*

Abstract. A discrete time stochastic feedback control system consisting of a non-linear plant, a sensor, a controller, and a noisy communication channel between the sensor and the controller is considered. The sensor has limited memory and at each time, it transmits an encoded symbol over the channel and updates its memory. The controller receives a noise-corrupted copy of the transmitted symbol. It generates a control action based on all its past observations and all its past actions. This control action is fed back to the plant. At each time instant the system incurs an instantaneous cost depending on the state of the plant and the control action. The objective is to choose encoding, memory update and control strategies to minimize: an expected total cost over a finite horizon, or an expected discounted cost over an infinite horizon, or an average cost per unit time over an infinite horizon. A solution methodology to obtain a sequential decomposition of the global optimization problem is developed. This solution methodology is extended to the case when the sensor makes an imperfect observation of the state of the plant.

Key words. optimal control over noisy communication, sequential stochastic control, decentralized optimal control, non-classical information structures, dynamic teams, common knowledge

AMS subject classifications. 93E03, 93E20, 93A14, 62B05, 49N30

1. Introduction.

1.1. Preliminaries and literature overview. Recent advances in network and communication technologies have led to an increasing interest in networked control systems (NCS) (see the papers in [1]), in particular, in understanding the limitations imposed upon a feedback control system by the presence of a communication channel in the loop. Most researchers have concentrated on stability analysis of the system. The problem of stabilization of a plant with finite data rate feedback was investigated in [4, 5, 7, 9, 11, 13, 15, 18, 22–25, 27, 47]. See [26] for a unified overview of stabilization with finite data rate feedback. LQG stability of various systems (deterministic, stochastic, stable, unstable) under various kinds of communication constraints (noisy and noiseless channel) was considered in [35–37, 39]. Stability of an unstable plant over AWGN channel subject to input power constraints was considered in [6]. Fundamental asymptotic limitations of feedback for a linear time invariant plant and arbitrary time-invariant causal feedback were investigated in [19, 20] using an information theoretic formulation. In retrospect, the connection between stability and information theory is not surprising since stability as well as the information theoretic notions of source entropy and channel capacity are asymptotic concepts.

Certain applications, like vehicular traffic control and biomedical applications, require performance metrics different from the asymptotic metrics of stability. In this paper we consider the class of additive performance metrics, where the total cost is the sum of costs along the entire path. In order to determine optimal performance NCS, both the transient and the steady state behaviours need to be chosen optimally. The asymptotic notions of source entropy and channel capacity, and the asymptotic

*The authors are with the department of EECS at the University of Michigan, Ann Arbor, MI 48109-2212, USA (email: {adityam,teneket}@eecs.umich.edu). A preliminary version of this paper appeared in the proceedings of the 45th IEEE Conference on Decision and Control, San Diego, December 2006.

results on stability are not appropriate for evaluating the transient performance and consequently not appropriate for performance evaluation.

The problem of optimal performance has received less attention than stabilization in the literature. The problems considered in the literature can be classified on the basis of their plant dynamics (linear or non-linear), the nature of the communication channel (rate-limited noiseless channel or noisy channel), and the information structure (classical or non-classical information structures). Optimal performance of a linear plant with rate-limited noiseless communication channel was considered in [21,30]: in [21] the plant disturbance is Gaussian and the controller is memoryless; in [30] the plant is undisturbed and the controller has perfect recall. Optimal performance of a linear plant with Gaussian disturbance, either a rate-limited noiseless channel or a Gaussian memoryless channel, and various information structures at the encoder was considered in [38]. Optimal performance of a non-linear plant and a noisy channel with noiseless feedback from the output of the channel to the encoder was considered in [40].

The most important feature in problems of optimal performance of NCS is whether the encoder knows the information available at the decoder/controller or not. We can classify problems into two cases on the basis of the presence or absence of this feature: case 1, when the encoder has access to all the information available at the decoder/controller, and case 2, when it does not. In case 1 the problem of determining optimal performance can be reduced to a centralized stochastic control problem from the encoder's point of view. Such a reduction is not possible in case 2. Consequently, in case 1 the encoder knows how the decoder/controller will interpret its messages; in case 2, it does not. So efficient communication between the encoder and decoder/controller is easier in case 1 than in case 2. Hence, in determining optimal strategies for the encoder and the controller in case 2 is a considerably more difficult problem than in case 1.

The models of [21,30,40] and the instances in [38] where there are noiseless channels as well as the instance of information pattern A (see [38, pg. 1550] for definition of information pattern A) belong to case 1. In all these situations optimal encoding and control strategies have been determined. The model in [38] with information pattern B (see [38, pg. 1550] for definition of information pattern B) belong to case 2. In this situation only sub-optimal encoding and control strategies have been proposed.

In this paper we consider a non-linear plant with a noisy communication channel. Our model belongs to case 2. We model the performance analysis of NCS as a stochastic control problem. We study the simplest NCS—a network with only two nodes with a noisy communication link between them. We identify the structure of optimal controllers and develop a general methodology for determining globally optimal encoding and control strategies for finite and infinite horizon problems. We show that even for “well behaved” infinite horizon problems, stationary designs are not necessarily globally optimal.

1.2. Features of the problem. We consider a discrete-time feedback control system with a communication channel between the sensor and the controller, as shown in Figure 2.1. Such problems arise when the plant and the controller are geographically separated. We assume that there is a noisy discrete memoryless channel between the sensor and the controller. (The rate limited communication channel is the degenerate case where the channel is noiseless). We model problems in which the sensor has limited resources in terms of the power at which it can transmit and the data it can store and process. The encoder connected with the sensor is assumed to have a finite

memory, thus it can not remember all its past observations and actions, and at each stage, must selectively shed some information. At each time, the sensor generates a symbol using its current observation and the contents of its memory, and transmits it over the noisy channel to the controller. We assume that there is no resource constraint at the controller. It has infinite memory and infinite power. Thus we assume that the controller has perfect recall—it remembers everything that it has seen and done in the past—and the communication channel between the controller and the plant is noiseless¹. At each stage t the system incurs an instantaneous cost depending on the state of the plant at t and the control action at t . The objective is to choose globally optimal encoding, memory update, and control strategies to minimize: the expected total cost over a finite horizon, or the expected discounted cost over an infinite horizon, or the expected average cost per unit time over an infinite horizon.

The problem has two decision-makers, the sensor and the controller. Due to the noise in the communication channel, the sensor and the controller have different information about what is happening in nature. Due to the finite memory at the sensor, the sensor forgets information and at any given time instant the sensor may not know what actions it took in the past and why it took those actions. These two considerations, the noise in the channel and the finite memory at the sensor, result in a *decentralized* control problem. There is no known solution methodology to solve infinite horizon decentralized stochastic control problems.

Markov decision theory [14] provides a solution methodology for *centralized* stochastic control problems. For centralized problems with imperfect observations, Markov decision theory shows that there is no loss of optimality in taking a control action based on the controller’s belief about the state of the plant. The belief is obtained using all the data available at the controller. Centralization of information and perfect recall at the controller are crucial for this idea to work. Consequently this idea does not extend to decentralized control problems: decentralization of information implies that one decision-maker cannot infer the data available with the other decision-makers and therefore cannot infer their beliefs. So if all decision-maker act according to their beliefs about the state of the plant, they will act in an inconsistent manner, and the system will not achieve the globally optimal performance. Hence, Markov decision theory is not appropriate for this problem.

Orthogonal search [29] techniques provide a solution methodology for decentralized stochastic control problem. There are different variations of orthogonal search algorithm, but the key idea is the following. Initialize by arbitrarily choosing the decision strategies of all agents; then pick an agent, say i , and determine the *best response* of agent i to the strategies of the other agents. Fix this best response as agent i ’s strategy. Next pick another agent j , $j \neq i$, and update agent j ’s strategy by its best response to the other agents’ strategies. Continue this way. If this procedure converges, the resultant strategies are member by member optimal, i.e., unilateral deviations by a single agent do not improve the system’s performance. Fictitious play techniques [8, 28, 34] are philosophically similar to orthogonal search and result in member by member optimal solutions. Since, decentralized stochastic control problems are, in general, non-convex in strategy space, the above procedure may not converge to globally optimal strategies; that is, it does not guarantee that there do not exist any other tuple of strategies for all agents that outperforms the member by member optimal strategies found by the above procedure. Thus, orthogonal

¹In the sequel we show that assuming a noiseless feedback channel does not entail any loss of generality.

search cannot be used to obtain globally optimal strategies for the problem under consideration.

Witsenhausen's standard form [44] is the only known solution methodology for general sequential decentralized stochastic control problems. It proceeds by converting the problem into a *standard form*, and then obtains a sequential decomposition for the standard form. The standard form is a finite horizon stochastic control problem whose state evolution satisfies some properties, the cost is a stopping cost incurred at the last time step, and the cost has a certain measurability properties. Since all the cost is incurred at the last time step in the standard form, infinite horizon problems cannot be converted into the standard form. Hence the solution methodology of [44] is not appropriate for the problem under consideration.

1.3. Contributions of the paper. The main contribution of this paper is in providing a solution methodology for sequentially determining globally optimal real-time encoding, memory update, and control strategies for feedback control systems with limited communication over noisy channels. To the best of our knowledge, the methodology developed in this paper is the first one to provide a sequential decomposition for the aforementioned class of problems. The methodology proceeds in two steps. In the first step, we obtain qualitative properties of optimal controllers. In the second step we use the qualitative properties of the first step to identify information states sufficient for performance evaluation (also called sufficient statistic for control) to obtain a sequential decomposition of the problem. *The main conceptual difficulty in the problem is identifying appropriate information states in the second step*; once appropriate information states are identified, obtaining a sequential decomposition is straight forward. We would like to emphasize that identifying appropriate information states for performance evaluation is nontrivial; the difficulty can be judged from the fact that decentralized stochastic control problems have been investigated since the early 1970s, and even now there is no known solution methodology to obtain information states sufficient for performance analysis for these problems. The results of this paper provide a solution methodology for a hitherto unsolved class of decentralized stochastic control problems and explains why this methodology works. This methodology may be useful for other classes of decentralized stochastic control problems.

1.4. Organization of the paper. The remainder of this paper is organized as follows. We formulate the performance analysis of feedback control systems with limited communication over a noisy channel as a decentralized stochastic optimization problem. To illustrate the key concepts associated with our solution methodology we first consider the finite horizon problem. In Section 2, we establish structural results of an optimal controller and obtain a methodology for sequentially global optimization of the encoding, memory update and control strategies for the finite horizon problem. We provide an explanation of the methodology in Section 3. In Section 4 we extend the methodology to infinite horizon problems. In Section 5 we consider the case of uncountable state space. In Section 6 we consider the feedback control problem when the encoder has imperfect observation of the state of the plant and extend the results of Sections 2 and 4 to this problem. We conclude in Section 7.

1.5. Notation. Throughout this paper we use the following notation. Uppercase letters (X, Y, Z) denote random variables, lowercase letters (x, y, z) denote their realizations, and calligraphic letters $(\mathcal{X}, \mathcal{Y}, \mathcal{Z})$ denote their alphabets. For random variables and functions, x^t is a short hand for x_1, \dots, x_t . $\mathbb{E}\{\cdot\}$ denotes the expectation.

tation of a random variable, $\Pr(\cdot)$ denotes the probability of an event, and $\mathbf{1}[\cdot]$ denotes the indicator function of a statement. To denote the expectation or probability of a random variable or an event that depends on a function φ , we use $\mathbb{E}\{\cdot | \varphi\}$ and $\Pr(\cdot | \varphi)$, respectively. We have chosen this slightly unusual notation because we want to keep track of all the functional dependencies and the conventional notation of $\mathbb{E}^\varphi\{\cdot\}$ and $\Pr^\varphi(\cdot)$ is too cumbersome.

2. The Finite Horizon Problem.




FIG. 2.1. Feedback control system with noisy communication

2.1. Problem formulation. Consider a discrete-time feedback control system of Figure 2.1 which operates for a horizon T . The state evolution is given by

$$(2.1) \quad X_{t+1} = f(X_t, U_t, W_t),$$

where f is the *plant evolution function* and the variables X_t, U_t, W_t denote the state of the plant, the control action and the plant disturbance respectively, at time t . We assume that all variables are finite valued. For all t , X_t takes values in a finite set \mathcal{X} , U_t takes values in a finite set \mathcal{U} and W_t takes values in a finite set \mathcal{W} . The initial state X_1 is a random variable with PMF P_{X_1} . The random variables W_1, \dots, W_T are i.i.d. (independent and identically distributed) with PMF P_W and are also independent of X_1 .

The sensor, consisting of an encoder and a memory, makes perfect observations of the state of the plant. At each time instant t the encoder generates an encoded symbol Z_t , taking values in a finite set \mathcal{Z} , as follows:

$$(2.2) \quad Z_t = c_t(X_t, M_{t-1}),$$

where c_t is the *encoding function* at time t and M_{t-1} denotes the content of the sensor's memory at $t-1$. M_t takes values in a finite set \mathcal{M} and is updated according to

$$(2.3) \quad M_t = l_t(X_t, M_{t-1}),$$

where l_t is the *memory update function* at time t . Observe that the sensor has a finite size memory and although it makes perfect observations of the state of the plant, it can not store all the past observations. Thus, it does not have perfect recall and at each stage it must selectively shed information.

The encoded symbol Z_t is transmitted over a noisy communication channel and a channel output Y_t is generated according to

$$(2.4) \quad Y_t = h(Z_t, N_t),$$

where h is the *channel function* and N_t denotes the channel noise. Y_t takes values in a finite set \mathcal{Y} and N_t takes values in a finite set \mathcal{N} . The sequence of random variables N_1, \dots, N_T is i.i.d. with given PMF P_N and is also independent of X_1, W_1, \dots, W_T .

The controller observes the channel outputs and generates a control action U_t as follows:

$$(2.5) \quad U_t = g_t(Y^t, U^{t-1}),$$

where g_t is the *control law* at time t . U_t takes values in a finite set \mathcal{U} . A uniformly bounded cost function $\rho : \mathcal{X} \times \mathcal{U} \rightarrow [0, K]$, where $K < \infty$ is given. At each t , an instantaneous cost $\rho(X_t, U_t)$ is incurred.

The collection $(\mathcal{X}, \mathcal{W}, \mathcal{M}, \mathcal{Z}, \mathcal{N}, \mathcal{Y}, \mathcal{U}, P_{X_1}, P_W, P_N, f, h, \rho, T)$ is called a *perfect observation system*. The choice of (C, L, G) , $C := (c_1, \dots, c_T)$, $L := (l_1, \dots, l_T)$, $G := (g_1, \dots, g_T)$, is called a *design*.

The performance of a design is quantified by the expected total cost under that design and is given by

$$(2.6) \quad \mathcal{J}_T(C, L, G) := \mathbb{E} \left\{ \sum_{t=1}^T \rho(X_t, U_t) \mid C, L, G \right\},$$

where the expectation in (2.6) is with respect to a joint measure on $(X_1, \dots, X_T, U_1, \dots, U_T)$ generated by P_W, P_N, f, h and the choice of design (C, L, G) . We are interested in the following optimization problem:

PROBLEM 2.1. *Given a perfect observation system $(\mathcal{X}, \mathcal{W}, \mathcal{M}, \mathcal{Z}, \mathcal{N}, \mathcal{Y}, \mathcal{U}, P_{X_1}, P_W, P_N, f, h, \rho, T)$, choose a design (C^*, L^*, G^*) such that*

$$(2.7) \quad \mathcal{J}_T(C^*, L^*, G^*) = \mathcal{J}_T^* := \min_{C, L, G \in \mathcal{C}^T \times \mathcal{L}^T \times \mathcal{G}^T} \mathcal{J}_T(C, L, G),$$

where $\mathcal{C}^T := \mathcal{C} \times \dots \times \mathcal{C}$ (T times), \mathcal{C} is the space of functions from $\mathcal{X} \times \mathcal{M}$ to \mathcal{Z} , $\mathcal{L}^T := \mathcal{L} \times \dots \times \mathcal{L}$ (T times), \mathcal{L} is the space of functions from $\mathcal{X} \times \mathcal{M}$ to \mathcal{M} , $\mathcal{G}^T := \mathcal{G}_1 \times \dots \times \mathcal{G}_T$, and \mathcal{G}_t is the space of functions from $\mathcal{Y}^t \times \mathcal{U}^{t-1}$ to \mathcal{U} .

Remarks.

1. There is no loss of generality in assuming a noiseless channel between the controller and the plant. Suppose that the channel between the controller and the plant is noisy. Let the input \hat{U}_t to the plant be a noise-corrupted version of U_t given by

$$(2.8) \quad \hat{U}_t = \hat{h}(U_t, \hat{N}_t),$$

where \hat{h} is the feedback channel, and \hat{N}_t denotes the noise in the feedback channel. $\hat{N}_1, \dots, \hat{N}_T$ is a sequence of independent variables that is also independent of X_1, W_1, \dots, W_T and N_1, \dots, N_T .² Then this model can be transformed into one equivalent to (2.1)–(2.5) by setting

$$(2.9) \quad \hat{W}_t = (W_t, \hat{N}_t),$$

$$(2.10) \quad X_{t+1} = f(X_t, \hat{h}(U_t, \hat{N}_t), W_t) := \hat{f}(X_t, U_t, \hat{W}_t).$$

²We only require $\hat{W}_1, \dots, \hat{W}_T$, where $\hat{W}_t = (W_t, \hat{N}_t)$, to be an independent process. So, \hat{N}_t need not be independent of W_t .




FIG. 2.2. Problem 2.1 as a sequential stochastic optimization problem. This figure shows the ordering relation between the system variables, design rules, and information states.

Thus, without any loss of generality we can assume a noiseless feedback channel.

2. A globally optimal design for Problem 2.1 always exists because there are finitely many designs and we can always choose one with the best performance.

2.2. Salient features of the problem. Problem 2.1 is a decentralized multi-agent stochastic optimization problem. There are two agents, the sensor and the controller, that have different information about the system and a common objective which is to minimize an expected total cost over a finite horizon. Multi-agent problems in which the agents have a common objective are called team problems [17]. Team problems are further classified as static teams or dynamic teams on the basis of their information structure. See [43] for a definition of information structure (also called information pattern). In static teams the actions taken by one agent do not affect the information structure of the other agents; in dynamic teams they do. In Problem 2.1, the actions taken by the sensor affect the observations of the controller and the actions taken by the controller affect the observations of the sensor; furthermore, the sensor and the controller have different information about the system. Moreover, due to the finite memory at the sensor and the noise in the channel, Problem 2.1 has a strictly non-classical information structure; thus Problem 2.1 is a dynamic team.

Determining globally optimal strategies for dynamic teams is difficult because they are, in general, non-convex functional optimization problems having a complex interdependence among their decision rules [12]. As pointed out in the introduction, Markov decision theory, orthogonal search, and standard form are not appropriate for solving infinite horizon dynamic team problems.

The solution concept that we are looking for is to decompose the global optimization problem into a sequence of nested optimization sub-problems where each sub-problem is easier to solve than the original problem. This is called *sequential decomposition* and it exponentially reduces the search complexity of finding an op-

timal strategy. A crucial step in obtaining a sequential decomposition of the global optimization problem is to identify information states that are sufficient for performance evaluation. Properties that such states must satisfy are explained in [16]. All the known techniques of identifying appropriate information states, *viz.*, Markov decision theory, orthogonal search and standard form are not appropriate for infinite horizon dynamic team problems. The information states in Markov decision theory — the conditional probability densities of the state given all the past observations and all the past control actions — works only when there is a single controller with perfect recall; so, they are inappropriate for dynamic teams. The information states in orthogonal search are obtained under the assumption that the strategies of other agents are fixed. These information states only determine member by member optimal strategies, so they are not appropriate for determining globally optimal strategies for dynamic teams. The information states in Witsenhausen’s standard form belong to a space that increases with time; hence, it is not appropriate for infinite horizon problems. Thus a new methodology for identifying information states is needed for the problem under consideration. We provide one such methodology in this paper.

The sequential order in which the system variables are generated is the key to understanding the solution methodology that we present in this paper. For this purpose we need to refine the notion of time. We call each step of the system a *stage*. At any stage t , we consider three time instants³ t^+ , $(t + 1/2)$, and $(t + 1)^-$. For ease of notation, we will denote these time instants by \underline{t} , t , and \bar{t} respectively. From now on, we will assume that the system has three agents—the encoder, the memory update, and the controller—even though the encoder and the memory update are located in the same device and have the same information. We assume that the sensor encodes just after \underline{t} , the sensor’s memory is updated just after t , and the controller takes a control action just after \bar{t} . The order in which the variables are generated in the system is shown in Figure 2.2. Since the ordering of the decision makers can be done independently of the realization of the system variables, the problem is a *sequential* stochastic optimization problem [45].

To obtain a sequential decomposition of Problem 2.1, we proceed in two steps. In step one, we derive structural properties of optimal controllers. In step two, we use the structural results of step one to identify an information state sufficient for performance evaluation, transform Problem 2.1 into an equivalent deterministic optimization problem and obtain a sequential decomposition for this equivalent problem. This sequential decomposition gives an algorithm to obtain an optimal design for Problem 2.1.

As pointed out in the introduction, step two is the crucial step. The key difficulty in step two is to identify an information state appropriate for performance evaluation. Even when the structural results of step one are available, identifying such an information state is a highly nontrivial task. Once an appropriate information state is identified, the transformation to a deterministic problem and the sequential decomposition follow.

2.3. Structure of optimal controllers. In this section we present structural properties of optimal controllers. We first define random variables that capture the information available just before the decision rules c_t , l_t , and g_t act on the system.

³The actual values of these time instants is not important; we just need three values in increasing order.

DEFINITION 2.1. Let \underline{I}_t , I_t and \bar{I}_t denote the information available at the controller at time \underline{t} , t , and \bar{t} respectively. Specifically

1. $\underline{I}_t := (Y^{t-1}, U^{t-1}, c^{t-1}, l^{t-1}, g^{t-1})$.
2. $I_t := (Y^t, U^{t-1}, c^t, l^{t-1}, g^{t-1})$.
3. $\bar{I}_t := (Y^t, U^{t-1}, c^t, l^t, g^{t-1})$.

We have included the past decision rules in the definition of information because the distribution of the random variables depends on the choice of the past decision rules. Observe that

$$(2.11) \quad \underline{I}_t = (\bar{I}_{t-1}, U_{t-1}, g_{t-1}), \quad I_t = (\underline{I}_t, Y_t, c_t), \quad \text{and} \quad \bar{I}_t = (I_t, l_t).$$

Next we define the belief of the controller about the state of the plant and the memory contents of the sensor at time \underline{t}^- , t^- , and \bar{t}^- .

DEFINITION 2.2. Let \underline{B}_t , B_t , and \bar{B}_t be random vectors defined as follows:

1. $\underline{B}_t(x, m) := \Pr(X_t = x, M_{t-1} = m | \underline{I}_t)$.
2. $B_t(x, m) := \Pr(X_t = x, M_{t-1} = m | I_t)$.
3. $\bar{B}_t(x, m) := \Pr(X_t = x, M_t = m | \bar{I}_t)$.

For any particular realization \underline{i}_t of \underline{I}_t , that is, for any particular realization y^{t-1}, u^{t-1} of Y^{t-1}, U^{t-1} and arbitrary (but fixed) choice of c^{t-1}, l^{t-1} and g^{t-1} , the realization \underline{b}_t of \underline{B}_t is a PMF on $\mathcal{X} \times \mathcal{M}$. If \underline{I}_t is a random vector, then \underline{B}_t is a random vector belonging to $\mathcal{P}^{\mathcal{X} \times \mathcal{M}}$, the space of PMFs on $\mathcal{X} \times \mathcal{M}$. Similar interpretations hold for B_t and \bar{B}_t .

The random vectors \underline{B}_t , B_t , and \bar{B}_t represent the belief of the controller about the state of the plant and the encoder's memory content at \underline{t} , t , and \bar{t} , respectively. The sequential ordering of these beliefs with respect to the other variables in the system are shown in Figure 2.2. The time evolution of these beliefs are coupled as follows.

LEMMA 2.3. For each stage t , there exist deterministic functions \underline{F} , F , and \bar{F} such that

1. $\underline{B}_t = \underline{F}(\bar{B}_{t-1}, U_{t-1})$.
2. $B_t = F(\underline{B}_t, Y_t, c_t)$.
3. $\bar{B}_t = \bar{F}(B_t, l_t)$.

Proof.

1. Consider a component of \underline{b}_t ,

$$(2.12) \quad \begin{aligned} \underline{b}_t(x_t, m_{t-1}) &= \Pr(X_t = x_t, M_{t-1} = m_{t-1} | \underline{i}_t) \\ &= \Pr(X_t = x_t, M_{t-1} = m_{t-1} | \bar{i}_{t-1}, u_{t-1}, g_{t-1}) \\ &= \frac{\Pr(X_t = x_t, M_{t-1} = m_{t-1}, U_{t-1} = u_{t-1} | \bar{i}_{t-1}, g_{t-1})}{\sum_{(x'_t, m'_{t-1}) \in \mathcal{X} \times \mathcal{M}} \Pr(X_t = x'_t, M_{t-1} = m'_{t-1}, U_{t-1} = u_{t-1} | \bar{i}_{t-1}, g_{t-1})}. \end{aligned}$$

Now consider

$$(2.13) \quad \begin{aligned} &\Pr(X_t = x_t, M_{t-1} = m_{t-1}, U_{t-1} = u_{t-1} | \bar{i}_{t-1}, g_{t-1}) \\ &= \Pr(x_t, m_{t-1}, u_{t-1} | \bar{i}_{t-1}, g_{t-1}) \\ &= \sum_{x_{t-1} \in \mathcal{X}} \Pr(x_{t-1}, m_{t-1} | \bar{i}_{t-1}, g_{t-1}) \\ &\quad \times \Pr(u_{t-1} | x_{t-1}, m_{t-1}, \bar{i}_{t-1}, g_{t-1}) \\ &\quad \times \Pr(x_t | x_{t-1}, m_{t-1}, u_{t-1}, \bar{i}_{t-1}, g_{t-1}) \end{aligned}$$

$$\begin{aligned}
& \stackrel{(a)}{=} \sum_{x_{t-1} \in \mathcal{X}} \Pr(x_{t-1}, m_{t-1} | \bar{i}_{t-1}) \mathbf{1} [u_{t-1} = g_{t-1}(y^{t-1}, u^{t-2})] \\
& \quad \times \Pr(x_t | x_{t-1}, u_{t-1}) \\
(2.14) \quad & = \mathbf{1} [u_{t-1} = g_{t-1}(y^{t-1}, u^{t-2})] \\
& \quad \times \sum_{x_{t-1} \in \mathcal{X}} \bar{b}_{t-1}(x_{t-1}, m_{t-1}) \Pr(x_t | x_{t-1}, u_{t-1}),
\end{aligned}$$

where equality (a) follows from (2.1) and (2.5) and $\mathbf{1}[\cdot]$ is the indicator function. Substitute equation (2.14) in equation (2.12) and cancel $\mathbf{1}[u_{t-1} = g_{t-1}(y^{t-1}, u^{t-2})]$ from the numerator and the denominator, giving

$$(2.15) \quad \underline{b}_t(x_t, m_{t-1}) = \frac{\sum_{x_{t-1} \in \mathcal{X}} \bar{b}_{t-1}(x_{t-1}, m_{t-1}) \Pr(x_t | x_{t-1}, m_{t-1})}{\sum_{(x'_t, x'_{t-1}, m'_{t-1}) \in \mathcal{X} \times \mathcal{X} \times \mathcal{M}} \bar{b}_{t-1}(x'_{t-1}, m'_{t-1}) \Pr(x'_t | x'_{t-1}, m'_{t-1})}.$$

Hence,

$$(2.16) \quad \underline{b}_t = \underline{F}(\bar{b}_{t-1}, u_{t-1}),$$

where \underline{F} is determined by (2.15).

2. Consider a component of b_t ,

$$\begin{aligned}
(2.17) \quad b_t(x_t, m_{t-1}) &= \Pr(X_t = x_t, M_{t-1} = m_{t-1} | i_t) \\
&= \Pr(X_t = x_t, M_{t-1} = m_{t-1} | \underline{i}_t, y_t, c_t) \\
&= \frac{\Pr(X_t = x_t, M_{t-1} = m_{t-1}, Y_t = y_t | \underline{i}_t, c_t)}{\sum_{(x'_t, m'_{t-1}) \in \mathcal{X} \times \mathcal{M}} \Pr(X_t = x'_t, M_{t-1} = m'_{t-1}, Y_t = y_t | \underline{i}_t, c_t)}.
\end{aligned}$$

Now consider

$$\begin{aligned}
(2.18) \quad & \Pr(X_t = x_t, M_{t-1} = m_{t-1}, Y_t = y_t | \underline{i}_t, c_t) \\
&= \Pr(x_t, m_{t-1}, y_t | \underline{i}_t, c_t) \\
&= \Pr(x_t, m_{t-1} | \underline{i}_t, c_t) \Pr(y_t | x_t, m_{t-1}, \underline{i}_t, c_t) \\
& \stackrel{(b)}{=} \Pr(x_t, m_{t-1} | \underline{i}_t) \Pr(y_t | x_t, m_{t-1}, c_t) \\
&= \underline{b}_t(x_t, m_{t-1}) \Pr(y_t | x_t, m_{t-1}, c_t),
\end{aligned}$$

where equality (b) follows from (2.1)–(2.4). Combining (2.17) and (2.18) we have

$$(2.19) \quad b_t(x_t, m_{t-1}) = \frac{\underline{b}_t(x_t, m_{t-1}) \Pr(y_t | x_t, m_{t-1}, c_t)}{\sum_{(x'_t, m'_{t-1}) \in \mathcal{X} \times \mathcal{M}} \underline{b}_t(x'_t, m'_{t-1}) \Pr(y_t | x'_t, m'_{t-1}, c_t)}.$$

Hence,

$$(2.20) \quad b_t = F(\underline{b}_t, y_t, c_t),$$

where F is given by (2.19).

3. Consider a component of \bar{b}_t ,

$$\begin{aligned}
(2.21) \quad \bar{b}_t(x_t, m_t) &= \Pr(x_t, m_t \mid \bar{i}_t) = \Pr(x_t, m_t \mid i_t, l_t) \\
&= \sum_{m_{t-1} \in \mathcal{M}} \Pr(x_t, m_t, m_{t-1} \mid i_t, l_t) \\
&= \sum_{m_{t-1} \in \mathcal{M}} \Pr(x_t, m_{t-1} \mid i_t, l_t) \Pr(m_t \mid x_t, m_{t-1}, i_t, l_t) \\
&\stackrel{(c)}{=} \sum_{m_{t-1} \in \mathcal{M}} \Pr(x_t, m_{t-1} \mid i_t) \Pr(m_t \mid x_t, m_{t-1}, l_t) \\
&= \sum_{m_{t-1} \in \mathcal{M}} b_t(x_t, m_{t-1}) \mathbb{1}[m_t = l_t(x_t, m_{t-1})]
\end{aligned}$$

where equality (c) follows from (2.1) and (2.3), and $\mathbb{1}[\cdot]$ is the indicator function. Hence,

$$(2.22) \quad \bar{b}_t = \bar{F}(b_t, l_t)$$

where \bar{F} is given by (2.21).

□

The above relationships between the controller's beliefs lead to the structural results of the optimal controllers.

THEOREM 2.4. *Consider Problem 2.1 for any arbitrary (but fixed) encoding and memory update strategies $C := (c_1, \dots, c_T)$ and $L := (l_1, \dots, l_T)$, respectively. Then, without loss of optimality, we can restrict attention to control laws of the form*

$$(2.23) \quad U_t = g_t(\bar{B}_t).$$

Proof. We will show that the process $\{\bar{B}_t, t = 1, \dots, T\}$ is a perfectly observed controlled Markov process with control action U_t .

The controller knows \bar{I}_t and hence \bar{B}_t is perfectly observed at the controller's site. Parts (i)–(iii) of Lemma 2.3 can be combined to obtain

$$\begin{aligned}
(2.24) \quad \bar{B}_t &= \bar{F}(F(\underline{F}(\bar{B}_{t-1}, U_{t-1}), Y_t, c_t), l_t) \\
&=: \hat{F}(\bar{B}_{t-1}, Y_t, U_{t-1}, c_t, l_t).
\end{aligned}$$

Let \bar{b} belong to $\mathcal{P}^{\mathcal{X} \times \mathcal{M}}$. For any realization \bar{b}^{t-1} of \bar{B}^{t-1} and u^{t-1} of U^{t-1} , consider

$$\begin{aligned}
(2.25) \quad \Pr(\bar{B}_t = \bar{b}_t \mid \bar{B}^{t-1} = \bar{b}^{t-1}, U^{t-1} = u^{t-1}; C, L, G) \\
&= \sum_{y_t \in \mathcal{Y}} \Pr(\bar{B}_t = b_t, Y_t = y_t \mid \bar{B}^{t-1} = \bar{b}^{t-1}, U^{t-1} = u^{t-1}; C, L, G) \\
&= \sum_{y_t \in \mathcal{Y}} \Pr(\bar{B}_t = b_t \mid Y_t = y_t, \bar{B}^{t-1} = \bar{b}^{t-1}, U^{t-1} = u^{t-1}; C, L, G) \\
&\quad \times \Pr(Y_t = y_t \mid \bar{B}^{t-1} = \bar{b}^{t-1}, U^{t-1} = u^{t-1}; C, L, G) \\
&= \sum_{y_t \in \mathcal{Y}} \mathbb{1} \left[\bar{b}_t = \hat{F}(\bar{b}_{t-1}, y_t, u_{t-1}, c_t, l_t) \right] \\
&\quad \times \Pr(Y_t = y_t \mid \bar{B}^{t-1} = \bar{b}^{t-1}, U^{t-1} = u^{t-1}; C, L, G).
\end{aligned}$$

Now consider

$$\begin{aligned}
(2.26) \quad & \Pr(Y_t = y_t \mid \bar{B}^{t-1} = \bar{b}^{t-1}, U^{t-1} = u^{t-1}; C, L, G) \\
&= \sum_{(x_t, x_{t-1}, z_t, m_{t-1}) \in \mathcal{X} \times \mathcal{X} \times \mathcal{Z} \times \mathcal{M}} \Pr(X_t = x_t, X_{t-1} = x_{t-1}, Z_t = z_t, Y_t = y_t, M_{t-1} = m_{t-1} \mid \\
&\quad \bar{B}^{t-1} = \bar{b}^{t-1}, U^{t-1} = u^{t-1}; C, L, G) \\
&= \sum_{(x_t, x_{t-1}, z_t, m_{t-1}) \in \mathcal{X} \times \mathcal{X} \times \mathcal{Z} \times \mathcal{M}} \Pr(X_{t-1} = x_{t-1}, M_{t-1} = m_{t-1} \mid \bar{B}^{t-1} = \bar{b}^{t-1}, U^{t-1} = u^{t-1}; C, L, G) \\
&\quad \times \Pr(X_t = x_t \mid X_{t-1} = x_{t-1}, M_{t-1} = m_{t-1}, \bar{B}^{t-1} = \bar{b}^{t-1}, U^{t-1} = u^{t-1}; C, L, G) \\
&\quad \times \Pr(Z_t = z_t \mid X_t = x_t, X_{t-1} = x_{t-1}, M_{t-1} = m_{t-1}, \bar{B}^{t-1} = \bar{b}^{t-1}, U^{t-1} = u^{t-1}; C, L, G) \\
&\quad \times \Pr(Y_t = y_t \mid Z_t = z_t, X_t = x_t, X_{t-1} = x_{t-1}, M_{t-1} = m_{t-1}, \bar{B}^{t-1} = \bar{b}^{t-1}, U^{t-1} = u^{t-1}; C, L, G) \\
&= \sum_{(x_t, x_{t-1}, z_t, m_{t-1}) \in \mathcal{X} \times \mathcal{X} \times \mathcal{Z} \times \mathcal{M}} \bar{b}_{t-1}(x_{t-1}, m_{t-1}) \Pr(X_t = x_t \mid X_{t-1} = x_{t-1}, U_{t-1} = u_{t-1}) \\
&\quad \times \mathbf{1}[z_t = c_t(x_t, m_{t-1})] \Pr(Y_t = y_t \mid Z_t = z_t) \\
&= \Pr(Y_t = y_t \mid \bar{B}_{t-1} = \bar{b}_{t-1}, U_{t-1} = u_{t-1}; c_t)
\end{aligned}$$

Substituting the value of (2.26) in (2.25) we get

$$\begin{aligned}
(2.27) \quad & \Pr(\bar{B}_t = \bar{b}_t \mid \bar{B}^{t-1} = \bar{b}^{t-1}, U^{t-1} = u^{t-1}; C, L, G) \\
&= \sum_{y_t \in \mathcal{Y}} \mathbf{1}[\bar{b}_t = \hat{F}(\bar{b}_{t-1}, y_t, u_{t-1}, c_t, l_t)] \\
&\quad \times \Pr(Y_t = y_t \mid \bar{B}_{t-1} = \bar{b}_{t-1}, U_{t-1} = u_{t-1}; c_t) \\
&= \Pr(\bar{B}_t = \bar{b}_t \mid \bar{B}_{t-1} = \bar{b}_{t-1}, U_{t-1} = u_{t-1}; c_t, l_t).
\end{aligned}$$

Thus for any fixed C and L , \bar{B}_t is a controlled Markov process with control action U_t . Further, the expected instantaneous cost can be written as

$$(2.28) \quad \mathbb{E}\{\rho(X_t, U_t) \mid \dot{i}_{t+1}\} = \sum_{x_t \in \mathcal{X}} \rho(x_t, u_t) \Pr(X_t = x_t \mid \dot{i}_{t+1}).$$

Now, by Bayes rule

$$(2.29) \quad \Pr(x_t \mid \dot{i}_{t+1}) = \Pr(x_t \mid \bar{i}_t, u_t, g_t) = \frac{\Pr(x_t, u_t \mid \bar{i}_t, g_t)}{\sum_{x'_t \in \mathcal{X}} \Pr(x'_t, u_t \mid \bar{i}_t, g_t)}.$$

Further,

$$\begin{aligned}
(2.30) \quad & \Pr(x_t, u_t \mid \bar{i}_t, g_t) = \Pr(u_t \mid x_t, \bar{i}_t, g_t) \Pr(x_t \mid \bar{i}_t, g_t) \\
&\stackrel{(d)}{=} \Pr(u_t \mid \bar{i}_t, g_t) \Pr(x_t \mid \bar{i}_t),
\end{aligned}$$

where equality (d) follows from (2.1) and (2.5). Combine (2.29) and (2.30), and cancel $\Pr(u_t \mid \bar{i}_t, g_t)$ from the numerator and the denominator to obtain

$$(2.31) \quad \Pr(X_t = x_t \mid \dot{i}_{t+1}) = \Pr(X_t = x_t \mid \bar{i}_t) = \sum_{m_t \in \mathcal{M}} \bar{b}_t(x_t, m_t).$$

Substituting back in (2.28) gives

$$\begin{aligned}
(2.32) \quad & \mathbb{E}\{\rho(X_t, U_t) \mid \dot{i}_{t+1}\} = \sum_{x_t, m_t \in \mathcal{X} \times \mathcal{M}} \rho(x_t, u_t) \bar{b}_t(x_t, m_t) \\
&=: \hat{\rho}(\bar{b}_t, u_t).
\end{aligned}$$

We can think of $\hat{\rho}(\cdot)$ as the instantaneous cost and write the total expected cost as

$$(2.33) \quad \mathbb{E} \left\{ \sum_{t=1}^T \rho(X_t, U_t) \mid C, L, G \right\} = \mathbb{E} \left\{ \sum_{t=1}^T \mathbb{E} \{ \rho(X_t, U_t) \mid I_{t+1} \} \mid C, L, G \right\} \\ = \mathbb{E} \left\{ \sum_{t=1}^T \hat{\rho}(\bar{B}_t, U_t) \mid C, L, G \right\}.$$

Hence the process $\{\bar{B}_t, t = 1, \dots, T\}$ is a perfectly observed controlled Markov process with control action U_t . The instantaneous cost $\hat{\rho}(\cdot)$ is a function of the controlled state \bar{B}_t and the control action U_t . From Markov decision theory [14] we know that there is no loss of optimality in restricting attention to control laws of the form (2.23). \square

2.3.1. Implication of the structural results. Theorem 2.4 implies that at each stage t , without loss of optimality, we can restrict attention to controllers belonging to the family \mathcal{G} of functions from $\mathcal{P}^{\mathcal{X} \times \mathcal{M}}$ to \mathcal{U} . With this modification Problem 2.1 is equivalent to the following problem:

PROBLEM 2.2. *Given a perfect observation system $(\mathcal{X}, \mathcal{W}, \mathcal{M}, \mathcal{Z}, \mathcal{N}, \mathcal{Y}, \mathcal{U}, P_{X_1}, P_W, P_N, f, h, \rho, T)$, choose a design (C^*, L^*, G^*) that is optimal with respect to the performance criterion of (2.6), i.e.,*

$$(2.34) \quad \mathcal{J}_T(C^*, L^*, G^*) = \mathcal{J}_T^* := \min_{C, L, G \in \mathcal{C}^T \times \mathcal{L}^T \times \mathcal{G}^T} \mathcal{J}_T(C, L, G),$$

where $\mathcal{G}^T := \mathcal{G} \times \dots \times \mathcal{G}$ (T times).

Using the structural results of Theorem 2.4, we can transform Problem 2.1 into an equivalent problem, Problem 2.2, in which the domain of all the decision rules, the encoding rules, the memory update rules, and the control rules, is not changing with time. This is in contrast to Problem 2.1 where the domain of the control rules was increasing with time. This reduction to a time-invariant domain is necessary for extending the solution methodology for the finite horizon problems to infinite horizon.

In the next section we provide a sequential decomposition of Problem 2.2.

2.4. Global optimization. As explained in Section 2.2, Problems 2.1 and 2.2 are dynamic teams with strictly non-classical information structure. To obtain a sequential decomposition we need to identify information states sufficient for performance evaluation, or equivalently, find sufficient statistics for performance evaluation. The sequential nature of the problem suggests choosing an information state for each decision rule. Suppose $\underline{\pi}_t$, π_t , and $\bar{\pi}_t$ are information states at time t for the encoder, memory update, and the controller respectively. Due to the decentralization of information, these information states should depend only on the decision rules (which are common knowledge) and not on the observation of any agent. For $\underline{\pi}_t$, π_t , and $\bar{\pi}_t$ to be information states in the sense of [14], at each instant of time π_t must be determined from $\underline{\pi}_t$ and c_t ; $\bar{\pi}_t$ must be determined from π_t and l_t ; and $\underline{\pi}_{t+1}$ must be determined from $\bar{\pi}_t$ and g_t . However, a system can have more than one information state, and not all of them are sufficient for performance evaluation (see [46]). To be sufficient for performance evaluation, the information states must *absorb/summarize* the effect

of past decision rules on the expected future cost⁴, that is they should satisfy

$$\begin{aligned}
 (2.35) \quad \mathbb{E} \left\{ \sum_{s=t}^T \rho(X_s, U_s) \mid C, L, G \right\} &= \mathbb{E} \left\{ \sum_{s=t}^T \rho(X_s, U_s) \mid \underline{\pi}_t, c_t^T, l_t^T, g_t^T \right\} \\
 &= \mathbb{E} \left\{ \sum_{s=1}^T \rho(X_s, U_s) \mid \pi_t, c_{t+1}^T, l_t^T, g_t^T \right\} \\
 &= \mathbb{E} \left\{ \sum_{s=t}^T \rho(X_s, U_s) \mid \bar{\pi}_t, c_{t+1}^T, l_{t+1}^T, g_t^T \right\}.
 \end{aligned}$$

or equivalently,

$$(2.36) \quad \mathbb{E} \{ \rho(X_t, U_t) \mid C, L, G \} = \mathbb{E} \{ \rho(X_t, U_t) \mid \bar{\pi}_t, g_t \}$$

These properties that information states sufficient for performance evaluation must satisfy are explained in more detail in [16].

For sequential problems, one way to obtain information states satisfying the above properties is by converting the model to Witsenhausen's standard form [44]. However in the standard form the space in which information states belong increases with time, so such a transformation to the standard form does not lead to a formulation that can be extended to infinite horizon problems. We want an information state that will be appropriate for both finite and infinite horizon problems. This is possible only when the space in which the information state belongs is time-invariant.

Thus information states sufficient for performance evaluation should satisfy the following properties:

- (P1) *They must be states*, that is, at each instant of time π_t should be a function of $\underline{\pi}_t$ and c_t ; $\bar{\pi}_t$ should be a function of π_t and l_t ; and $\underline{\pi}_{t+1}$ should be a function of $\bar{\pi}_t$ and g_t .
- (P2) *They must be sufficient for performance evaluation*, that is, they should satisfy (2.35) or (2.36).
- (P3) They should take values in a time-invariant space.

Next we present information states that have the above properties and show how these information states lead to a sequential decomposition of Problem 2.2. *We want to reemphasize that the hardest part in our solution methodology is to identify the appropriate information states; there are no known solution methodologies in identifying information states for decentralized stochastic control problems like Problem 2.2.*

The information states defined below have all the above-discussed desired features.

DEFINITION 2.5. *Let Π be the space of probability measure on $\mathcal{X} \times \mathcal{M} \times \mathcal{P}^{\mathcal{X} \times \mathcal{M}}$. Define $\underline{\pi}_t, \pi_t, \bar{\pi}_t, t = 1, \dots, T$, as follows:*

1. $\underline{\pi}_t := \Pr(X_t, M_{t-1}, \underline{B}_t)$.
2. $\pi_t := \Pr(X_t, M_{t-1}, B_t)$.
3. $\bar{\pi}_t := \Pr(X_t, M_t, \bar{B}_t)$.

Here $\underline{\pi}_t, \pi_t$, and $\bar{\pi}_t$ are probability measures (or probability laws) on the probability space $(\mathcal{X} \times \mathcal{M} \times \mathcal{P}^{\mathcal{X} \times \mathcal{M}}, \mathcal{B}(\mathcal{P}^{\mathcal{X} \times \mathcal{M}}))$, where $\mathcal{B}(\mathcal{P}^{\mathcal{X} \times \mathcal{M}})$ is the Borel

⁴Note that in problems with classical information structure, we can find an information state that is independent of the control law [14]. For problems with strictly non-classical information structures it is not always possible find information states that are independent of the control law. However, as long as the expected future cost conditioned on the information state is conditionally independent of the past control laws, a sequential decomposition can be obtained using that information state. See [44] for a proof.

σ -algebra on $\mathcal{P}^{\mathcal{X} \times \mathcal{M}}$. These probability measures are information states sufficient for performance evaluation of Problem 2.2. Specifically, they satisfy the following properties:

LEMMA 2.6. $\underline{\pi}_t, \pi_t, \bar{\pi}_t$ are information states for the encoder, the memory update and the controller respectively, i.e.,

1. there is a linear transformation $\underline{Q}(c_t)$ such that

$$(2.37) \quad \pi_t = \underline{Q}(c_t)\underline{\pi}_t.$$

2. there is a linear transformation $Q(l_t)$ such that

$$(2.38) \quad \bar{\pi}_t = Q(l_t)\pi_t.$$

3. there is a linear transformation $\bar{Q}(g_t)$ such that

$$(2.39) \quad \underline{\pi}_{t+1} = \bar{Q}(g_t)\bar{\pi}_t.$$

4. the conditional expected instantaneous cost can be expressed as

$$(2.40) \quad \mathbb{E}\{\rho(X_t, U_t) \mid c^t, l^t, g^t\} = \tilde{\rho}(\bar{\pi}_t, g_t),$$

where $\tilde{\rho}$ is a deterministic function.

Proof.

1. Consider a component of π_t ,

$$(2.41) \quad \begin{aligned} \pi_t(x, m, b) &= \int_{\underline{A}(b, c_t)} \underline{\pi}_t(x, m, \underline{b}) d\underline{b} \\ &=: \underline{Q}_t(c_t)\underline{\pi}_t, \end{aligned}$$

where $\underline{A}(b, c) = \{\underline{b} \in \mathcal{P}^{\mathcal{X} \times \mathcal{M}} : b = F(\underline{b}, c)\}$.

2. Consider a component of $\bar{\pi}_t$,

$$(2.42) \quad \begin{aligned} \bar{\pi}_t(x, m, \bar{b}) &= \sum_{\{m' \in \mathcal{M} : m' = l_t(x, m)\}} \int_{A(\bar{b}, l_t)} \pi_t(x, m', b) db \\ &=: Q(l_t)\pi_t, \end{aligned}$$

where $A(\bar{b}, l) = \{b \in \mathcal{P}^{\mathcal{X} \times \mathcal{M}} : \bar{b} = \bar{F}(b, l)\}$.

3. Consider a component of $\underline{\pi}_{t+1}$,

$$(2.43) \quad \begin{aligned} \underline{\pi}_{t+1}(x, m, \underline{b}) &= \sum_{x_t \in \mathcal{X}} \int_{\bar{A}(\underline{b}, g_t)} \bar{\pi}_t(x_t, m, \bar{b}) \\ &\quad \times \Pr(X_{t+1} = x \mid X_t = x_t, U_t = g_t(\bar{b})) d\bar{b} \\ &=: \bar{Q}(g_t)\bar{\pi}_t, \end{aligned}$$

where $\bar{A}(\underline{b}, g) = \{\bar{b} \in \mathcal{P}^{\mathcal{X} \times \mathcal{M}} : \underline{b} = \underline{F}(\bar{b}, g(\bar{b}))\}$.

4. Consider $\mathbb{E}\{\rho(X_t, U_t) \mid c^t, l^t, g^t\}$. By the problem formulation $\underline{\pi}_1$ is known to all agents. For specified c^t, l^t and g^{t-1} , the information state $\bar{\pi}_t$ can be evaluated using the transformations of previous steps of this Lemma. Thus,

$$(2.44) \quad \begin{aligned} \mathbb{E}\{\rho(X_t, U_t) \mid c^t, l^t, g^t\} &= \mathbb{E}\{\rho(X_t, U_t) \mid c^t, l^t, g^t, \bar{\pi}_t\} \\ &= \sum_{x_t \in \mathcal{X}} \int_{\mathcal{P}^{\mathcal{X} \times \mathcal{M}}} \bar{\pi}_t(x_t, \bar{b}_t) \rho(x_t, g_t(\bar{b}_t)) d\bar{b}_t := \tilde{\rho}(\bar{\pi}_t, g_t), \end{aligned}$$

where $\bar{\pi}_t(x_t, \bar{b}_t)$ is the marginal of $\bar{\pi}_t(x_t, m_t, \bar{b}_t)$.

□

Points 1, 2, and 3 of Lemma 2.6 shows that the information states $\underline{\pi}_t$, π_t , and $\bar{\pi}_t$ satisfy property (P1); point 4 shows that these information states satisfy property (P2). Property (P3) is satisfied by definition. Thus, $\underline{\pi}_t$, π_t , and $\bar{\pi}_t$ are information states sufficient for performance evaluation. In order to obtain a sequential decomposition, first reconsider the performance criterion of (2.6), which can be rewritten as

$$(2.45) \quad \mathbb{E} \left\{ \sum_{t=1}^T \rho(X_t, U_t) \mid C, L, G \right\} = \sum_{t=1}^T \mathbb{E} \{ \rho(X_t, U_t) \mid c^t, l^t, g^t \} \\ =: \sum_{t=1}^T \tilde{\rho}(\bar{\pi}_t, g_t),$$

where the sequence $\{\bar{\pi}_1, \dots, \bar{\pi}_T\}$ depends on the choice of (C, L, G) . Hence, Problem 2.2 is equivalent to the following deterministic problem:

PROBLEM 2.3. *Consider a deterministic system with states $\underline{\pi}_t, \pi_t, \bar{\pi}_t$. The initial state $\underline{\pi}_1$ is known and for $t \geq 1$, the system evolves as follows:*

$$(2.46) \quad \pi_t = \underline{Q}(c_t)\underline{\pi}_t,$$

$$(2.47) \quad \bar{\pi}_t = Q(l_t)\pi_t,$$

$$(2.48) \quad \underline{\pi}_{t+1} = \bar{Q}(g_t)\bar{\pi}_t,$$

where c_t, l_t, g_t belong to $\mathcal{C}, \mathcal{L}, \hat{\mathcal{G}}$ respectively, and $\underline{Q}, Q, \bar{Q}$ are known linear transformations given by Lemma 2.6. At time t , an instantaneous cost $\tilde{\rho}(\bar{\pi}_t, g_t)$ is incurred.

The optimization problem is to determine design (C, L, G) , where $C := (c_1, \dots, c_T)$, $L := (l_1, \dots, l_T)$, and $G := (g_1, \dots, g_T)$, to minimize the total cost over horizon T , i.e.,

$$(2.49) \quad \min_{(C, L, G) \in \mathcal{C}^T \times \mathcal{L}^T \times \hat{\mathcal{G}}^T} \sum_{t=1}^T \tilde{\rho}(\bar{\pi}_t, g_t).$$

This is a classical deterministic optimal control problem in function space; optimal functions (C^*, L^*, G^*) can be determined as follows:

THEOREM 2.7. *An optimal design (C^*, L^*, G^*) for Problem 2.3 (and consequently for Problem 2.2 and thereby for Problem 2.1) is given the following nested optimality equations:*

$$(2.50) \quad \bar{V}_T(\bar{\pi}) = \inf_{g_T \in \hat{\mathcal{G}}} \tilde{\rho}(\bar{\pi}, g_T),$$

and for $t = 1, \dots, T$

$$(2.51) \quad \underline{V}_t(\underline{\pi}) = \min_{c_t \in \mathcal{C}} V_t(\underline{Q}(c_t)\underline{\pi}),$$

$$(2.52) \quad V_t(\pi) = \min_{l_t \in \mathcal{L}} \bar{V}_t(Q(l_t)\pi),$$

$$(2.53) \quad \bar{V}_t(\bar{\pi}) = \inf_{g_t \in \hat{\mathcal{G}}} \{ \tilde{\rho}(\bar{\pi}, g_t) + \underline{V}_{t+1}(\bar{Q}(g_t)\bar{\pi}) \}.$$

The arg min (or arg inf) at each step determines the corresponding optimal design for that stage. Furthermore, the optimal performance is given by

$$(2.54) \quad \mathcal{J}_T^* = \underline{V}_1(\underline{\pi}_1).$$

Proof. This is a standard result, see [14, Chapter 2]. \square

2.5. Discussion of problem 2.3. We present an alternative look at Problem 2.3 which will be useful when we study the infinite horizon version of Problem 2.1. As pointed out in Section 2.4, Problem 2.3 is a deterministic control problem with state space Π and action space alternating between \mathcal{C} , \mathcal{L} and \mathcal{G} . We now introduce a sequence of *meta-functions* $\underline{\Delta}_t$, Δ_t , and $\overline{\Delta}_t$, $t = 1, \dots, T$, where $\underline{\Delta}_t$ is a function from Π to \mathcal{C} , Δ_t is a function from Π to \mathcal{L} , and $\overline{\Delta}_t$ is a function from Π to \mathcal{G} . These meta-functions describe the rationale used to select the “action” (i.e. the design c^t, l^t, g^t) at time t . The choice of all meta-functions for horizon T is called *meta-design*. Problem 2.3 is equivalent to the following feedback control problem.

PROBLEM 2.4. Consider a deterministic system with states $\underline{\pi}_t, \pi_t, \overline{\pi}_t \in \Pi$, and “control actions” $c_t \in \mathcal{C}$, $l_t \in \mathcal{L}$, and $g_t \in \mathcal{G}$. The initial state $\underline{\pi}_1$ is known and for $t \geq 1$, the system evolves as follows:

$$(2.55) \quad \pi_t = \underline{Q}(c_t)\underline{\pi}_t, \quad \overline{\pi}_t = Q(l_t)\pi_t, \quad \text{and} \quad \underline{\pi}_{t+1} = \overline{Q}(g_t)\overline{\pi}_t,$$

where \underline{Q} , Q , and \overline{Q} are known transformations given by Lemma 2.6. The “control actions” c_t , l_t , and g_t are chosen according to the meta-functions $\underline{\Delta}_t$, Δ_t , and $\overline{\Delta}_t$ as follows

$$(2.56) \quad c_t = \underline{\Delta}_t(\underline{\pi}_t), \quad l_t = \Delta_t(\pi_t), \quad \text{and} \quad g_t = \overline{\Delta}_t(\overline{\pi}_t).$$

At each time, an instantaneous cost $\tilde{\rho}(\overline{\pi}_t, g_t)$ is incurred. The optimization problem is to determine the meta-design $\tilde{\Delta}^T := (\underline{\Delta}_1, \Delta_1, \overline{\Delta}_1, \dots, \underline{\Delta}_T, \Delta_T, \overline{\Delta}_T)$ to minimize the total cost over horizon T , i.e.,

$$(2.57) \quad \min \sum_{t=1}^T \tilde{\rho}(\underline{\pi}_t, g_t)$$

where the minimization is over the choice of $\tilde{\Delta}^T$

The nested optimality equations of Theorem 2.7 determine the globally optimal meta-functions $\underline{\Delta}_t, \Delta_t, \overline{\Delta}_t$, for $t = 1, \dots, T$, i.e., the optimal feedback laws for Problem 2.4. Since Problem 2.4 is a deterministic control problem with a known initial state we only need to specify the control “actions” c_t, l_t, g_t , for $t = 1, \dots, T$. This is why we have considered Problem 2.3 instead of Problem 2.4. Nevertheless, Problem 2.4 will be useful in clarifying the nature of the solution of the infinite horizon problem corresponding to Problem 2.1.

3. Explanation of the Solution Methodology. The sequential decomposition obtained above can be interpreted as follows. Suppose that before the system is started the sensor and the controller get together to determine an optimal design that they will use. Instead of testing the performance of each design one by one, they decide to choose the design sequentially. So, they need to agree on a mechanism (or an algorithm) that will, at each time instant and for any choice of past design rules,⁵ determine the future design rules optimally. To do so, for any choice of past design rules, the sensor and the controller must be able to consistently evaluate the optimal future performance. To be consistent in their evaluation, each agent must

⁵In this description, we use design rule to refer to either the encoding rule, the memory update rule, or the control law.

“know” what the other agent is “thinking”. Suppose the design rules until time t , denoted by $\gamma^t := (c_1, c_2, \dots, c_t, l_1, l_2, \dots, l_{t-1}, g_1, g_2, \dots, g_{t-1})$ have been agreed upon (by some mechanism) and the sensor and the controller want to determine the next design rule l_t . If they allow the system to run until time t , the sensor will know the values of X_t and M_{t-1} while the controller will know the values of Y^t and U^{t-1} . However, they do not know the other agent’s observations. They can form a belief on the other agent’s observations, but then they do not know the other agent’s belief on their observations. If they form a belief on the other agent’s belief on their observation, they will not know this belief on the belief. This process of forming a belief on what the other agent is “thinking” will continue until the sensor and the controller agree upon what they are thinking. In [3] Aumann showed that such an agreement will occur in the “common knowledge” between the two agents. Formally, suppose (Ω, \mathcal{F}, P) is the probability space of the primitive random variables of the system. For any fixed γ^t , (X_t, M_{t-1}) and (Y^t, U^{t-1}) are random vectors on $(\mathcal{X} \times \mathcal{M}, 2^{\mathcal{X} \times \mathcal{M}})$ and $(\mathcal{Y}^t \times \mathcal{U}^{t-1}, 2^{\mathcal{Y}^t \times \mathcal{U}^{t-1}})$, respectively. Let $\sigma(X_t, M_{t-1})$ and $\sigma(Y^t, U^{t-1})$ denote the smallest subfields of \mathcal{F} with respect to which (X_t, M_{t-1}) and (Y^t, U^{t-1}) are, respectively, measurable. Then the common knowledge between (X_t, M_{t-1}) and (Y^t, U^{t-1}) is $\sigma(X_t, M_{t-1}) \cap \sigma(Y^t, U^{t-1}) =: K_t(\gamma^t)$. Thus, to do a sequential decomposition, the agents should decide what to do for all $K_t(\gamma^t)$ obtained by varying γ^t over all possible values. However, it is difficult to identify the space of all possible realizations of $K_t(\gamma^t)$. So instead of using $K_t(\gamma^t)$ as an (information) state, the agents can use $\sigma(X_t, M_{t-1}, B_t) =: \hat{K}_t(\gamma^t)$, which is a super-field of $K_t(\gamma^t)$ (see Appendix A for proof). $\hat{K}_t(\gamma^t)$ also depends on γ^t and it is difficult to evaluate the space of realization of $\hat{K}_t(\gamma^t)$ obtained by varying γ^t over all possible values. However, if we go to the image space of the random variables, we can obtain an “over-approximation” of $\hat{K}_t(\gamma^t)$. Consider the image space of the random vectors $(X_t, M_{t-1}, B_t) : (\Omega, \mathcal{F}, P) \rightarrow (\mathcal{X} \times \mathcal{M} \times \mathcal{B}(\mathbb{R}^2), \mathcal{B}(2^{\mathcal{X} \times \mathcal{M}} \times \mathcal{B}(\mathbb{R}^2)), \hat{P}_t(\gamma^t)) =: \Lambda_t(\gamma^t)$, where $\mathcal{B}(\cdot)$ denotes the Borel set. In $\Lambda_t(\gamma^t)$ only the measure $\hat{P}_t(\gamma^t)$ depends on the choice of past design rules. Although it is difficult to evaluate all reachable realizations of $\hat{P}_t(\gamma^t)$ obtained by varying γ^t over all possible values, the space of all realizations of $\hat{P}_t(\gamma^t)$ is known and is equal to all probability measures on $(\mathcal{X} \times \mathcal{M} \times \mathcal{B}(\mathbb{R}^2), \mathcal{B}(2^{\mathcal{X} \times \mathcal{M}} \times \mathcal{B}(\mathbb{R}^2)))$. So the sensor and the controller can decide on what action to take for each probability measure \hat{P}_t , that is for any probability space $\Lambda_t := (\mathcal{X} \times \mathcal{M} \times \mathcal{B}(\mathbb{R}^2), \mathcal{B}(2^{\mathcal{X} \times \mathcal{M}} \times \mathcal{B}(\mathbb{R}^2)), \hat{P}_t)$, and not worry whether the space is reachable or not. Notice that the information state π_t is equivalent to Λ_t defined here. In the definition of π_t the sample space and the σ -algebra are implicitly specified. Similar interpretations hold for $\underline{\pi}_t$ and $\bar{\pi}_t$.

If the rules for breaking ties are made common knowledge, the nested optimality equations of Theorem 2.7 allow the sensor and the controller (or anyone who knows the model and rules for breaking ties, henceforth referred to as the *designer*) to sequentially and consistently determine optimal design rules in two stages. In the first stage, for each time instant and for each realization of the information state determine an optimal design rule to be used if that information state is actually realized. In the second stage, sequentially determine for every t optimal design rules c_t, l_t, g_t to be implemented as follows. For the first time instant using the information state $\underline{\pi}_1$, which is part of the model, the sensor and the controller (and the designer) can determine an optimal c_1^* . This choice of c_1^* is common knowledge between the sensor and the controller since the model and the rule for breaking ties are common knowledge. For these values of $\underline{\pi}_1$ and c_1^* , Lemma 2.6 gives the value of the realization of π_1 . This value is common knowledge between the sensor and the controller (and the designer).

Now, using the result of the first stage, an optimal l_1^* can be determined which in turn gives the realization of $\bar{\pi}_1$. This realization is common knowledge between the sensor and controller (and the designer). This processes can be continued until all the design rules $c_1^*, l_1^*, g_1^*, \dots, c_T^*, l_T^*, g_T^*$ are determined. This design is optimal and common knowledge between the sensor and the controller (and the designer).

In view of the discussion in Section 2.5 the first stage corresponds to determining optimal meta-functions $\underline{\Delta}_t$, Δ_t , and $\bar{\Delta}_t$, $t = 1, \dots, T$, while the second stage corresponds to determining optimal design rules c_t, l_t, g_t , $t = 1, \dots, T$ that are implemented. These design rules correspond to the control actions in Problem 2.4; since the problem is deterministic, they can be specified before the system starts running.

The nested optimality equations of Theorem 2.7 are functional optimization problems: for each realization of the information state we need to determine an optimal design rule (a function) to be used if that state is actually realized. Contrast this with the centralized stochastic optimization problems where the dynamic programming equations result in parameter optimization problems: for each realization of the information state we need to determine an optimal control action (a parameter) to be taken if that state is actually realized. Functional optimization problems are an order of magnitude harder to solve than parameter optimization problems. The cardinality of the function space (e.g. \mathcal{C}) increases exponentially with a linear increase in the cardinality of the “action” space (\mathcal{Z}). Moreover the function space (e.g. \mathcal{G}) can be uncountable even when the action space (\mathcal{U}) is finite because the size of the function space is determined by both its domain and range. This increase in complexity makes decentralized stochastic control problems harder to solve than centralized stochastic optimization problems.

We believe that this is a fundamental feature of decentralized optimization problems and not something specific to our solution. Since an agent does not know other agent’s observations, in order to consistently interpret the other agent’s action, it should know the design rule of the other agent. So, in any sequential decomposition, at each instant of time an agent needs to determine its design rule and not just its control action. So any sequential decomposition will result in functional optimization problems.

Our solution is “simpler” than the only other known methodology for sequentially solving dynamic teams—Witsenhausen’s Standard Form [44]. In [44] Witsenhausen showed how to convert *any* sequential optimization problem to “standard form” and showed how to obtain a sequential decomposition for the standard form. If Problem 2.1 is converted into standard form, the information state at time t will be $\sigma(X_t, M_{t-1}, Y^t, U^{t-1}) =: \tilde{K}_t(\gamma^t)$. Observe that $\hat{K}_t(\gamma^t) \subset \tilde{K}_t(\gamma^t)$. So, our information state is a subfield of the information state in the standard form, and a sufficient statistic for the decomposition presented in standard form. However, the image space $\Lambda_t(\gamma^t)$ in our decomposition is bigger than the corresponding image space in the standard form. But the image space in the standard form increases with time, so the standard form can not be used to solve the infinite horizon problem. The image space in our decomposition does not change with time, which enables us to tackle infinite horizon problems as shown in the next section.

4. The Infinite Horizon Problem. In this section we extend the model of Section 2.1 to an infinite horizon ($T \rightarrow \infty$) using two performance criteria: the expected discounted cost and the average cost per unit time. Let (C, L, G) , $C := (c_1, c_2, \dots)$, $L := (l_1, l_2, \dots)$, $G := (g_1, g_2, \dots)$ denote an infinite horizon policy. The two performance criteria that we consider are:

1. THE EXPECTED DISCOUNTED COST where the performance of a design is determined by

$$(4.1) \quad \mathcal{J}^\beta(C, L, G) = \mathbb{E} \left\{ \sum_{t=1}^{\infty} \beta^{t-1} \rho(X_t, U_t) \mid C, L, G \right\},$$

where $0 < \beta < 1$ is called the discount factor.

2. THE AVERAGE COST PER UNIT TIME where the performance of a design is determined by

$$(4.2) \quad \bar{\mathcal{J}}(C, L, G) = \limsup_{T \rightarrow \infty} \frac{1}{T} \mathbb{E} \left\{ \sum_{t=1}^T \rho(X_t, U_t) \mid C, L, G \right\}.$$

We take the limsup rather than lim as for some designs (C, L, G) the limit may not exist.

Ideally, while implementing a solution for infinite horizon problems, we would like to use time-invariant designs. This motivates the following definitions.

DEFINITION 4.1. A design (C, L, G) , $C := (c_1, c_2, \dots)$, $L := (l_1, l_2, \dots)$, $G := (g_1, g_2, \dots)$ is called stationary (or time-invariant) if $c_1 = c_2 = \dots = c$, $l_1 = l_2 = \dots = l$, $g_1 = g_2 = \dots = g$.

DEFINITION 4.2. Let $\tilde{\Delta}_t := (\underline{\Delta}_t, \Delta_t, \bar{\Delta}_t)$. A meta-design $\tilde{\Delta}^\infty := (\tilde{\Delta}_1, \tilde{\Delta}_2, \dots)$ is called stationary (or time-invariant) if $\tilde{\Delta}_1 = \tilde{\Delta}_2 = \dots = \tilde{\Delta}$.

In centralized stochastic control problems with time-homogenous evolution and time-homogenous cost function, one can restrict attention to stationary designs without any loss of optimality. This greatly simplifies the search for an optimal design. It is natural to wonder if such a result also holds for dynamic teams. It is not known whether, in general, stationary designs are optimal for dynamic teams or not. In this section we show that for the problem under consideration, stationary designs may not be optimal. However, there exist stationary meta-designs that are optimal: for the discounted cost problem one can restrict attention to stationary meta-designs without any loss of optimality; for the average cost per unit time problem, under a technical condition, one can restrict attention to stationary meta-designs. The optimal design corresponding to an optimal stationary meta-design is, in general, time-varying.

4.1. The expected discounted cost problem. Consider the infinite horizon problem with expected discounted cost criterion given by (4.1). For this problem the relations of Lemma 2.3 hold, hence the structural result of Theorem 2.4 is valid, and we can restrict attention to encoders belonging to \mathcal{G} . Consider $\underline{\pi}_t, \pi_t, \bar{\pi}_t$ as in Definition 2.5. Lemma 2.6 can be proved as before. The transformations $\underline{Q}, Q, \bar{Q}$ and the expected instantaneous cost $\tilde{\rho}$ are the same as in the finite horizon case. Let $\gamma_t := (c_t, l_t, g_t)$ denote the design at time t , and Γ denote the function space $\mathcal{C} \times \mathcal{L} \times \mathcal{G}$. We can combine (2.55) and (2.56) as

$$(4.3) \quad \underline{\pi}_{t+1} = \tilde{Q}(\gamma_t) \underline{\pi}_t, \quad \gamma_t = \tilde{\Delta}_t(\underline{\pi}_t)$$

where $\tilde{Q}(\gamma_t) := \bar{Q}(g_t) \circ Q(l_t) \circ \underline{Q}(c_t)$ and $\tilde{\Delta}_t(\underline{\pi}_t) = (\underline{\Delta}(\underline{\pi}_t), \Delta_t(\pi_t), \bar{\Delta}_t(\bar{\pi}_t))$. The instantaneous cost at time t can be rewritten as

$$(4.4) \quad \tilde{\rho}(\underline{\pi}_t, \gamma_t) := \hat{\rho}((Q(l_t) \circ \underline{Q}(c_t)) \underline{\pi}_t, g_t).$$

Hence, the infinite horizon problem with the expected discounted cost criterion given by (4.1) is equivalent to the following deterministic optimization problem.

PROBLEM 4.1. Consider a deterministic system with state space Π and action space Γ . The system dynamics are given by

$$(4.5) \quad \pi_{t+1} = \tilde{Q}(\gamma_t)\pi_t, \quad \gamma_t = \tilde{\Delta}_t(\pi_t)$$

where \tilde{Q} is a known transformation and $\tilde{\Delta} : \Pi \rightarrow \Gamma$ for all t . At each instant of time an instantaneous cost $\bar{\rho}(\pi_t, \gamma_t)$ is incurred. The objective is to choose meta-design $\tilde{\Delta}^\infty := (\tilde{\Delta}_1, \tilde{\Delta}_2, \dots)$ so as to minimize the infinite horizon cost given by

$$(4.6) \quad \mathcal{J}^\beta(\tilde{\Delta}^\infty) := \sum_{t=1}^{\infty} \beta^{t-1} \bar{\rho}(\pi_t, \gamma_t).$$

Problem 4.1 is a standard infinite horizon discounted cost feedback control problem. Since we have assumed that $0 \leq \rho < K$, where $K < \infty$, which in-turn implies $0 \leq \bar{\rho} < K$, an optimal meta-design is guaranteed to exist and we have the following result:

THEOREM 4.3. For Problem 4.1 and consequently for the infinite horizon expected discounted cost problem with the performance criterion given by (4.1) one can restrict attention to stationary meta-designs without any loss of optimality. Specifically there exists a stationary meta-design $\tilde{\Delta}^{*,\infty} := (\tilde{\Delta}^*, \tilde{\Delta}^*, \dots)$, and a corresponding infinite horizon design (C^*, L^*, G^*) , $C^* := (c_1^*, c_2^*, \dots)$, $L := (l_1^*, l_2^*, \dots)$, $G := (g_1, g_2, \dots)$ such that

$$(4.7) \quad \mathcal{J}^\beta(\tilde{\Delta}^{*,\infty}) = V(\pi_1),$$

where V is the unique uniformly bounded fixed point of

$$(4.8) \quad V(\pi) = \min_{\gamma \in \Gamma} \{ \bar{\rho}(\pi, \gamma) + \beta V(\tilde{Q}(\gamma)(\pi)) \},$$

and $\tilde{\Delta}^*$ satisfies

$$(4.9) \quad V(\pi) = \bar{\rho}(\pi, \tilde{\Delta}^*(\pi)) + \beta V(\tilde{Q}(\tilde{\Delta}^*(\pi))(\pi)).$$

An optimal design (c_t^*, l_t^*, g_t^*) to be implemented at time t is given by

$$(4.10) \quad (c_t^*, l_t^*, g_t^*) =: \gamma_t^* = \tilde{\Delta}^*(\pi_t).$$

Proof. This is a standard result, see [10, Chapter 6]. \square

4.2. The average cost per unit time problem. Consider the infinite horizon problem with average cost per unit time criterion given by (4.2). Using the argument of the first paragraph of Section 4.1, this problem is equivalent to the following deterministic problem.

PROBLEM 4.2. Consider a deterministic system with state space Π and action space Γ . The system dynamics are given by

$$(4.11) \quad \pi_{t+1} = \tilde{Q}(\gamma_t)\pi_t, \quad \gamma_t = \tilde{\Delta}_t(\pi_t)$$

where \tilde{Q} is a known transformation and $\tilde{\Delta}_t : \Pi \rightarrow \Gamma$ for all t . At each instant of time an instantaneous cost $\bar{\rho}(\pi_t, \gamma_t)$ is incurred. The objective is to choose meta-design $\tilde{\Delta}^\infty := (\tilde{\Delta}_1, \tilde{\Delta}_2, \dots)$ so as to minimize the infinite horizon cost given by

$$(4.12) \quad \bar{\mathcal{J}}(\tilde{\Delta}^\infty) := \limsup_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T \bar{\rho}(\pi_t, \gamma_t).$$

For this problem an optimal meta-design may not exist. However, under suitable conditions, we can guarantee the existence of ε -optimal meta-designs. Specifically, we have the following result:

THEOREM 4.4. *For Problem 4.2 and consequently for the infinite horizon average cost per unit time problem with the performance criterion given by (4.2), assume **(A1)** for any $\varepsilon > 0$ there exist bounded measurable functions $v(\cdot)$ and $r(\cdot)$ and meta-function $\tilde{\Delta}^* : \Pi \rightarrow \Gamma$ such that for all $\underline{\pi}$,*

$$(4.13) \quad v(\underline{\pi}) = \min_{\gamma \in \Gamma} v\left(\tilde{Q}(\gamma)\underline{\pi}\right) = v\left(\tilde{Q}(\tilde{\Delta}^*(\underline{\pi}))\underline{\pi}\right),$$

and

$$(4.14) \quad \min_{\gamma \in \Gamma} \left\{ \bar{\rho}(\underline{\pi}, \gamma) + r(\tilde{Q}(\gamma)\underline{\pi}) \right\} \leq v(\underline{\pi}) + r(\underline{\pi}) \leq \bar{\rho}(\underline{\pi}, \tilde{\Delta}^*(\underline{\pi})) + r(\tilde{Q}(\tilde{\Delta}^*(\underline{\pi}))\underline{\pi}) + \varepsilon.$$

Then for any horizon T and any meta-design $\tilde{\Delta}^T := (\tilde{\Delta}_1, \dots, \tilde{\Delta}_T)$, the stationary meta-design $\tilde{\Delta}^{*,T} := (\tilde{\Delta}^*, \dots, \tilde{\Delta}^*)$ (T -times) satisfies

$$(4.15) \quad \mathcal{J}_T(\tilde{\Delta}^{*,T}) = r(\underline{\pi}_1) + Tv(\underline{\pi}_1) \leq \mathcal{J}_T(\tilde{\Delta}^T) + T\varepsilon$$

Further, the stationary meta-design $\tilde{\Delta}^{*,\infty} := (\tilde{\Delta}^*, \tilde{\Delta}^*, \dots)$ is ε -optimal. That is, for any infinite horizon meta-design $\tilde{\Delta}^\infty := (\tilde{\Delta}_1, \tilde{\Delta}_2, \dots)$ we have

$$(4.16) \quad \overline{\mathcal{J}}(\tilde{\Delta}^{*,\infty}) = v(\underline{\pi}_1) \leq \underline{\mathcal{J}}(\tilde{\Delta}^\infty) + \varepsilon,$$

where

$$(4.17) \quad \overline{\mathcal{J}}(\tilde{\Delta}^{*,\infty}) := \limsup_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T \bar{\rho}(\underline{\pi}_t, \tilde{\Delta}^*(\underline{\pi}_t))$$

with $\underline{\pi}_{t+1} = \tilde{Q}(\tilde{\Delta}^*(\underline{\pi}_t)\underline{\pi}_t)$ and

$$(4.18) \quad \underline{\mathcal{J}}(\tilde{\Delta}^\infty) := \liminf_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T \bar{\rho}(\underline{\pi}_t, \tilde{\Delta}_t(\underline{\pi}_t))$$

with $\underline{\pi}_{t+1} = \tilde{Q}(\tilde{\Delta}_t(\underline{\pi}_t)\underline{\pi}_t)$.

Proof. This is a standard result, see [10, Chapter 7]. \square

4.3. Discussion of the results. The results of this section show that for infinite horizon problems stationary designs are not necessarily optimal (or ε -optimal). In view of the discussion in Section 2.5, this result is not surprising. The design rules c_t, l_t, g_t of the problems under consideration correspond to the control actions and the meta-functions correspond to the control law in classical deterministic optimization problems. In classical infinite horizon deterministic optimization problems restricting attention to stationary *control laws* does not entail any loss in optimality; however even for a stationary control law, control actions change with time. By analogy, in the infinite horizon problems considered in this section, restricting attention to stationary meta-designs does not entail any loss in optimality; however even for a stationary meta-design, optimal design rules change with time. In the absence of

a systematic framework, the task of finding and implementing an optimal infinite horizon design is intractable. Conveniently, the methodology and results presented in this section suggest a method to obtain and implement time-varying optimal designs i.e., obtain and implement optimal stationary meta-designs. The off-line problem simplifies to obtaining the fixed point of a functional equation, which also gives an optimal stationary meta-design. This meta-design can be implemented at the sensor and the controller. When the system is running, the sensor and the controller need to keep track of the information state of the system, and use the meta-design and the current information state to determine the current optimal design rules. This greatly simplifies the on-line implementation of a time-varying optimal design.

4.4. Some additional remarks. The remarks made in Section 3.4 of [16] also apply to Problems 4.1 and 4.2. For completeness, we present them here.

1. In Theorem 4.3 the fixed point equation (4.8) can be simplified as

$$(4.19) \quad \underline{V}(\underline{\pi}) = \min_{c \in \mathcal{C}} V'(Q(c)\underline{\pi}),$$

$$(4.20) \quad V'(\pi) = \min_{l \in \mathcal{L}} \bar{V}(Q(l)\pi),$$

$$(4.21) \quad \bar{V}(\bar{\pi}) = \inf_{g \in \mathcal{G}} \{ \tilde{\rho}(\bar{\pi}, g) + \underline{V}_{t+1}(\bar{Q}(g)\bar{\pi}) \},$$

with \underline{V} being equivalent to V of (4.8). Here we are further decomposing the problem into its “natural” sequential form. This system of equations (4.19)–(4.21) is the infinite horizon analogue of the optimality equations (2.51)–(2.53) of Theorem 2.7. The system (4.19)–(4.21) may be easier to solve (numerically) than (4.8).

2. In Theorem 4.3 taking the limit $\beta \rightarrow 1$ in the expected discounted cost problem *does not* lead to a solution of average cost per unit time problem. Such a result is valid only when the problem has finite state and action space (see [41, Theorem 31.5.2]), which is not the case here. For partially observed Markov decision problems (POMDP), one needs further assumptions for such a result to hold. See [2] for a survey of various results connecting the expected discounted cost problem with the average cost per unit time problem.
3. The standard dynamic program for average cost per unit time (see [14, Chapter 8]) assumes that v of Theorem 4.4 is a finite constant. This assumption can be too restrictive for problems with uncountable action spaces. We allow v to be a bounded function as long as it is a fixed point of (4.13). This is a generalization of the standard assumption (see [10, Chapter 7] for details).
4. Relation (4.16) in Theorem 4.4 is a stronger result than we were seeking as a solution of Problem 4.2. An interpretation of (4.16) is that the “most pessimistic” average performance under (C^*, L^*, G^*) is no more than ε worse than the “most optimistic” performance under any other design.
5. Conditions that guarantee that assumption (A1) of Theorem 4.4 is satisfied are fairly technical and do not provide much insight into the properties of the plant, the channel, and the cost functions that will guarantee the existence of such policies. The interested reader may look at [10, Chapter 7, §10]. It may be possible to extend the sufficiency conditions of [31–33] to uncountable action spaces.

5. Uncountable State Space. Consider the model of Section 2.1 with the following differences: the state of the plant X_t , the plant disturbance W_t and the

control action U_t belong to uncountable spaces, i.e., $\mathcal{X} = \mathbb{R}^{d_x}$, $\mathcal{W} = \mathbb{R}^{d_w}$ and $\mathcal{U} = \mathbb{R}^{d_u}$ where d_x , d_w and d_u are positive integers. The initial state X_1 is a random variable belonging to $(\mathbb{R}^{d_x}, \mathcal{B}(\mathbb{R}^{d_x}), \mu_{X_1})$, where $\mathcal{B}(\mathbb{R}^{d_x})$ is the Borel σ -algebra on \mathbb{R}^{d_x} and the probability law μ_{X_1} is given. The plant disturbance W_1, \dots, W_T are i.i.d. random variables belonging to $(\mathbb{R}^{d_w}, \mathcal{B}(\mathbb{R}^{d_w}), \mu_W)$ where the probability law μ_W is given. The rest of the model is the same as that of Section 2.1. The sensor has finite memory \mathcal{M} and the channel is a discrete memoryless channel with input \mathcal{Z} and output \mathcal{Y} . The plant function f , the design (C, L, G) , $C := (c_1, \dots, c_T)$, $L := (l_1, \dots, l_T)$, $G := (g_1, \dots, g_T)$ and the cost ρ are Borel measurable with respect to appropriate σ -algebras. The objective is to choose a design (C, L, G) that minimizes the total expected cost under that design.

The fact that the state of the plant, the plant disturbance, and the control action belong to uncountable spaces does not change the problem fundamentally. The methodology of Section 2 applies here—the technical details are a bit more involved. Notice that the existence of an optimal design is not guaranteed for this problem. However, since the function spaces are compact, there exist ε -optimal design.

5.1. Solution methodology. For a fixed encoder, the design of an optimal controller is a centralized stochastic control problem as in the case of Problem 2.1. We need to modify the definition of beliefs, given by Definition 2.2, to take the uncountable state space into account.

DEFINITION 5.1. *For any $A_X \in \mathcal{B}(\mathbb{R}^{d_x})$, $A_M \in 2^{\mathcal{M}}$, where $2^{\mathcal{M}}$ denotes the power set of \mathcal{M} , define the measurable transforms \underline{B}_t , B_t and \overline{B}_t as follows:*

1. $\underline{B}_t(A_X, A_M) := \Pr(X_t \in A_X, M_{t-1} \in A_M \mid \underline{I}_t)$.
2. $B_t(A_X, A_M) := \Pr(X_t \in A_X, M_{t-1} \in A_M \mid I_t)$.
3. $\overline{B}_t(A_X, A_M) := \Pr(X_t \in A_X, M_t \in A_M \mid \overline{I}_t)$.

Lemma 2.3 can be proved as before by using Bayes rule for continuous valued random variables. Lemma 2.3 implies that Theorem 2.4 also holds in this case. Thus without loss of optimality, we can restrict attention to controllers of the form

$$(5.1) \quad U_t = g_t(\overline{B}_t).$$

that is controller belonging to \mathcal{G}_S , the family of functions from $\mathcal{P}^{\mathcal{X} \times \mathcal{M}}$ to \mathbb{R}^{d_u} that are $\mathcal{B}(\mathcal{P}^{\mathcal{X} \times \mathcal{M}})/\mathcal{B}(\mathbb{R}^{d_u})$ measurable. Thus, at each stage we can optimize over a fixed (rather than a time varying) domain.

With this reduction, we can define information states $\underline{\pi}_t$, π_t , $\overline{\pi}_t$ as in Definition 2.5, with the beliefs given by Definition 5.1. It is easy to show that these information states satisfy Lemma 2.6. Thus they are sufficient for performance evaluation and lead to a sequential decomposition of the problem. An ε -optimal design can be obtained by the nested optimality equations (2.50)–(2.53). Similar results extend to infinite horizon problems using the ideas of Section 4.

5.2. Computational issues. Numerically, problems where the state space is uncountable are much harder than the problems with finite state space. This increase in complexity does not arise from the increase in dimensionality of the information state; as a matter of fact, the information states for finite and uncountable state spaces problems belong to isomorphic spaces. The uncountable state space problems are harder to solve due to the increase in the complexity of the action space. Let us first consider some results from probability theory [10, Appendices 1–5] to show that these information states belong to isomorphic spaces.

DEFINITION 5.2 (Borel Space). *A measurable space B is called Borelian or a Borel space if it is isomorphic to a measurable subset of a Polish (i.e., a complete separable metric) space E .*

Consider the following Borel spaces

1. A finite or countable space D , with the σ -algebra of all subsets.
2. The unit interval J with the σ -algebra of all open subintervals.

THEOREM 5.3. *Every Borel space is isomorphic to either D or to J .*

THEOREM 5.4. *Suppose that \mathcal{P}^E is the set of all probability measures on the space E . If E is a Borel space, then \mathcal{P}^E is also a Borel space.*

For the finite state space problem, let E denote the space $\mathcal{X} \times \mathcal{M}$ with σ -algebra $2^{\mathcal{X} \times \mathcal{M}}$; for the uncountable state space problem, let E denote the space $\mathbb{R}^{d_x} \times \mathcal{M}$ with σ -algebra $\mathcal{B}(\mathbb{R}^{d_x}) \times 2^{\mathcal{M}}$. Then E is Borelian and by Theorem 5.4, the space \mathcal{P}^E of probability measures on E is a Borel space. By the same argument, the space Π of probability measures on $(E \times \mathcal{P}^E, \mathcal{B}(E \times \mathcal{P}^E))$ is a Borel space. Thus the information state for the finite state space problem and the information state for the infinite state space problem are isomorphic; each being isomorphic to J , the unit interval with σ -algebra of all open subintervals.

The dimensionality of the information state is only one component that determines the complexity of the numerical solution, the dimensionality of the action space being the other one. In our problem the action spaces alternate between \mathcal{C} , \mathcal{L} , and $\hat{\mathcal{G}}$. For the finite state space problem \mathcal{C} , \mathcal{L} and $\hat{\mathcal{G}}$ are the family of functions from $\mathcal{X} \times \mathcal{M}$ to \mathcal{Z} , $\mathcal{X} \times \mathcal{M}$ to \mathcal{M} , and Π to \mathcal{U} , respectively. For the uncountable state space problem \mathcal{C} , \mathcal{L} and $\hat{\mathcal{G}}$ are the family of functions from $\mathbb{R}^{d_x} \times \mathcal{M}$ to \mathcal{Z} , $\mathbb{R}^{d_x} \times \mathcal{M}$ to \mathcal{M} , and Π to \mathbb{R}^{d_v} , respectively. Thus the complexity of all three function spaces increases when we go to the uncountable state space problems; this increase in complexity makes it harder to obtain numerical solutions in the case of uncountable state space problems.

6. Imperfect Observations. So far we have assumed that the sensor perfectly observes the state of the plant. However in many practical systems, the sensor observations are noisy due to external disturbances and the intrinsic noise in the measurement hardware. In this section we model this scenario and show that noisy observations by the sensor do not alter the nature of the problem. We first consider the finite horizon case.

6.1. Problem formulation. Consider a discrete time imperfect observation system as shown in Figure 6.1 which operates for T time steps. The state of the plant X_t evolves according to (2.1). The observations S_t made by the observer at time t are noise-corrupted version of the state of the plant and are given by

$$(6.1) \quad S_t = \hat{h}(X_t, \hat{N}_t),$$

where \hat{N}_t denotes the observation noise and \hat{h} is the observation channel. S_t takes values in $\mathcal{S} := \{1, \dots, |\mathcal{S}|\}$ and \hat{N}_t takes values in $\hat{\mathcal{N}} := \{1, \dots, |\hat{\mathcal{N}}|\}$. The sequence of random variables $\hat{N}_1, \dots, \hat{N}_T$ are i.i.d. with PMF $P_{\hat{\mathcal{N}}}$. The sequence $\hat{N}_1, \dots, \hat{N}_T$ is also independent of $X_1, W_1, \dots, W_T, N_1, \dots, N_T$.

The sensor is modeled as in Section 2.1 and operates as follows

$$(6.2) \quad Z_t = c_t(S_t, M_{t-1}),$$

$$(6.3) \quad M_t = l_t(S_t, M_{t-1}).$$

7. Conclusion. We have presented a methodology for determining globally optimal (or globally ε -optimal) encoding and control strategies for feedback control systems with limited communication over noisy channels. The methodology is applicable to finite horizon problems with expected total cost criterion, to infinite horizon problems with expected discounted cost criterion, and to infinite horizon problems with average cost per unit time criterion. We have extended this methodology to problems where the encoder/sensor makes imperfect observations of the state of the system. The resulting optimality equations can be viewed as POMDPs where the state space is a real valued vector and the action space is uncountable. There are very few results on efficient computational techniques for this class of POMDPs. We hope that the problem of optimal control over a noisy communication channel will motivate researchers to investigate numerical methods for optimization problems that are of the type of Problem 2.3.

For the problems considered in this paper, the action space is uncountable because of the assumption of perfect recall at the controller's site. In light of the sequential decomposition for decentralized team problems presented in this paper, this assumption of perfect recall needs to be reconsidered. For most applications the assumption of perfect recall, that is, the assumption that an agent remembers everything that it has seen and everything that it has done in the past, is impractical. Nevertheless, in centralized stochastic control problems perfect recall is assumed since it implies a classical information structure, and simplifies the solution methodology. In decentralized problems, the information structure is non-classical, and remains non-classical even with the unrealistic assumption of perfect recall at each agent's site. Further, the assumption of perfect recall makes it harder to obtain a numerical solution of the resultant nested optimality equations. In the problems considered in this paper, the sensor/encoder has finite memory while the controller has perfect recall. In the nested optimality equations of Theorem 2.7, to obtain an optimal encoder and memory update rule in (2.51) and (2.52) we need to choose c_t and l_t belonging to \mathcal{C} and \mathcal{L} respectively; both \mathcal{C} and \mathcal{L} are finite spaces. On the other hand, to obtain an optimal controller in (2.53) we need to choose g_t belonging to \mathcal{G} , which is an uncountable space; even though the action space \mathcal{U} is finite, to choose an optimal g_t we have to search over an uncountable space. If we had assumed a finite memory at the controller, we would have obtained equations where we need to choose a control law and a memory update rule at the controller from a finite set, and this problem is similar to a POMDP with finite action space. Thus, the unrealistic assumption of perfect recall at any agent's site does not simplify the analysis rather makes the problem numerically more difficult to solve while the realistic assumption of a finite memory at all agent's site results in a solution algorithm that is easier to solve.

It is important to identify special cases in which the information states $\bar{\pi}_t$, π_t and $\bar{\pi}_t$ can be restricted to a parametric family of distributions. In centralized stochastic control problems, LQG systems possess such a property—the information state can be restricted to Gaussian distributions. This is because in LQG systems with classical information pattern, without any loss of optimality we can restrict attention to affine control laws, which implies that the state of the plant is always Gaussian. Thus the information state—which is the conditional probability of the state of the plant, conditioned on all the past observations and all the past control actions of the controller—is also Gaussian. This simplifies the search for an optimal design. Unfortunately, in decentralized systems (more precisely, in systems with non-classical information structure) non-linear control laws can outperform affine control laws even

in linear systems where all primitive random variables are Gaussian, as illustrated by the Witsenhausen counterexample [42]. So, the state of the plant may not be Gaussian and hence the information state need not be Gaussian. However, there may be other special cases for which information states in a decentralized system belong to a parametric family of distributions. Finding such special cases remains a challenging open problem.

The results of Section 4 show that for infinite horizon problems stationary designs are not optimal and in order to implement a time-varying optimal design, we need to implement an optimal stationary meta-design. Thus, implementing optimal designs for decentralized systems is an order of magnitude more complicated as compared to centralized systems. Traditionally, for infinite horizon decentralized control problems, performance limitations of only stationary designs is considered. It will be worthwhile to characterize the performance difference between an optimal time-varying design and the best stationary design. It will also be important to obtain performance limitations of time-varying optimal designs.

Acknowledgments. This research was supported in part by ONR Grant N00014-03-1-0232, NSF Grant CCR-0325571 and NASA Grant NNX06AD47G. We are grateful to the referees and the associate editor for their comments which helped us to significantly improve the presentation of the paper.

Appendix A. Nested σ algebras. We first present a general lemma and then use its result to justify the statement made in the discussion in Section 3.

LEMMA A.1. *Consider a probability space (Ω, \mathcal{F}, P) . Let X and Y be real-valued random variables defined on (Ω, \mathcal{F}, P) , and let $g : (\mathbb{R}, \mathcal{B}(\mathbb{R})) \rightarrow (\mathbb{R}, \mathcal{B}(\mathbb{R}))$ be a measurable real-valued function. Then*

$$(A.1) \quad \sigma(X) \cap \sigma(Y) \subseteq \sigma(X, g(Y)).$$

Proof. Consider any set A belonging to $\sigma(X) \cap \sigma(Y)$. Then, there exist sets B_1 and B_2 belonging to $\mathcal{B}(\mathbb{R})$ such that $A = X^{-1}(B_1)$ and $A = Y^{-1}(B_2)$. Define a real-valued random variable Z on (Ω, \mathcal{F}, P) by $Z(\omega) = g(Y(\omega))$. Let $B_3 := g(B_2)$. Now, $g^{-1}(B_3) \supseteq B_2$, so $Z^{-1}(B_3) := Y^{-1}(g^{-1}(B_3)) \supseteq Y^{-1}(B_2) = A$. Thus,

$$(A.2) \quad X^{-1}(B_1) \cap Z^{-1}(B_3) = A.$$

Hence, $A \in \sigma(X, Z)$, and thus

$$(A.3) \quad \sigma(X) \cap \sigma(Y) \subseteq \sigma(X, Z) = \sigma(X, g(Y)).$$

□

Now, in the discussion in Section 3, we claimed that

$$(A.4) \quad K_t(\gamma^t) := \sigma(X_t, M_{t-1}) \cap \sigma(Y^t, U^{t-1}) \subseteq \sigma(X_t, M_{t-1}, \underline{B}_t) =: \hat{K}_t(\gamma^t).$$

This follows by taking $X = (X_t, M_{t-1})$, $Y = (Y^t, U^{t-1})$ and $g(Y^t, U^{t-1}) = \Pr(X_t, M_{t-1} | Y^t, U^{t-1}, \gamma^t)$ in Lemma A.1.

REFERENCES

- [1] P. J. ANTSAKLIS AND J. BAILLIEUL, eds., *IEEE Trans. Autom. Control: Special Issue on Networked Control Systems*, vol. 49, Sept. 2004.

- [2] A. ARAPOSTATHIS, V. S. BORKAR, E. FERNANDEZ-GAUCHERAND, M. K. GHOSH, AND S. MARCUS, *Discrete-time controlled Markov processes with average cost criterion - a survey*, SIAM Journal of Control and Optimization, 31 (1993), pp. 282–344.
- [3] R. J. AUMANN, *Agreeing to disagree*, Annals of Statistics, (1976), pp. 1236–39.
- [4] J. BAILLIEUL, *Feedback designs in information based control*, in Proc. Workshop in Stochastic Theory and Control, B. Pasik-Duncan, ed., Springer Verlag, Oct. 2001, pp. 35–27.
- [5] ———, *Feedback coding for information based control — operating near the data-rate limit*, in Proceedings of 41st IEEE Conference on Decision and Control, Dec. 2002, pp. 3229–3236.
- [6] J. H. BRASLAVSKY, R. H. MIDDLETON, AND J. S. FREUDENBERG, *Feedback stabilization over signal-to-noise ratio constrained channels*, submitted to the IEEE Trans. Autom. Control, (2005).
- [7] R. W. BROCKETT AND D. LIBERZON, *Quantized feedback stabilization of linear systems*, IEEE Trans. Autom. Control, 45 (2000), pp. 1279–1289.
- [8] G. W. BROWN, *Iterative solutions of games by fictitious play*, in Activity analysis of production and allocation, T. C. Koopmans, ed., Wiley, 1951, pp. 374–376.
- [9] D. F. DELCHAMPS, *Stabilizing a linear system with quantized state feedback*, IEEE Trans. Autom. Control, 35 (1990), pp. 916–924.
- [10] E. B. DYNKIN AND A. A. YUSHKEVICH, *Controlled Markov Processes*, A Series of Comprehensive Studies in Mathematics, Springer-Verlag, 1975.
- [11] N. ELIA AND S. K. MITTER, *Stabilization of linear systems with limited information*, IEEE Trans. Autom. Control, 46 (2001), pp. 1384–1400.
- [12] Y.-C. HO, *Team decision theory and information structures*, Proc. IEEE, 68 (1980), pp. 644–654.
- [13] H. ISHII AND B. A. FRANCIS, *Quadratic stabilization of sampled-data systems with quantization*, Automatica, 39 (2003), pp. 1793–1800.
- [14] P. R. KUMAR AND P. VARAIYA, *Stochastic Systems: Estimation Identification and Adaptive Control*, Prentice Hall, 1986.
- [15] D. LIBERZON, *On stabilization of linear systems with limited information*, IEEE Trans. Autom. Control, 48 (2003), pp. 304–307.
- [16] A. MAHAJAN AND D. TENEKETZIS, *On jointly optimal encoding, decoding and memory update for noisy real-time communication systems*, submitted to IEEE Trans. Inf. Theory, (2006). Available as Control Group Report CGR-06-03, Department of EECS, University of Michigan, Ann Arbor, MI 48109-2122.
- [17] J. MARSCHAK AND R. RADNER, *Economic Theory of Teams*, Yale University Press, New Haven, 1972.
- [18] N. MARTINS, M. DAHLEH, AND N. ELIA, *Stabilization of uncertain systems in the presence of a stochastic digital link*, submitted to IEEE Trans. Autom. Control, (2004).
- [19] N. C. MARTINS AND M. A. DAHLEH, *Feedback control in the presence of noisy channels: “bode-like” fundamental limitations of performance*, submitted to IEEE Trans. Autom. Control, (2004).
- [20] N. C. MARTINS, M. A. DAHLEH, AND J. C. DOYLE, *Fundamental limitations of feedback in the presence of side information*, submitted to IEEE Trans. Autom. Control, (2005).
- [21] A. S. MATVEEV AND A. V. SAVKIN, *Problem of LQG optimal control via a limited capacity communication channel*, System and Control Letters, (2004), pp. 51–64.
- [22] G. N. NAIR AND R. J. EVANS, *Stabilization with data-rate limited feedback: Tightest attainable bounds*, Systems Control Letters, 41 (2000), pp. 304–307.
- [23] ———, *Exponential stabilisability of finite-dimensional linear systems with limited data rates*, Automatica, 39 (2003), pp. 585–593.
- [24] ———, *Stabilization of stochastic linear systems with finite feedback data rates*, SIAM Journal of Optimal Control, 43 (2004), pp. 413–436.
- [25] G. N. NAIR, R. J. EVANS, I. M. Y. MAREELS, AND W. MORAN, *Topological feedback entropy and nonlinear stabilization*, IEEE Trans. Autom. Control, 49 (2004), pp. 1585–1597.
- [26] G. N. NAIR, F. FAGNANI, S. ZAMPIERI, AND R. J. EVANS, *Feedback control under data rate constraints: An overview*, Proc. IEEE, 95 (2007), pp. 108–137.
- [27] I. R. PETERSON AND A. V. SAVKIN, *Multi-rate stabilization of multivariable discrete-time linear systems via a limited capacity communication channel*, in Proceedings of 40th IEEE Conference on Decision and Control, 2001, pp. 304–309.
- [28] J. ROBINSON, *An iterative method for solving a game*, Annals of mathematics, 54 (1951), pp. 296–301.
- [29] N. R. SANDELL, JR., *Control of Finite-State, Finite-Memory Stochastic Systems*, ph.d. thesis, Massachusetts Institute of Technology, Cambridge, MA., 1974.
- [30] A. V. SAVKIN, *Analysis and synthesis of networked control systems: Topological entropy, ob-*

- servability, robustness and optimal control*, Automatica, 42 (2006), pp. 51–62.
- [31] L. I. SENNOTT, *The computation of average optimal policies in denumerable state Markov decision chains*, Advances in Applied Probability, 29 (1997), pp. 114–137.
- [32] ———, *On computing average cost optimal policies with application to routing parallel queues*, ZOR Mathematical Methods in Operations Research, 45 (1997), pp. 45–62.
- [33] ———, *Stochastic dynamic programming and the control of queueing systems*, Wiley-Interscience, New York, NY, USA, 1999.
- [34] L. S. SHAPLEY, *Some topics in two-person games*, in Advances in game theory, M. Dresher, L. S. Shapley, and A. W. Tucker, eds., vol. 52 of Annals of Mathematics Studies, Princeton University Press, 1964, pp. 1–28.
- [35] S. TATIKONDA, *Control Under Communication Constraints*, ph.d. thesis, Department of EECS, Massachusetts Institute of Technology, Cambridge, MA, 2000.
- [36] S. TATIKONDA AND S. K. MITTER, *Control over noisy channels*, IEEE Trans. Autom. Control, 49 (2004), pp. 1196–1201.
- [37] ———, *Control under communication constraints*, IEEE Trans. Autom. Control, 49 (2004), pp. 1056–1068.
- [38] S. TATIKONDA, A. SAHAI, AND S. MITTER, *Stochastic linear control over a communication channel*, IEEE Trans. Autom. Control, 49 (2004), pp. 1549–1561.
- [39] S. TATIKONDA, A. SAHAI, AND S. K. MITTER, *Stochastic linear control over a communication channel*, IEEE Trans. Autom. Control, 49 (2004), pp. 1549–1561.
- [40] J. C. WALRAND AND P. VARAIYA, *Causal coding and control of markov chains*, System and Control Letters, 3 (1983), pp. 189–192.
- [41] P. WHITTLE, *Optimization Over Time*, vol. 2 of Wiley Series in Probability and Mathematical Statistics, John Wiley and Sons, 1983.
- [42] H. S. WITSENHAUSEN, *A counterexample in stochastic optimum control*, SIAM Journal of Optimal Control, 6 (1968), pp. 131–147.
- [43] ———, *Separation of estimation and control for discrete time systems*, Proc. IEEE, 59 (1971), pp. 1557–1566.
- [44] ———, *A standard form for sequential stochastic control*, Mathematical Systems Theory, 7 (1973), pp. 5–11.
- [45] ———, *The intrinsic model for discrete stochastic control: Some open problems*, in Control Theory, Numerical Methods and Computer System Modelling, A. Bensoussan and J. L. Lions, eds., vol. 107 of Lecture Notes in Economics and Mathematical Systems, Springer Verlag, 1975, pp. 322–335.
- [46] ———, *Some remarks on the concept of state*, in Directions in Large-Scale Systems, Y. C. Ho and S. K. Mitter, eds., Plenum, 1976, pp. 69–75.
- [47] W. S. WONG AND R. W. BROCKETT, *Systems with finite communication bandwidth constraints II: Stabilization with limited information feedback*, IEEE Trans. Autom. Control, 44 (1999), pp. 1049–1053.