Proceedings of the 35th
Conference on Decision and Control
Kobe, Japan • December 1996

Decentralized Control of a Multiple Access Broadcast Channel: Performance Bounds

James M. Ooi and Gregory W. Wornell
Research Laboratory of Electronics
Massachusetts Institute of Technology
Cambridge, Massachusetts 02139 USA

Abstract

Optimal decentralized control of the multiple access broadcast channel is considered. A technique is presented for upper bounding the throughput of a slotted multiple access system with a finite number of users, immediate ternary feedback, retransmission of collisions, and no buffering. The upper bound is calculated for the two- and three-user cases, and it is shown that Hluchyj and Gallager's optimized window protocol is effectively optimal for these cases.

1 Introduction

The multiple access broadcast channel (MABC) is a useful model for a variety of packet switched communication systems. For this channel, there has been considerable interest in the development of efficient protocols for coordinating the data transmissions of the users. Although several variations of the MABC have been considered in the literature (see, e.g., [1]—[7]), we focus on a finite-user slotted system with immediate ternary feedback, retransmission of collisions, no buffering, and no communication among users; we refer to this system as the *canonical* system.

Previous research has exploited the fact that the MABC protocol design problem can be analyzed as a decentralized control problem [2]—[7]. However, it has generally been necessary to adopt simplifications that make the problem tractable but also removed from the canonical problem. For example, Schoute [6] and Varaiya [7] both consider decentralized control of the MABC under a delayed sharing pattern and under the assumption that colliding packets incur a fixed cost rather than requiring retransmission; Rosberg [5] also assumes a fixed collision cost and no retransmissions, but differs by assuming no information sharing among controllers, as well as control inputs that depend only on broadcast feedback. Although these two simplified problems are easier to analyze, their relationships to the canonical problem are unclear. On the other hand, the simplified problems considered by Hluchyj and Gallager [4], Grizzle et al. [2], and Paradis [3] yield solutions that can be used to find lower and upper bounds on the throughput of the canonical system. Hluchyj and Gallager consider the canonical system and find protocols that are optimal in the class of protocols known as the window protocols. Since window protocols are a subset of the set of all protocols, the throughput of Hluchyj and Gallager's optimized window protocol provides a *lower* bound to the throughput achievable in the canonical system. Conversely, Grizzle et al. and Paradis attack the problem of optimally controlling a MABC that is canonical except for a one-step delay sharing (OSDS) information pattern. Because the canonical system does not allow any communication among users, its throughput is *upper* bounded by the throughput achievable under OSDS. For two users, the OSDS bound is very close to the throughput of the Hluchyj-Gallager optimized window protocol [2]—[4]. Unfortunately, for more than two users, the OSDS bound is close to neither the throughput of the Hluchyj-Gallager protocol nor any other known protocol.

We introduce new upper bounds on the throughput of the canonical system via the K-step delay state information pattern, which is similar to the previously considered K-step delay sharing pattern [8]. That is, we calculate the throughput of a MABC that is canonical except for a K-step delay state information pattern, and this provides an upper bound on the throughput of the canonical system. We use the dynamic programming method developed by Aicardi et al. [9] as a starting point for performing this calculation. However, for the class of systems that includes the MABC with K-step delay state information, this method is unnecessarily expensive in terms of computation. For this class, we develop a more computationally efficient version of the algorithm that that extends its practical utility. As a result of the more efficient algorithm, we are able to find upper bounds on the throughput of the canonical system for two and three users that are tighter than the OSDS bound. The new bounds show that the performance of Hluchyj and Gallager's optimized window protocol is optimal for two users and at least nearly optimal for three users, where optimality is with respect to maximizing throughput. In fact, the Hluchyj-Gallager protocol meets the upper bound for three users when the packet arrival probability is moderately large.

The bounding technique we present is quite flexible and can be adapted to handle a greater number of users as well as other variations on the problem. Indeed, the bounding technique applies to the general class of decentralized control problems with no information sharing and may be useful in a broad range of other decentralized control and multiple-access communication problems. For this reason, we first describe the bounding technique in this general setting and later focus attention on the MABC.

This work has been supported in part by the Advanced Research Projects Agency monitored by ONR under Contract No. N00014-93-1-0686, the Office of Naval Research under Grant No. N00014-95-1-0834, and an AT&T Doctoral Fellowship.

0-7803-3590-2/96 $5.00 © 1996 IEEE

2 Decentralized Systems with No Sharing

In this section, we describe the class of systems for which the bounding technique to be described in Section 4 is applicable. This class of systems consists of decentralized systems with no sharing, i.e., systems in which no information is communicated among controllers. We show in Section 5 that the MABC can be regarded as such a system.

Consider a discrete-time stochastic system that is regulated by M decentralized controllers, each with an associated measurement station. The system state variable x_t and the mth measurement station's observations y_t^m evolve for $t = 0, 1, \cdots, T-1$ according to the equations

$$x_{t+1} = f_t(x_t, u_t^1, \cdots, u_t^M, v_t), \qquad (1)$$

$$y_t^m = g_t^m(x_t, w_t^m),\ m = 1, \ldots, M, \qquad (2)$$

where f_t and g_t^m are given functions, u_t^m represents the mth controller's input at time t, and the quantities $x_0, (v_0, w_0^1, \cdots, w_0^M), \cdots, (v_{T-1}, w_{T-1}^1, \cdots, w_{T-1}^M)$ are mutually independent primitive random variables whose distributions are known. At a particular time step t, the random variables $v_t, w_t^1, \cdots, w_t^M$ are allowed to be statistically dependent. The vectors u_t^m, v_t, w_t^m, x_t, and y_t^m take values in the given *finite* sets U_t^m, V_t, W_t^m, X_t, and Y_t^m, respectively. We assume that these sets are disjoint, e.g., $X_1 \cap X_2 = \emptyset$, and $W_1^1 \cap Y_1^1 = \emptyset$.[1]

Each controller produces, according to a pre-designed control law, an input based only on local observations from its own measurement station. If the mth controller is governed at time t by a control law γ_t^m then $u_t^m = \gamma_t^m(y_0^m, y_1^m, \cdots, y_t^m)$. It is this functional dependence of u_t^m on only the history of the mth measurement station's measurements that makes the problem one of no sharing.

The set of admissible control laws, Γ_t^m, is the set of all functions mapping $Y_0^m \times \cdots \times Y_t^m$ to U_t^m. The design objective is to choose the control laws $\{\gamma_t^m \in \Gamma_t^m\}_{m=1,t=0}^{m=M,t=T-1}$ to minimize the total expected cost per stage

$$\frac{1}{T} \sum_{t=0}^{T-1} E[h_t(x_{t+1}, u_t^1, \cdots, u_t^M)], \qquad (3)$$

where h_t is a given real-valued cost function.

3 Notational Conventions

We adopt the convention of using context to distinguish between values assumed by random variables and the random variables themselves. The dependence of densities and expectations on a choice of control laws $\gamma_{s:t}$ is indicated by $p(\cdot; \gamma_{s:t})$ and $E[\cdot; \gamma_{s:t}]$, respectively. Domains and ranges of functions are to be inferred from context, but sometimes may be explicitly given for emphasis or clarity.

To consolidate lists of related symbols, we define $y_t = (y_t^1, \cdots, y_t^M)$, $y_{s:t}^m = (y_s^m, \cdots, y_t^m)$, and $y_{s:t} = (y_s, \cdots, y_t)$ for $t \geq s$. If $t < s$, then $y_{s:t}$ and $y_{s:t}^m$ are empty tuples. Moreover, we denote the range of y_t by $Y_t = \prod_{m=1}^M Y_t^m$, the range of $y_{s:t}^m$ by $Y_{s:t}^m = \prod_{j=s}^t Y_j^m$, and the range of $y_{s:t}$ by $Y_{s:t} = \prod_{j=s}^t Y_j$. Analogous notation will be used for other variables and their ranges.

It is convenient to define set operations on tuples as follows. Suppose that $\mathcal{A} = \{A_1, \cdots, A_n\}$ is an ordered collection of disjoint sets such that $A_1 < \cdots < A_n$, i.e., A_1 is the "smallest" element of $\mathcal{A}$, while A_n is the "largest." This ordering of the collection $\mathcal{A}$ is only important for the systematic construction of tuples from sets. Let $a = (a_1, \cdots, a_n) \in \prod_{i=1}^n A_i$ be a tuple. Let $b = (a_{j_1}, \cdots, a_{j_p})$, and $c = (a_{k_1}, \cdots, a_{k_q})$ be tuples which may be called "sub-tuples" of a. Then define $b \setminus c$ to be a tuple consisting of the elements of the set difference $\{a_{j_i}\}_{i=1}^p \setminus \{a_{k_i}\}_{i=1}^q$, ordered so that an element in the tuple precedes another if it belongs to a set that is "smaller" than the set to which the other belongs. Analogous notation holds for $b \cup c$, and $b \cap c$. Finally, if $d = (a_{\ell_1}, \cdots, a_{\ell_r})$, we denote the set product $\prod_{i=1}^r A_{\ell_i}$ by $\mathcal{S}(d)$. In particular, the collection $\mathcal{X} = \{U_t^m, V_t, W_t^m, X_t, X_T, Y_t^m \mid m = 1, \cdots, M,\ t = 0, \cdots, T-1\}$ is, by assumption, a collection of disjoint sets. Let $\vartheta = (u_{0:T-1}, v_{0:T-1}, w_{0:T-1}, x_{0:T}, y_{0:T-1})$, and let the members of $\mathcal{X}$ have an ordering corresponding to the ordering of the elements in the tuple ϑ. We may now consider unions, intersections, and differences of the sub-tuples of ϑ and will do so in the following sections.

4 Bounding Techniques

The only known solution to the decentralized control problem with no sharing described in Section 2 involves exhaustive search, which is computationally expensive in practice and infeasible when the time horizon is infinite. While we do not develop an explicit solution to this problem, we present a technique for lower bounding the optimal cost achievable by the system in Section 2. The bounds are useful for the evaluation of suboptimal control laws; also, when a control law happens to achieve the bound, we know it is optimal.

The optimal cost achievable by a system with the K-step delay state information pattern, which will be described shortly, is a lower bound to the optimal cost achievable in the no-sharing system. This is because more information is available to controllers under the K-step delay state information pattern than under the no sharing pattern. Furthermore, control laws for systems with the K-step delay state information pattern can be found relatively efficiently [9]. As we will see, however, the algorithm in [9] is more computationally complex than necessary in the special case when all measurement stations have some observations in common. For this special case, we present a more computationally efficient version of the algorithm.

4.1 Systems with K-Step Delay State Information

Under the K-step delay state information pattern, all measurement stations observe the state of the system with a K-step delay. This information pattern is similar to the K-step delay sharing pattern considered in, e.g., [8]. Indeed, K-

[1] This assumption is not essential and is present only to facilitate notation in subsequent sections. Indeed, we will focus on the stationary case, which is at odds with this assumption, but this difficulty can be ignored.

step delay state and K-step delay sharing are equivalent in the case of noiseless observations, i.e., the case in which x_t can be determined from y_t. For this noiseless case, Aicardi et al. [9] show how to efficiently find optimal control laws when the input and state spaces are finite. We adapt the algorithm from [9] to the general K-step delay state control problem.

A system with the K-step delay state information pattern is the same as the system described in Section 2 except for the following modifications. With $\delta_t = (y_{0:t-1}, x_{0:t})$ and the range of δ_t denoted by Δ_t, an admissible control law γ_t^m for the K-step delay state problem must have the form $u_t^m = \gamma_t^m(y_{t-K:t}^m, \delta_{t-K})$. For convenience, we denote $(y_{t-K:t}^m, \delta_{t-K})$ by z_t^m and $Y_{t-K:t}^m \times \Delta_{t-K}$ by Z_t^m. The set of possible control laws for the mth controller at time t is denoted $\check{\Gamma}_t^m$, and consists of all possible maps from Z_t^m to U_t^m. We also change the cost criterion as follows. Since we are interested primarily in the infinite horizon scenario, we assume for convenience that the first K optimal control laws $\gamma_{0:K-1}$ are given and that the goal is to minimize the expected cost per stage starting from stage K:

$$\min_{\gamma_{K:T-1} \in \check{\Gamma}_{K:T-1}} \frac{1}{T-K} E[\sum_{\tau=K}^{T-1} h_\tau(x_{\tau+1}, u_\tau); \gamma_{0:T-1}] \tag{4}$$

Since the first K steps will not affect the limiting behavior, we can just as well choose an arbitrary set of starting control laws $\gamma_{0:K-1}$.

It will be useful to introduce the notion of a "sub-law" as follows. Let $\gamma_t^m \in \check{\Gamma}_t^m$ be a control law, and let η be a sub-tuple of ϑ, and let $N = \mathcal{S}(\eta)$. A sub-law with respect to N is a map from $\mathcal{S}(z_t^m \setminus \eta)$ to U_t^m, and the set of all such maps is denoted $\check{\Gamma}_{t|N}^m$. Define $\gamma_{t|\eta}^m$ to be the sub-law in $\check{\Gamma}_{t|N}$ satisfying $\gamma_{t|\eta}^m(\beta) = \gamma_t^m(\beta \cup \eta)$, $\forall \beta \in \mathcal{S}(z_t^m \setminus \eta)$, where the elements in the tuple $\beta \cup \eta$ are assumed to be in the order required by γ_t^m. Similarly, if $\psi_t^m \in \check{\Gamma}_{t|N}^m$ is sub-law with respect to N, $\hat{\eta}$ is a sub-tuple of ϑ, and $\hat{N} = \mathcal{S}(\hat{\eta})$, then $\psi_{t|\hat{\eta}}^m$ is defined to be a sub-law in $\check{\Gamma}_{t|N\times\hat{N}}^m$ satisfying $\psi_{t|\hat{\eta}}^m(\beta) = \psi_t^m(\beta \cup \hat{\eta})$, $\forall \beta \in \mathcal{S}(z_t^m \setminus (\eta \cup \hat{\eta}))$. Also define the expansion of a sub-law $\psi_t^m \in \check{\Gamma}_{t|N}^m$ to be $G_{t|N}^m(\psi_t^m) = \mu \in \check{\Gamma}_t^m$ such that $\mu(\beta \cup \bar{\eta}) = \psi_t^m(\beta)$, $\forall \beta \in \mathcal{S}(z_t^m \setminus \check{\eta})$, $\forall \bar{\eta} \in S(\check{\eta})$, where $\check{\eta}$ is any element of N. For convenience, we denote $G_{t|N}^m$ by G, assuming that a sub-law always expands into a control law for the corresponding time and controller. Also, we denote $(G(\gamma_{t|\eta}^1), \cdots, G(\gamma_{t|\eta}^M))$ by $G(\gamma_{t|\eta})$ and $(G(\gamma_{s|\eta}), \cdots, G(\gamma_{t|\eta}))$ by $G(\gamma_{s:t|\eta})$.

We can now state a theorem that characterizes the optimal control law for a system with K-step delay state. Note that this theorem is similar to a theorem presented in [9], which is the analogous theorem for the noiseless observation case.

Theorem 1 *Consider the following recursive equations that characterize optimal control laws for the K-step delay state system:*

$$J_T^*(x_{T-K}, \psi_{T-K:T-1}) = 0, \quad \forall \psi_{T-K:T-1} \in \check{\Gamma}_{T-K:T-1|\Delta_{T-K}} \tag{5}$$

$$\begin{aligned} J_t(x_{t-K}, \psi_{t-K:t-1}, \psi_t) &= E[h_t(x_{t+1}, u_t) + J_{t+1}^*(x_{t-K+1}, \psi_{t-K+1:t|y_{t-K}}) \\ &\qquad |x_{t-K}; G(\psi_{t-K:t})], \\ &\forall \psi_{t-K:t} \in \check{\Gamma}_{t-K:t|\Delta_{t-K}} \end{aligned} \tag{6}$$

$$\begin{aligned} J_t^*(x_{t-K}, \psi_{t-K:t-1}) &= \min_{\psi_t \in \check{\Gamma}_{t|\Delta_{t-K}}} J_t(x_{t-K}, \psi_{t-K:t-1}, \psi_t), \\ &\forall \psi_{t-K:t-1} \in \check{\Gamma}_{t-K:t-1|\Delta_{t-K}} \end{aligned} \tag{7}$$

Let $\gamma_{0:K-1} \in \check{\Gamma}_{0:K-1}$ be given starting optimal control laws. If control laws $\gamma_{K:T-1} \in \check{\Gamma}_{K:T-1}$ satisfy

$$J_t(x_{t-K}, \gamma_{t-K:t-1|\delta_{t-K}}, \gamma_{t|\delta_{t-K}}) = J_t^*(x_{t-K}, \gamma_{t-K:t-1|\delta_{t-K}}) \tag{8}$$

for every $\delta_{t-K} \in \Delta_{t-K}$, and every $t = K, \cdots, T-1$, then $\gamma_{K:T-1}$ are optimal.

Proof: The proof of this theorem follows along the same lines as the proof of Theorem 2 in Section 4.2. A proof of a similar theorem is provided in [9]. □

We can interpret the equations in Theorem 1 as the equations resulting from applying the dynamic programming algorithm to the following centralized stochastic control problem. Using the notation of [10], the state for the centralized problem is $x_t' = (x_{t-K}, \psi_{t-K:t-1}) \in X_{t-K} \times \check{\Gamma}_{t-K:t-1|\Delta_{t-K}}$, the input is $u_t' = \psi_t \in \check{\Gamma}_{t|\Delta_{t-K}}$, and the disturbance is $w_t' = (x_{t-K+1}, y_{t-K}) \in X_{t-K+1} \times Y_{t-K}$. The state transition function f_t' is defined by

$$\begin{aligned} f_t'(x_t', u_t', w_t') &= (x_{t-K+1}, \psi_{t-K+1:t|y_{t-K}}), \\ &\forall x_{t-K:t-K+1} \in X_{t-K+1}, \forall y_{t-K} \in Y_{t-K}, \\ &\forall \psi_{t-K+1:t} \in \check{\Gamma}_{t-K+1:t|\Delta_{t-K}}, \end{aligned} \tag{9}$$

and the cost function g_t' is defined by

$$\begin{aligned} g_t'(x_{t+1}', u_t') &= E[h_t(x_{t+1}, u_t)|x_{t-K+1}; G(\psi_{t-K+1:t|y_{t-K}})], \\ &\forall x_{t-K+1} \in X_{t-K+1}, \forall y_{t-K} \in Y_{t-K}, \\ &\forall \psi_{t-K+1:t} \in \check{\Gamma}_{t-K+1:t|\Delta_{t-K}}. \end{aligned} \tag{10}$$

Note that the disturbance satisfies the properties of a disturbance variable, namely that it is independent of previous disturbances given the current state and the input.

A stationary optimal undiscounted infinite horizon control law for the K-step delay state problem will exist if this equivalent centralized control problem has such a stationary optimal infinite horizon solution. The conditions under which such a solution exists are described in [10]. If the conditions are satisfied, then an optimal control law can be found efficiently by known methods, such as Howard's policy iteration algorithm [11].

Immediately, we see from this equivalent centralized stochastic control problem that the state space will likely be enormous, a situation that makes even the dynamic programming algorithm computationally expensive. More precisely, the size of the state space is

$$|X_{t-K} \times \check{\Gamma}_{t-K:t-1|\Delta_{t-K}}| = |X_{t-K}| \cdot \prod_{m=1}^{M} |U_t^m|^{|Y_{t-K:t-1}^m|}, \tag{11}$$

where $|Y_{t-K:t-1}^m|$ grows exponentially with K, implying that the state space grows *doubly* exponentially with K. For a certain class of systems that includes the MABC, the complexity can be substantially reduced, although the growth rate remains doubly exponential. We describe this class and the method by which complexity can be reduced in the following subsection.

4.2 Complexity Reduction

In this section, we demonstrate how substantial computational savings in the bounding algorithm can be realized by exploiting the common information inherent in the MABC problem (viz., the ternary feedback).

A system with common information and K-step delay state retains the elements of the system in Section 4.1 with one modification. Namely, the observation y_t^m can be partitioned into a local observation and a common observation. Specifically, we can write $y_t^m = (\lambda_t^m, \xi_t)$, for all $m = 1, \cdots, M$, where ξ_t is the common observation since every controller observes it, and λ_t^m is the mth controller's local observation. We denote their ranges by Ξ_t and Λ_t^m, respectively. The shared information δ_t can then be written $\delta_t = (\xi_{0:t-1}, \lambda_{0:t-1}, x_{0:t})$, and the control laws take the form

$$u_t^m = \gamma_t^m(y_{t-K:t}^m, \delta_{t-K}) \tag{12}$$

$$= \gamma_t^m(\lambda_{t-K:t}^m, \xi_{t-K:t}, \delta_{t-K}). \tag{13}$$

For future convenience, we denote $(\delta_{t-K}, \xi_{t-K:t})$ by θ_t and its range by Θ_t.

Exploiting the common information, we arrive at the following modified theorem characterizing optimal control laws.

Theorem 2 *Consider the following recursive equations that characterize optimal control laws for the K-step delay state system with common information:*

$$J_T^*(x_{T-K}, \xi_{T-K:T}, \psi_{T-K:T-1}) = 0, \quad \forall \psi_{T-K:T-1} \in \check{\Gamma}_{T-K:T-1|\Theta_T} \tag{14}$$

$$\begin{aligned} J_t(x_{t-K}, \xi_{t-K:t}, \psi_{t-K:t-1}, \psi_t) &= E[h_t(x_{t+1}, u_t) + \\ &\quad J_{t+1}^*(x_{t-K+1}, \xi_{t-K+1:t+1}, \psi_{t-K+1:t|\lambda_{t-K}}) \\ &\quad |x_{t-K}, \xi_{t-K:t}; G(\psi_{t-K:t})], \\ &\forall \psi_{t-K:t} \in \check{\Gamma}_{t-K:t|\Theta_t} \end{aligned} \tag{15}$$

$$\begin{aligned} J_t^*(x_{t-K}, \xi_{t-K:t}, \psi_{t-K:t-1}) &= \min_{\psi_t \in \check{\Gamma}_{t|\Theta_t}} J_t(x_{t-K}, \xi_{t-K:t}, \psi_{t-K:t-1}, \psi_t), \\ &\forall \psi_{t-K:t-1} \in \check{\Gamma}_{t-K:t-1|\Theta_t} \end{aligned} \tag{16}$$

Let $\gamma_{0:K-1} \in \check{\Gamma}_{0:K-1}$ be given optimal control laws. If control laws $\gamma_{K:T-1} \in \check{\Gamma}_{K:T-1}$ satisfy

$$J_t(x_{t-K}, \xi_{t-K:t}, \gamma_{t-K:t-1|\theta_t}, \gamma_{t|\theta_t}) = J_t^*(x_{t-K}, \xi_{t-K:t}, \gamma_{t-K:t-1|\theta_t}) \tag{17}$$

for all $\theta_t = (\lambda_{0:t-K-1}, \xi_{0:t}, x_{0:t-K}) \in \Theta_t$ and $t = K, \cdots, T-1$, then $\gamma_{K:T-1}$ are optimal.

Proof: See reference [12]. □

Again, we can interpret the above equations as the result of applying dynamic programming to the following centralized stochastic control problem. Let the state for the centralized problem be

$$x_t'' = (x_{t-K}, \xi_{t-K:t}, \psi_{t-K:t-1}) \tag{18}$$

$$\in X_{t-K} \times \Xi_{t-K:t} \times \check{\Gamma}_{t-K:t-1|\Theta_t}, \tag{19}$$

let the input be $u_t'' = \psi_t \in \check{\Gamma}_{t|\Theta_t}$, and let the disturbance be

$$w_t'' = (x_{t-K+1}, \xi_{t-K+1:t+1}, \lambda_{t-K}) \tag{20}$$

$$\in X_{t-K+1} \times \Xi_{t-K+1:t+1} \times \Lambda_{t-K}. \tag{21}$$

Define the transition function f_t'' by

$$f_t''(x_t'', u_t'', w_t'') = (x_{t-K+1}, \xi_{t-K+1:t+1}, \psi_{t-K+1:t|\lambda_{t-K}}), \tag{22}$$

and define the cost function g_t'' by

$$g_t''(x_{t+1}'', u_t'') = E[h_t(x_{t+1}, u_t)|x_{t-K+1}, \xi_{t-K+1:t+1}; G(\psi_{t-K+1:t|\lambda_{t-K}})]. \tag{23}$$

Again, conditions under which a stationary optimal undiscounted infinite horizon control law exists for this centralized problem are given in [10].

Note that the size of the state space for this centralized stochastic control problem is now

$$|X_{t-K} \times \Xi_{t-K:t} \times \check{\Gamma}_{t-K:t-1|\Theta_t}| = |X_{t-K}| \cdot |\Xi_{t-K:t}| \cdot \prod_{m=1}^{M} |U_t^m|^{|\Lambda_{t-K:t-1}^m|}, \tag{24}$$

and that $|\Xi_{t-K:t}|$ is not in the exponent of $|U_t^m|$ as it would be had we used the method in Section 4.1. However, the growth rate of the state space remains doubly exponential in K since $|\Lambda_{t-K:t-1}^m|$ grows exponentially with K. Nevertheless, even for small problems, the efficiency of the above algorithm allows an enormous reduction in required computation when compared to the algorithm in Section 4.1.

5 Multiple Access Broadcast Channel

In this section, we focus on the canonical MABC and formulate the problem of designing protocols as a decentralized control problem with no information sharing.

The multiple access broadcast channel is defined as follows. Time is discrete and is known by all users; at the start of each time slot, each of M users makes a decision regarding whether to transmit or not transmit; if two or more users in the system transmit a packet at the same time, then a collision results, and no packet is successfully received; if only a single user transmits a packet, a packet is received successfully, while if no user transmits, the channel is wasted; after every attempt at using the channel, all users are immediately informed as to whether a collision (two or more users transmit), success (one user transmits), or idle (no users transmit) occurred; packets arrive during each time slot with probability p to each user, and each user has no buffer (sometimes called a single buffer), i.e., he can hold only one packet waiting to be transmitted. All of the above are common and useful assumptions that model the essence of the multiple access problem. The goal is to choose protocols so that the probability of a successfully received packet in each slot, i.e., the throughput, is maximized.

We introduce notation to facilitate discussion of the MABC. Let q_t^{m-} be the number of packets in the mth user's buffer (either 1 or 0) prior to the arrivals for that time slot; let q_t^{m+} be the number of packets in the mth users' buffer (either 1 or 0) after the arrivals for that slot. Let b_t be the feedback after the tth slot. Let a_t^m be the number of arrivals (either 1 or 0) to user m during the tth slot. The control applied by the mth user is denoted u_t^m and specifies whether the user will ($u_t^m = 1$) or will not ($u_t^m = 0$) transmit a packet in his buffer. The number of packets transmitted by user m (although not necessarily received by the receiver) is s_t^m (either 1 or 0).

Denoting the Boolean logic operators "and" and "or" by $\wedge$ and $\vee$, respectively, we describe the operation of the channel as follows:

$\vdots$

$t.1$: The pre-arrival buffer state is q_t^{m-}.

$t.2$: Arrivals a_t^m to each user occur independently with probability p.

$t.3$: Post-arrival queue state is $q_t^{m+} = a_t^m \vee q_t^{m-}$.

$t.4$: Number of packets transmitted is $s_t^m = q_t^{m+} \wedge u_t^m$.

$t.5$: Feedback is $b_t = \min\{\sum_{m=1}^{M} s_t^m, 2\}$.

$(t+1).1$: The next pre-arrival buffer state is

$$q_{t+1}^{m-} = \begin{cases} q_t^{m+} - s_t^m & \text{if } b_t = 1 \\ q_t^{m+} & \text{otherwise} \end{cases}$$

$\vdots$

By assumption, each user may use the history of feedback broadcasts $b_{0:t-1}$ and its history of local post-arrival buffer states $q_{0:t}^{m+}$ to decide its control input u_t^m.

Let us now put the canonical MABC design problem in the decentralized control framework of Section 2. Let the state be $x_t = (q_t^-, b_{t-1})$. Let the primitive random variables be $v_t = a_t$ and $w_t^m = a_t^m$. Then the state transition function f_t is given by $x_{t+1} = f_t(x_t, u_t, v_t) = (\phi(q_t^-, u_t, a_t), \beta(q_t^-, u_t, a_t))$, where the functions ϕ and β are defined as follows:

$$\beta(q_t^-, u_t, a_t) = \min\{\sum_{m=1}^{M}((q_t^{m-} \vee a_t^m) \wedge u_t^m), 2\} \qquad (25)$$

$$\phi(q_t^-, u_t, a_t) = \begin{cases} (q_t^- \vee a_t) - ((q_t^- \vee a_t) \wedge u_t) & \text{if } \beta(q_t^-, u_t, a_t) = 1 \\ q_t^- \vee a_t & \text{otherwise} \end{cases}, \qquad (26)$$

where binary operations on tuples are performed element by element. The function g_t^m relating observations to the state is defined as $y_t^m = (\lambda_t^m, \xi_t) = g_t^m(x_t, a_t^m) = (q_t^{m-} \vee a_t^m, b_{t-1})$, and the cost function h_t is defined by

$$h_t(x_{t+1}, u_t) = \begin{cases} -1 & \text{if } b_t = 1 \\ 0 & \text{otherwise} \end{cases}, \qquad (27)$$

so that minimizing the expected cost will maximize the probability of successful packet reception in a slot. As mentioned earlier, the control input u_t^m must be a function of $b_{0:t-1}$ and $q_{0:t}^{m+}$, i.e., it must be a function of $y_{0:t}^m$, indicating that no sharing is allowed.

With these definitions, we have formulated the canonical MABC protocol design problem as a decentralized control problem with common information but with no sharing. The corresponding problem with K-step delay state from Section 4.2 can now be solved, and the resulting optimal throughput will provide an upper bound to the throughput of the canonical MABC.

5.1 Results for Two and Three Users

We use Theorem 2 with Howard's policy iteration algorithm [11] to find the bound described in Section 4.2. The tightness of the bound increases as the delay K increases, so it is desirable to compute the bound for the largest K possible. For reference, with respect to the decentralized control formulation of the MABC given in Section 5, the OSDS problem considered by Grizzle et al. [2] and Paradis [3] is equivalent to a K-step delayed state problem with $K = 0$. For two users, we choose the delay to be $K = 1$, while for three users, we choose a delay of $K = 2$. These choices were made because for two users, with $K = 1$, the bound meets the performance of the Hluchyj-Gallager protocol, while for three users, computation prohibits choosing $K > 2$.

Before using the policy iteration algorithm, we eliminate all self-contradictory states. For example, the state with sublaws that never transmit cannot have a success or collision feedback associated with it. The remaining states meet a sufficient condition (the "weak accessibility condition" given in [10]) that guarantees that an optimal stationary infinite-horizon solution exists.

p	*Probability of Success*		Ratio (≤ 1)
	H-G Protocol	$K = 2$ Bound	
.05	.1486	.1486	1.000
.10	.2881	.2882	.9996
.15	.4099	.4103	.9989
.20	.5081	.5101	.9961
.25	.5905	.5926	.9965
.28	.6301	.6327	.9959
.29	.6421	.6446	.9961
.30	.6570	.6570	1.000
.40	.7840	.7840	1.000
.50	.8750	.8750	1.000

Table 1: Three-user case: bounds on probability of success with $K = 2$ compared to probability of success of Hluchyj and Gallager's optimized window protocol (H-G protocol).

In the two-user case, numerical calculation of the bound with $K = 1$ shows that Hluchyj and Gallager's optimized window protocol achieves the bound for $K = 1$ to within machine precision. In the three-user case, numerical calculations of the bound for $K = 2$ show that Hluchyj and Gallager's optimized window protocol approaches the bound for $p < .3$ and meets the bound for $p > .3$. Note that for $p > .2891$, the Hluchyj-Gallager protocol is the same as time-division multiple access (TDMA) [4]. Because much computation is required for the $K = 2$ bound, the bound is only computed for the values of p in Table 1. For these values of p, the performance of the optimized window protocol is at least 99.59% of the bound, suggesting near optimality. Numerical calculation of the bound for three users with $K = 1$ show that for $p > .5$, the Hluchyj-Gallager (TDMA) protocol is optimal. Thus, we conclude that for $p > .3$, the Hluchyj-Gallager (TDMA) protocol is optimal.

6 Discussion and Conclusions

We have presented a bounding technique for decentralized control problems with no sharing of information and showed how complexity of the bound calculation can be reduced in the special case of common information. The complexity reduction has allowed us to apply the bound to the canonical MABC design problem for two and three users. We have developed results that show that the Hluchyj and Gallager's optimized window protocol is effectively optimal in the two- and three- user cases.

The results open several avenues for further inquiry. That the Hluchyj-Gallager optimized window protocol is at least nearly optimal for two and three users suggests that this protocol may be nearly optimal for more than three users; it may be fruitful to try to show this analytically. Alternatively, since the optimized window protocol uses only the common information [4], i.e., the broadcast feedback, it may be possible instead to show that no performance loss results from restricting protocols to the class of protocols that use only common information. Another task is to determine p such that TDMA is the optimal protocol for $M > 3$ users. It is known that TDMA is optimal as $p \rightarrow 1$, but our results suggest that TDMA is optimal for values of p that are substantially smaller than 1.

References

[1] D. Bertsekas and R. Gallager, *Data Networks*. Prentice Hall, 1992.

[2] J. W. Grizzle, S. I. Marcus, and K. Hsu, "Decentralized control of a multiaccess broadcast network," in *Proceedings of the 20th IEEE Conference on Decision and Control*, pp. 390–391, 1981.

[3] A. Paradis, "Application optimal control to the multiple access channel," Master's thesis, Massachusetts Institute of Technology, 1981.

[4] M. G. Hluchyj and R. G. Gallager, "Multiaccess of a slotted channel by finitely many users," in *Proc. National Telecommunications Conference*, pp. D4.2.1–D4.2.7, 1981.

[5] Z. Rosberg, "Optimal decentralized control in a multiaccess channel with partial information," *IEEE Transactions on Automatic Control*, vol. 28, pp. 187–193, February 1983.

[6] F. C. Schoute, "Decentralized control in packet switched satellite communication," *IEEE Transactions on Automatic Control*, vol. 23, pp. 362–371, April 1979.

[7] P. Varaiya and J. Walrand, "Decentralized control in packet switched satellite communication," *IEEE Transactions on Automatic Control*, vol. 24, pp. 794–796, October 1979.

[8] H. S. Witsenhausen, "Separation of estimation and control for discrete time systems," *Proceedings of the IEEE*, vol. 59, pp. 1557–1566, November 1971.

[9] M. Aicardi, F. Davoli, and R. Minciardi, "Decentralized optimal control of markov chains with a common past information set," *IEEE Transactions on Automatic Control*, vol. 32, pp. 1028–1031, November 1987.

[10] D. Bertsekas, *Dynamic Programming: Deterministic and Stochastic Models*. Prentice Hall, 1987.

[11] R. Howard, *Dynamic Programming and Markov Processes*. The M.I.T. Press, 1960.

[12] J. Ooi, Ph.D Thesis, Massachusetts Institute of Technology, Cambridge, MA, in preparation.