

This PDF is a selection from an out-of-print volume from the National Bureau of Economic Research

Volume Title: Annals of Economic and Social Measurement, Volume 4, number 2

Volume Author/Editor: Sanford V. Berg, ed.

Volume Publisher: NBER

Volume URL: <http://www.nber.org/books/aesm75-2>

Publication Date: 1975

Chapter Title: Survey of Decentralized Control Methods

Chapter Author: Michael Athans

Chapter URL: <http://www.nber.org/chapters/c12708>

Chapter pages in book: (p. 345 - 355)

## THEORY AND APPLICATIONS SURVEY OF DECENTRALIZED CONTROL METHODS\*

BY MICHAEL ATHANS

*The purpose of this paper is to present a nonmathematical overview of trends in the general area of decentralized control strategies, which are suitable for hierarchical systems. Special emphasis is given to pitfalls in the mathematical problem formulation. In addition, recent results that appear promising are indicated.*

### 1. INTRODUCTION

The purpose of this paper is to present an overview of the types of problems that are being considered by control theorists in the area of dynamic large scale systems with emphasis on decentralized control strategies. A similar paper by Varaiya [1] indicated the interplay between static notions drawn from the mathematical economics, management, and programming areas and the attempts by control theorists to extend the static notions into the stochastic dynamic case. In this paper we shall not elaborate upon the dynamic or team aspects of large scale systems. Rather we shall concentrate upon approaches that deal directly with decentralized decision making for large scale systems.

Although a survey paper, the number of surveyed results is relatively small. This is due to the fact that there is still not a unified theory for decentralized control. What is available is set of individual contributions that point out both "blind alleys" as well as potentially fruitful approaches.

What we shall attempt to point out is that future advances in decentralized system theory are intimately connected with advances in the so-called *stochastic control problem with non-classical information pattern*. To appreciate the difference of this problem relative to the *classical stochastic control problem* it is useful to briefly summarize the basic assumptions and mathematical tools associated with the latter. This is done in Section 2 of this paper. Section 3 is concerned with certain pitfalls that arise when one attempts to impose a decentralized structure right at the start, but the mathematics "wipes out" the original intent. Hence, one can draw certain conclusions about the proper mathematical formulation of decentralized control problems. Section 4 surveys some research (primarily carried out by the author and his students) that attempts to circumvent some of the pitfalls discussed in Section 3. Section 5 presents some conclusions about future research.

### 2. THE CLASSICAL STOCHASTIC CONTROL PROBLEM

In this section we review in an informal way the classical stochastic control problem or the problem of stochastic control with classical information structure.

\* This research has been conducted in the Decision and Control Sciences Group of the MIT Electronic Systems Laboratory with support from AFOSR grant 72-2273, NASA/AMES grant NGL-22-009-124, and NSF grant GK-41647.

Our main purpose is to indicate the types of assumptions that one makes in this class of problems, the nature of the mathematical tools available, and the general structure of the solution. This overview is necessary so that one can see that the solution to the classical stochastic control problem leads to a completely centralized system.

From a technical point of view, the best survey of the issues associated with classical stochastic control is, in the opinion of the author, the paper by Witsenhausen [2]. The classical stochastic control problem is the subject of many standard texts and monographs (see, for example, references [3] to [7]).

Figure 2.1 shows an abstract block diagram of a centralized control system. One deals with a stochastic dynamical system (usually described by a set of stochastic difference or differential equations) and one makes possibly noisy measurements of certain of its variables. The time evolution of the system variables is influenced by stochastic disturbances and decisions (or control) variables generated in closed-loop feedback form by a *single* controller or decision maker.




Figure 2.1

The assumptions that are made in the classical stochastic control problem are as follows.

- A2.1 There is a single controller or decision maker.
- A2.2 The controller knows the mathematical form of the system dynamics (i.e. the stochastic differential or difference equations).
- A2.3 The controller knows the relationship of the measurements to the system variables.
- A2.4 The controller knows a priori the probability densities of all underlying stochastic variables (i.e. exogenous disturbances, uncertain system parameters, measurement errors, etc.).
- A2.5 The controller wishes to minimize a well defined scalar deterministic cost functional.
- A2.6 The controller at any time  $t$  has *instant and perfect recall* of
  - (1) all past applied inputs or decisions.
  - (2) all past and present measurements.

Under these assumptions classical stochastic control provides a well defined rule which translates all information available to the controller at time  $t$  (i.e., the contents of assumptions A2.1 to A2.6) to an optimal control or decision at time  $t$ .

From a technical point of view

- (a) the state variable (causal) description of the dynamical system.

(b) the use of a Bayesian rule for dealing with the stochastic elements.

(c) the use of stochastic dynamic programming.

blend well together to yield the optimal stochastic control in a relatively straightforward *conceptual* manner; actual calculations are in general very complex with the exception of the so-called Linear-Quadratic-Gaussian (LQG) problem [4]-[8].

### 3. DECENTRALIZED CONTROL: PITFALLS

In this section we shall present, once more in an informal way, the types of issues that arise when the basic assumptions of Section 2 associated with classical centralized stochastic control are modified.

There are several ways that one can depart from the basic framework of classical stochastic control. In this paper we shall adopt the viewpoint of examining the issues when we wish to analyze some sort of hierarchical multilevel decentralized system [9].

One does not have to examine a complex hierarchical structure to understand the issues associated with decentralized control. Figure 3.1 presents the simplest possible case involving a two-level structure. We shall elaborate upon the structure implied in Figure 3.1 to point out its general characteristics.

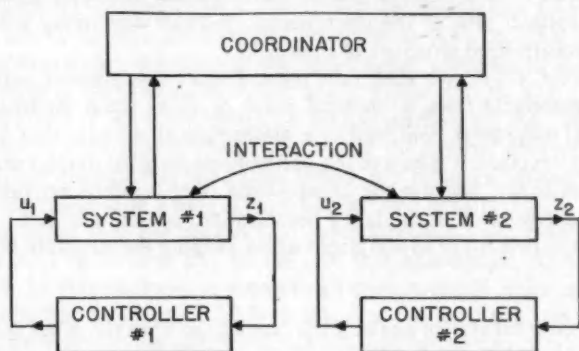


Figure 3.1

*Case 3.1.* Imagine for the time being that the "interaction" channel and the "coordinator" did not appear in Figure 3.1. Under this assumption we are left with two "uncoupled" dynamical systems. If we adopt the framework of Section 2 we can postulate that each controller solves a classical stochastic control system.

*Case 3.2.* Next, let us still leave the "coordinator" out of Figure 3.1, but restore the "interaction" channel. What we mean by "interaction" is that certain decision and/or state and/or output variables of each system influence the dynamic evolution of the other system. If this interaction is "weak", then it is possible for both systems to operate non-optimally, but still satisfactorily, without altering the basic control strategy of Case 3.1, because the interactions are viewed as exogenous unknown disturbances, and the inherent use of feedback tends to make the overall system response somewhat insensitive to weak unmodelled disturbances.

This situation, namely with weak interaction and the absence of coordination, has been analyzed by Chong, Kwong, and Athans in references [10]–[12]. This research attempted to replace the weak interaction disturbances, which are actually correlated in time, with equivalent “fake white noise” inputs, which are uncorrelated in time, and evaluate techniques by which, in the LQG context, the covariance of the “fake white noise” can be selected.

*Case 3.3.* If the dynamic interaction between the systems is not negligible, then severe deterioration in the performance of each system in Figure 3.1 can be expected. To “cure” this performance degradation one introduces the “coordinator” in Figure 3.1. Intuitively, in any physical large scale system, the role of the coordinator is to receive some sort of information from the local subsystems and make some decisions which improve the performance as compared to that under Case 3.2.

The crucial question then is to make precise the role of the coordinator, as a function of postulated strategies for the lower level subsystems. It is possible, in any given physical situation, to specify the task of the coordinator in a reasonably good, but *ad hoc* way, so that the overall system performance is satisfactory for the specific application. However, what lies in the heart of decentralized control theory research is to formulate precise mathematical problems whose solution defines the *optimal task* of the coordinator, without destroying the intuitively appealing decentralized structure of Figure 3.1.

*Some Blind Alleys.* We shall now present some assumptions which, though seemingly reasonable from a physical point of view, when incorporated in a mathematical framework, will lead to a mathematical solution that destroys the decentralized structure. It is easy to see, without resorting to complex mathematics, what the pitfalls are, because the appreciation of the pitfalls provides valuable knowledge on how *not* to formulate a decentralized control problem.

We start with a list of assumptions again keeping the structure of Figure 3.1 in mind.

- A3.1 Each local system neglects the interaction from the other system.
- A3.2 Under A3.1 each local controller knows the dynamics and probabilistic information associated with his own system, has his own performance index, and has perfect recall of his own past measurements and controls. It follows (see Section 2) that each local system can solve its own well formulated classical stochastic control problem, and we assume that each local system applies in real time the optimal stochastic control obtained under these assumptions.
- A3.3 The coordinator knows the dynamics of both subsystems, including the interaction, as well as all prior probabilistic information available to each local subsystem.
- A3.4 The coordinator's cost functional (performance index, utility function) is a well defined function of the cost functionals of each local subsystem (e.g., a weighted sum).
- A3.5 At each instant of time, the coordinator can apply a dynamic control to each local system of the same nature as the local control.
- A3.6 At each instant of time, each local subsystem transmits *instantly* and *without error* its

(a) measurements.

(b) controls.

to the coordinator; furthermore, the coordinator has perfect recall.

The key question is then: *Under assumptions A3.1 to A3.6 what is the optimal decision rule for the coordinator?*

The answer is exceedingly simple. Under assumptions A3.3 to A3.6 the coordinator has a classical stochastic control problem for the entire system. Hence, as far as the coordinator is concerned he must solve the overall optimal stochastic control problem (see Section 2), and his optimal strategy is to

(a) cancel the locally computed controls (see assumption A3.2).

(b) substitute the global optimal controls.

Thus, the essential decentralized nature of the problem is destroyed. This points out that even in the stochastic case, one cannot allow full knowledge of everything to the coordinator, because the mathematically optimal solution allows the coordinator to completely take over. This problem is even more serious if a complete deterministic framework for decentralized control is adopted as in reference [9].

*Control Sharing Strategies.* To avoid having the coordinator take over completely, one has to modify some of the assumptions A3.1 to A3.6 so that full knowledge of everything is denied to the coordinator. Needless to say, there are a myriad of ways in which one can modify assumptions A3.1 to A3.6 and attempt to analyze the role of the coordinator. In this subsection we shall examine one variation, because it has received some attention in the control literature, although not precisely in the context of this paper. Hence, our remarks represent a reinterpretation of the research of Aoki [13] and Sandell and Athans [14].

One can argue that the reason that the coordinator can take over under assumptions A3.1 to A3.6 is due to the fact that assumption A3.6 allows the coordinator to have instantaneous access to all measurements of the local systems. This allows him (see assumption A3.3) to calculate the local controls that will be used by the subsystems (see assumption A3.2), and so he can cancel them (see assumption A3.5). Hence, one may think that one way of preventing the coordinator from taking over is to deny to him the actual measurements of the local systems. Thus, we seek to modify assumption A3.6.

Assumption A3.6, however, not only deals with measurements but controls. We can adopt the philosophy "do not flood your boss with day-to-day occurrences, but let him know your day-to-day decisions." If we adopt this framework, we can replace assumption A3.6 with the following

A3.6(M) At each instant of time, each local subsystem transmits *instantly* and *without error* ONLY its controls, BUT NOT its measurements to the coordinator; furthermore, the coordinator has perfect recall.

One can then pose a mathematical problem under assumptions A3.1 to A3.5 and A3.6(M) in order to find the optimal decision rule for the coordinator.

The answer [13], [14], is both surprising and interesting. It turns out that

(a) the stochastic control problem for the coordinator is *not well defined*, in the sense that an optimal solution *does not exist*.



- (b) although an optimal solution does not exist, one can find  $\varepsilon$ -optimal solutions in the sense that one can approach the unattainable optimal solution arbitrarily closely.

The way by which these  $\varepsilon$ -optimal solutions are obtained is interesting, and instructive because they indicate once more on how *not* to formulate a decentralized control problem. We shall attempt to explain how this happens by a simple example.

In reference to Figure 3.1 let us suppose that the system operates in discrete time so that measurements and decisions are made at the values of the time index

$$t = 0, 1, 2, 3, \dots$$

Let  $z_1(t)$  denote the measurement and let  $u_1(t)$  denote the control of system # 1; for the sake of simplicity assume that both  $z_1(t)$  and  $u_1(t)$  are scalars.

Suppose that the following sequence of measurements has been made by system # 1

$$(3.1) \quad \left. \begin{array}{l} t = 0, \quad z(0) = 6 \\ t = 1, \quad z(1) = 7 \\ t = 2, \quad z(2) = 8 \\ t = 3, \quad z(3) = 9 \end{array} \right\}$$

Under assumption A3.2, system # 1 has a well defined rule for generating its own optimal control. Suppose that at the basis of the measurements (3.1) the optimal local control for system # 1 at time  $t = 3$  turned out to be

$$(3.2) \quad u_1^*(3) = 1.234$$

Under assumption A3.6(M) the control (3.2) can be transmitted instantly and without error to the coordinator. However, the nature of the  $\varepsilon$ -optimal solution indicates that system # 1 *should not transmit* the control (3.2) to the coordinator.

Rather, it should transmit, and apply to the system, a control of the following form

$$(3.3) \quad \hat{u}_1(3) = 1.234000000 \dots 00006789$$

The information conveyed to the coordinator when he receives the control (3.3) *without error* is very different from that contained in (3.2). Examination of Equations (3.1) to (3.3) indicates that the past measurements (6, 7, 8, 9) have been *coded* in  $\hat{u}_1(3)$ . From the front part of (3.3) the coordinator knows the control  $u_1^*(3)$  of Equation (3.2); from the tail end of (3.3) the coordinator knows exactly the past measurements of system # 1. The string of zeroes between the control (1.234) and the coded measurements (6789) is simply to guarantee that the application of  $\hat{u}_1(3)$ , rather than  $u_1^*(3)$ , to the system results in infinitesimal loss in system performance (i.e. the 000...006789 part of the control will be wiped out by the system uncertainty).

Hence, assumption A3.6(M) is not strong enough to prevent the coordinator from obtaining all the information that he needs in order to take over for all practical cases.

*Conclusions.* The above discussion points out that in stochastic decentralized problems *instantaneous error-free* transmissions of either both controls and measurements, or controls alone is *not* a realistic mathematical assumption, because this allows the coordinator to take over, and destroy the decentralized nature of the problem.

#### 4. DECENTRALIZED CONTROL: PROMISING AVENUES

In this section we discuss some recent results which appear to be useful toward building some of the elements of a theory for decentralized control. Once more the reader is referred to Varaiya [1] for additional concepts and discussion.

##### 4.1. Decentralization with Fixed Structure

Figure 4.1 depicts a specific decentralized structure which is somewhat different from the one discussed in Section 3. One is given an  $n$ -th order linear stochastic dynamic system, with two sets of measurements ( $z_1$  and  $z_2$ ) and two sets of controls ( $u_1$  and  $u_2$ ). It is decided *a priori* to select two dynamic controllers that generate stochastic controls in the manner illustrated by Figure 4.1. It is assumed that there is a single cost functional to be minimized, and one can view the job of the coordinator as defining the characteristics of the two controllers. This represents a variation on the dynamic team problem (see for example Ho and Chu [15], [16]).

Because of the non-classical information pattern and the general lack of knowledge for solving dynamic team problems, some additional assumptions have to be made in order to be able to solve the problem of designing the two controllers in Figure 4.1.

For the LQG continuous-time case Chong and Athans [10], [17] fixed the structure of each controller to be linear and of the same dimension as the order

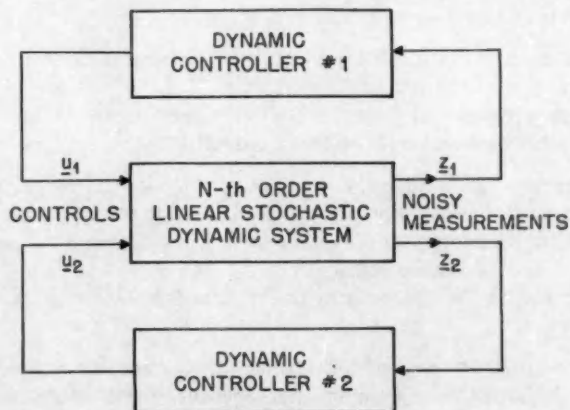


Figure 4.1



of the dynamic system which was controlled. Furthermore, each controller was constrained so that its internal Kalman-Bucy filter would produce unconditional zero mean estimates of the state, ignoring the actions of the other controller. The parameter matrices of each dynamic controller could then be globally optimized by solving a deterministic matrix optimal control problem, through the use of the matrix minimum principle [18]. The discrete-time version of this problem was considered by Carpenter [19].

There are two basic conclusions that can be drawn from the above studies (see also Varaiya [1]).

- (a) The off-line computational effort for solving such decentralized problems is greater than that required for the centralized case.
- (b) Even in the LQG context the *separation theorem* or *certainty equivalence principle* fails to hold.

#### 4.2. Periodic Coordination

The discussion in Section 3 indicated that for decentralized systems, in the context of Figure 3.1, one *cannot* provide the coordinator with instantaneous and error free transmission of the local subsystem measurements and/or decisions to the coordinator; otherwise, the coordinator takes over, and substitutes the globally optimal stochastic controls thus overriding the local controller's decisions.

One way of bypassing this problem is to assume that the coordinator is allowed to "interfere" only once in a while. This notion of periodic coordination has been considered by Chong and Athans [20]-[22].

To understand the intuitive notion of periodic coordination, suppose that the assumptions A3.1 to A3.4, given in Section 3, are still valid but assumptions A3.5 and A3.6 are replaced by the following (informal) one.

*Periodic Coordination Structure:* Suppose that the entire system operates in discrete time. For the sake of concreteness we assume that the basic time unit is a day. Then the basic system operation is as follows.

- (1) We assume that each lower-level system makes its measurements and generates its controls (decisions) *once a day*.
- (2) *Once a month*, all lower-level system measurements and controls are "mailed" without error to the coordinator.

The basic question is: *what is the job of the coordinator at the beginning of each month?* The mathematical approach adopted, and the results obtained, [20]-[22], have the following interpretation which we feel has certain intuitively appealing aspects, discussed informally below.

Once a month the coordinator has a three-fold task with respect to each lower-level system.

- (1) *Sets it straight:* in a technical context he corrects the estimates generated by the lower-level Kalman filters, because these estimates are in error due to the fact that each lower-level system neglects the interactions from the other lower-level systems.

- (2) *Changes its directives*: in the sense that new time paths for the lower-level controls are given in an open loop sense.
- (3) *Changes its incentives*: in the sense that additional terms are added to each lower-level system cost functional to compensate for the fact that the global cost functional is different from each lower-level cost functional.

The main advantage of these results is that the mathematical theory itself suggests the tasks that will have to be performed by the coordinator on a periodic basis.

The main *disadvantage* is that the coordinator still has to solve a very large scale stochastic optimization problem, although not as often as in the basic time frame of the lower-level. Although for certain applications this approach may be feasible, nonetheless it *lacks the capability of somehow aggregating the information flow from the lower-level systems to the coordinator*.

Nonetheless, because the theory itself suggests this mode of coordination (by changing directives and incentives) it provides strong motivation to postulate a specific framework for the operation of the lower level systems. A preliminary formulation along these lines can be found in a recent paper by Athans [23].

The notion of delaying the information exchanged between different portions of a hierarchical system is an intuitively appealing idea. Much more research is needed to understand its impact on decentralized control theory. However, the results of Sandell and Athans [14], in which it was shown that LQG problems with a unit-time step delay of information exchange admit a *linear optimal decision rule*, which can be calculated explicitly, appear to be promising as far as their applicability is concerned to decentralized control theory.

#### 4.3. Remarks

Most of the results surveyed attempt in one way or another to present to both the coordinator as well as the local subsystems a classical stochastic control problem. Although research along these lines is useful, there is a great lack of any definitive theory which deals directly with issues of aggregated information, decision making with partial information, or decision making with finite memory.

To adequately deal with these issues in the context of decentralized control much additional research is needed in the area of *stochastic control with non-classical information patterns*. The famous Witsenhausen counter-example [24], in which a simple LQG problem with nonclassical information pattern was shown to have a *nonlinear* optimal decision rule, points out the immense difficulties associated with this class of problems. Witsenhausen [25], [26] has continued his fundamental investigations in this class of problems, but their implications in the context of decentralized control theory remain largely unexplored. The work on finite-state finite-memory control of Sandell and Athans [27], [28] may be useful for aggregating the flow of information between the different levels of a hierarchy, and limiting the computational complexity available to the coordinator. In addition, the recent results surveyed by Ho [29] pertaining to approaches in information structures when many decision makers are involved, is of direct importance to decentralized control problems.

## 5. CONCLUSIONS

The main conclusions that one can draw are as follows:

- (1) Any purely deterministic approach to multilevel hierarchical dynamic systems is not apt to lead into realistic theories or designs.
- (2) The flow of measurements and decisions in a decentralized system should not be instantaneous and error-free.
- (3) Delays in information exchange in a decentralized system lead to reasonable approaches to decentralized control.
- (4) A mathematically precise notion of aggregating information is not as yet available.
- (5) Research in non-classical information structures is directly relevant to problems of decentralized control.

*MIT Electronic Systems Laboratory*

## REFERENCES

- [1] P. P. Varaiya, "Trends in the Theory of Decision-Making in Large Systems," *Annals of Economic and Social Measurement*, Vol. 1, No. 4, pp. 493-500, October 1972.
- [2] H. S. Witsenhausen, "Separation of Estimation and Control for Discrete Time Systems," *Proceedings IEEE*, Vol. 59, 1971.
- [3] R. Bellman, "Dynamic Programming," Princeton University Press, 1957.
- [4] M. Aoki, "Optimization of Stochastic Systems," Academic Press, 1967.
- [5] K. Hinderer, "Foundations of Non-stationary Dynamic Programming with Discrete Time Parameter," Springer-Verlag, Berlin, 1970.
- [6] K. Astrom, "Introduction to Stochastic Control Theory," Academic Press, 1970.
- [7] H. J. Kushner, "Introduction to Stochastic Control," Holt, Rinehart, and Winston, 1971.
- [8] M. Athans, "The Discrete Time Linear-Quadratic-Gaussian Stochastic Control Problem," *Annals of Economic and Social Measurement*, Vol. 2, pp. 449-91, October 1972.
- [9] M. D. Mesarovic, D. Macko, and Y. Takahara, "Theory of Hierarchical Multilevel Systems," Academic Press, 1970 (see also the review of this book in *IEEE Transactions on Automatic Control*, Vol. AC-17, pp. 280-81, 1972).
- [10] C. Y. Chong, "On the Stochastic Control of Coupled Linear Systems," S.M. Thesis, MIT, 1972.
- [11] R. H. Kwong, "On a Stochastic Control Method for Weakly Coupled Systems," S.M. Thesis, MIT, 1972.
- [12] R. H. Kwong, C. Y. Chong, and M. Athans, "On Stochastic Control System Design Methods for Weakly Coupled Large Scale Systems," *Proceedings 1972 Joint Automatic Control Conference*, pp. 220-29.
- [13] M. Aoki, "On Decentralized Linear Stochastic Control Problems with Quadratic Cost," *IEEE Transactions on Automatic Control*, Vol. AC-18, pp. 243-50, June 1973.
- [14] N. R. Sandell, Jr. and M. Athans, "Solution of Some Nonclassical LQG Stochastic Decision Problems," *IEEE Transactions on Automatic Control*, Vol. AC-19, pp. 108-116, April, 1974.
- [15] Y. C. Ho and K. C. Chu, "Team Decision Theory and Information Structures in Optimal Control Problems; Part I," *IEEE Transactions on Automatic Control*, Vol. AC-17, 1972.
- [16] K. C. Chu, "Team Decision Theory and Information Structures in Optimal Control Problems; Part II," *IEEE Transactions on Automatic Control*, Vol. AC-17, 1972.
- [17] C. Y. Chong and M. Athans, "On the Stochastic Control of Linear Systems with Different Information Sets," *IEEE Transactions on Automatic Control*, Vol. AC-16, pp. 423-30, October, 1971.
- [18] M. Athans, "The Matrix Minimum Principle," *Information and Control*, Vol. 11, pp. 592-606, November 1967.
- [19] C. H. Carpenter, Jr., "Some Control and Communications Aspects of Stochastic Systems," S.M. Thesis, MIT, 1972.
- [20] C. Y. Chong, "On the Decentralized Control of Large-Scale Systems," Ph.D. Thesis, MIT, 1972.
- [21] C. Y. Chong and M. Athans, "On the Periodic Coordination of Linear Stochastic Systems," July 1974 (to appear).

- [22] C. Y. Chong and M. Athans, "Hierarchical Decomposition for a Stochastic Optimization Problem," *Proceeding 1973 Princeton Conference on Information Sciences and Systems*, Princeton, N.J., March 1973.
- [23] M. Athans, "Theoretical and Practical Problems in the Design of Stochastic Large Scale Systems," *Proceedings Fourth Iranian Conference on Electrical Engineering*, Shiraz, Iran, May 1974.
- [24] H. S. Witsenhausen, "A Counterexample in Stochastic Optimal Control," *SIAM Journal on Control*, Vol. 6, 1968.
- [25] H. S. Witsenhausen, "On Information Structures, Feedback, and Causality," *SIAM Journal on Control*, Vol. 9, 1971.
- [26] H. S. Witsenhausen, "A Standard Form for Sequential Stochastic Control," *J. Mathematical System Theory*, Vol. 7, 1973.
- [27] N. R. Sandell, Jr. and M. Athans, "Control of Finite State Finite Memory Stochastic Systems," *Proceedings 1974 Princeton Conference on Information Sciences and Systems*, Princeton, N.J., March 1974.
- [28] N. R. Sandell, Jr., "Control of Finite-State, Finite-Memory Stochastic Systems," Ph.D. Thesis, MIT, May 1974.
- [29] Y. C. Ho, "Information Structures in Many-Person Optimization Problems," paper presented before the third NBER Stochastic Control Conference, Washington, D.C., May, 1974.