Multi-Agent Stochastic Control: Models Inspired from Quantum Physics

John S. Baras
Department of Electrical and Computer Engineering
and the Institute for Systems Research
University of Maryland
College Park, MD 20742, USA
email: baras@isr.umd.edu.

Abstract

In this paper we consider multi-agent stochastic optimization and control problems, with partial information. The agent can operate in a distributed and asynchronous fashion. We investigate new problems that arise out of the interaction between observations and control actions by the agent. We show that new non-classical and non-commutative probability models are needed in order to properly formulate such problems. The models we develop here are inspired by models developed for dynamical physics problems. We establish a series of fundamental results for the trade-off between information and control patterns in distributed stochastic control, detection and estimation.

1 Introduction

In stochastic control problems, the main objective is to achieve satisfactory performance of a system operating in an uncertain environment. Typically, performance is measured by the expected value of a performance criterion (or cost function). In classical stochastic control, there is one controller (control station, control agent) and one performance measure. The controller employs sensors to perform measurements on the system, collects and stores the resulting data (observations), which are subsequently processed in actuators to produce the decisions (inputs) that optimize the performance criterion. This process is customarily identified as employing feedback (more precisely, information feedback) from the system.

The mathematical formulation of this classical stochastic control problem is thought to be well understood to date [14, 15], but explicit solutions are known only for few special cases. The case where the information available to the controller is explicitly generated by some sort of noisy observation of the "state" (i.e., the case of so-called partially observable stochastic systems) is treated in [8, 14, 15, 16]. In classical stochastic control, there is no concern for the interaction between information (and its transmission) and control. Furthermore, there is no interaction between measurement process and system dynamics, and typically selection of measurements (or observations) is not part of the problem. If we can obtain an explicit solution, there is not much difficulty in implementing the optimal controller which does not depend explicitly on the observation model (in the sense that the latter is fixed). Available information is modelled by σ-algebras, the system's "state" by a vector valued stochastic process, but is not in general well defined and understood.

Serious complications arise, however, when one considers the control of a stochastic system by various controllers with different available information and possibly different criteria. Such problems go under the categories of non-classical information patterns, stochastic control of large systems, decentralized or hierarchical control, etc. Thus seemingly simple problems lead to major departures from classical stochastic control results (c.f., Witsenhausen's well known and often quoted counterexample [25]). It is fair to state that despite many worthwhile and enlightening contributions by several people [2, 3, 10, 23, 24], the major questions still remain unanswered. In our opinion, there are two fundamental problems whose resolution is widely recognized as key to further progress:

(a) *The interreaction between information and control.* Under this heading are such problems as communication between controllers via "signaling strategies", "information neighborhoods" for controllers, cost of information versus cost of control [2, 3, 23, 24, 26]. Despite the pioneering work of Witsenhausen [25, 26], who devel-

0-7803-7939-X/03/$17.00 © 2003 IEEE

oped a number of important formulations and results about the separation of the use of information (i.e., estimation) and control, it is the author's opinion that there does not exist to date an agreed upon and satisfactory formulation of the joint "optimization" problem in information flow and control. It is important to develop theories that treat control strategies and information patterns in a balanced manner.

(b) *The concept of "state" for such a system is not well understood,* although Witsenhausen gave results towards a resolution of this problem. The problem stems primarily from the fact that in multi-agent stochastic problems there need not be a preassigned total time order of actions. In fact, action times may depend on observations and controls by the same or other agents. This problem is also related to the availability (theoretically, despite the obvious computational complexity) of a dynamic programming algorithm for solution. It is important to develop local state descriptions. Global state descriptions are equivalent to centralized control. Local state models, and the associated local times, permit the consideration of asynchronous actions by various controllers.

A third major problem that may appear in stochastic control problems with many agents, which has not been emphasized to date, is centered around the possible *interactions between measurements by different agents and between system dynamics and measurements.* We shall see that these concepts are related to some of the difficulties encountered to date and are akin to very strong interaction between information and control. This is typically the case where one cannot prove existence of an optimal control law (or strategy or design). We pointed out in [7] that similar problems appear in communication problems with quantum mechanical signal and noise models. Pursuing further the similarities between these two different problems, we develop certain new formulations for the problems of interaction between information and control, and system dynamics and measurement inspired by the methodologies used in quantum communication theory [4, 9, 18]. We show that a "non-commutative" probability theory (in the sense the term is used in the axiomatic foundation of quantum mechanics [1, 18, 19, 21, 22]) is necessary for some of these problems. Motivation for this work comes from the following problems: control of networked systems, distributed and asynchronous cooperative control, sensor selection and scheduling in distributed state estimation schemes, efficient coordination of mobile wireless networks.

2 Multi-Agent Stochastic Control Problems

We are interested in decentralized control problems where one wishes to specify an optimal design, in Witsenhausen's terminology, but with the additional characteristic that there is some flexibility over the information pattern. By that, we mean that there are several alternatives for the information available to each controller, on which decisions have to be based. What can be said abstractly about the joint selection of information and control patterns? Obviously, we are not interested in an exhaustive search between all possible information patterns. We are also interested in problems where there is no strict preassigned order of action times to be followed by the various controllers (or agents). Thus in economic systems, each agent (which can be an individual or an organization) has ample choice of sets of data on which to base decisions (consider, for example, the various economic indicators or statistical data reduction results available to the public and the government). Furthermore, it does not appear that there exists a strict preassigned order of action times in economic systems. In systems with such properties, we encounter a new kind of difficulty. Mainly, since control actions by one controller affect the measurements (or observations) of another, there may very well exist situations where efforts by two agents (by choice of information and control) to obtain as accurate as possible values for two critical (for their actions) variables will be in conflict, resulting in the impossibility of such simultaneous accuracy. This certainly requires very strong *information-control interaction.*

Finally, we are interested in systems where the agents can anticipate certain control actions by other agents when they are informed of the type, and not the results, of measurements performed by these other agents. Such so-called anticipatory systems have been studied by others. Often it is possible to give a statistical description of such a system's reaction to measurement. Roughly, the system's "state" changes due to the measurement performed. It is quite interesting to note that such properties were used by Wigner to produce a well known (but rather easily resolved) paradox in quantum physics. Such phenomena can be seen again in economic systems where, for example, measurement of income levels for taxation may have adverse effects on productivity. In another setting, the knowledge by a driver that a traffic detector exists nearby may cause changes in his velocity in order to, for example, catch the green or yellow at an intersection.

Some immediate questions of interest are:

(i) What are the implications of such phenomena on the probabilistic models used in stochastic control?

(ii) Are there any assumptions that will permit a reasonable definition of a "state"?

(iii) How can we formulate optimization problems for such control systems?

(iv) Can we solve any?

The problems of interaction between information and control described above have been observed and discussed in the literature under such code words as "signalling strategies", "information value" versus "control value", "deadlock problems in distributed estimation", "optimal selection of measurement data upon which to base a decision".

In a stochastic system with several agents acting asynchronously, *action times* may depend on observations and controls by the same or other agents; that is, there is "asynchrony". So the usually made assumption about preassigned time order of actions is not a realistic assumption. It is clear that once we accept a central clock according to which each controller times his actions, we have really accepted a centralized control algorithm. So in this paper, we would like to abandon the concept of global time and global state and develop systematic means of generating "local state" models. An important point related with this issue is that these "local states" (or "local models") must be supported by *locally collectable data*. To elaborate further on the concept of interaction between measurements performed by different agents, it is easy to see that when two agents operate in a system, they can often be in "conflict", meaning here that it can be impossible to coordinate their actions so that each obtains high quality observational data *simultaneously*.

We show in this paper that it is possible to develop a general mathematical framework that has several of these desirable features built into its basic mathematical constructs. Furthermore, it allows the treatment of information flow and control actions (strategies) as "dual" concepts. For example, an appropriate information pattern may facilitate the control task and vice versa. We need then to develop methodologies that establish this duality in a prescriptive accurate sense. This is a clear prerequisite for balanced treatment of information patterns and control strategies. Furthermore, once this "duality" is established in a mathematical sense, one can proceed to develop an optimization framework which makes feasible the "evaluation" of proposed control and information flow strategies jointly.

Given our main thrust just described, it was natural to focus on understanding information collection and information flow in multi-agent stochastic systems. This led us quickly to a fundamental problem: *Are the currently used probability models adequately equipped to analyze these critical issues?* In plain language, this translates: Do we currently have adequate representations of data collected by the sensoring elements in a multi-agent stochastic system? Here we noted certain similarities between some of the critical problems mentioned above and the axiomatic development of quantum mechanics? This led us to certain postulates about structure of the probability models in such systems. We would like to emphasize here, however, that there is one fundamental difference. Namely, the non-commutative models in quantum mechanics had as objective the description of the *passive* interaction between measurement and "local states". What we have in stochastic multi-agent systems is actually an *active* version of these models.

2.1 The Need for a Non-Commutative Structure

The heuristic discussion of the previous section suggests that a careful examination of the information available for decisions and a precise description of its place in the mathematical formulation are necessary for further understanding of these problems. Witsenhausen proposed a model for doing this. According to that model, the system's dynamics are determined by the realizations of the noise variables and all control variables. The performance measure can also be expressed as a function of these same variables (by solving the system's equations). Finally, for each decision, the data available for that decision are functions of these same variables and define a σ-field in an appropriate space. The complete specification of the control problem consists of the specification of the performance measure and these σ-fields.

Following Witsenhausen, let (Ω, B, P) be the probability space for the intrinsic random variables of the system and suppose we have a finite set A of agents acting on the system. In this formulation, a controller acting at two different times will be considered as two different agents. Let (Υ_a, F_a) be the measurable space in which agent a selects his control action u_a.

Considering product sets and product σ-fields, to a subset B of agents we associate the set

$$H_B = \Omega \times \Pi_{a\epsilon B}\ \Upsilon_a$$

and the σ-field

$$F_B = B \times \Pi_{a\epsilon B}\ F_a$$

All σ-fields are considered as subfields of F_a using the natural projection of H_A onto H_B. The information available to agent a is characterized by a subfield T_a of F_A. The possible control laws for agent a are the functions $\gamma_a : H_a \to \Upsilon_a$ which are measurable from T_a to F_a; they form the set Γ_a. Then the control for the subset of agents $B \subset A$ can be chosen from $\Gamma_B = \Pi_{a \epsilon B} \Gamma_a$ and the whole design in Γ_A. Witsenhausen then goes on to characterize various types of *information patterns* (i.e., the collection of T_a), such as causal, classical, quasi-classical, without self-information, etc.

Our first point is that in a well-posed multi-agent stochastic control problem, the agents will make inferences about system variables based on their own information fields T_a (appropriately pulled back in B). That is, the agents will compute conditional expectations (and/or probabilities) either implicitly or explicitly. To us, one important difference between classical and nonclassical information patterns is that these operations *commute* in a classical pattern and *do not commute* in a nonclassical pattern. Indeed, according to [26], an information pattern is *classical* if it is *sequential* (i.e., there is an ordering $(a_1, a_2, \cdots, a_n)$ of A such that $T_{a_k} \subset F_{\{a_1,\cdots,a_{k-1}\}}$ for $1 \leq k \leq n$ and $T_{a_k} \subset F_\phi, T_{a_{k-1}} \subset T_{a_k}$ for $k = 2, \cdots, n$. Then the commutativity of conditional expectations is just a consequence of the smoothing property of conditional expectations [15]. The statement for the nonclassical patterns is also obvious. The non-commutative modifier of our title refers to the corresponding property of conditional expectations. We adopt the point of view that the σ-fields T_a forming the information pattern are generated (or correspond to) measurements (for observations) performed on the system. The natural question is then: Does there exist a different model than the one described above which can describe statistically the events (observations) associated with a multi-agent stochastic control problem, including an intrinsically non-commutative conditional expectation operation? We shall see in a later section that there is an affirmative answer and that a prime example of such probability models is the von-Neumann model of quantum mechanics [18, 19] and extensions.

Second, we firmly believe that one of the deficiencies in the current development of decentralized stochastic control is that the information pattern is assumed fixed and given apriori. Little effort has been directed towards "optimal" selection of information pattern. The articles that address such problems (see references in [2, 3, 24]) use a formulation that represents the choice of information pattern as an optimization problem over a finite set of parameters. The usual method of attacking such problems is to solve a parameterized family of stochastic control problems, then to select the parameters which result in better performance, and thus choose the corresponding information pattern. This approach has been followed, for example, in problems of optimizing sensor locations in distributed systems.

2.2 Representative Problems

We start from problems that are static. Basically, these are distributed detection and estimation problems, with many sensors. We discuss extensions to simple dynamics, still with static information flow (i.e., no measurement-dynamics interaction). These problems basically have event driven transitions as their dynamics. We will assume simple control cost structure, information (computation) cost structure, in order to facilitate the formulation of the joint problem in an optimization framework. We shall comment on extensions and formulations of more complex system problems, but we shall not analyze them. In particular, we shall discuss duality of control and information patterns and hierarchies.

A concrete application is the so-called *Distributed M-ary Detection* problem [4, 7, 18]. There is a finite set of hypotheses $H_1, H_2, \cdots, H_M$, each of course affecting in some way the "system". We deliberately do not want to specify what we mean by system. There are also N agents; let us denote them by $A_1, \cdots, A_N$. The agents collect "data". They can store, make inference, communicate data or communicate inferences between themselves according to certain rules.

The next question is: What exactly is meant by the influence of $H_1, H_2, \cdots, H_M$ or the "system"? Here, we can pose two kinds of problems whose treatment differs by *at least* an order of magnitude in the mathematical sophistication required. In the *first category of problems*, we assume existence of a global system state. In the *second category of problems*, the hypotheses $H_1, H_2, \cdots, H_M$ affect directly the statistics of the data observed by each agent. Each agent can interpret this effect through a *local state model*. We shall describe results for this second problem.

One important concern to us is the following: Do not fix information pattern apriori and then optimize decision rules. Rather, determine the information flow structure that should be employed given some minimal but realistic constraints.

A rather challenging question is: How do we formulate such problems so dynamic information exchange occurs as a result of it? We expect the resultant decision strategies to force agents to act *asynchronously*. For detailed descriptions of the results presented here we refer to [5].

3 Non-commutative Probability Models: Algebraic Structures

In this section, we provide an axiomatic foundation of several (increasing in complexity) non-commutative probability models that can be used to represent the data observed by agents in a large scale system.

We want to emphasize that our approach begins from fundamental requirements that will later enable proper formulation of distributed communication problems. The primary benefit from such works is that the various physical assumptions, engineering intuition, etc., are built into the algebra of the model, and this process mechanizes the subsequent derivations.

In quantum mechanics, such an approach has been originated by Birkoff and von Neumann. The starting point was the structure of *propositions*, that is, yes-no measurements [19]. Due to its similarity to a logical system, the set of propositions is called *quantum logic*. This set can be easily given the structure of a lattice, and the basic question addressed by physicists was: Is there any set of phenomenological axioms that can allow one to identify the quantum logic with the set of orthogonal projections on a complex Hilbert space H ? Details of such theories can be found in [18-22]. There is one major disadvantage in this school of thought, however, as pointed out by Pool [21, 22]. It tacitly assumes the structure of quantum logic is sufficient in itself to determine the mathematical formalism which should be employed in the quantum theory. This is not true, however: It is a fact that quantum mechanics is used not so much to reproduce the logical properties of simple yes-no experiments (and sequences of those) but rather to compute transition probabilities, cross sections, etc. Therefore, the probabilistic aspects must be unified with the logical aspects. The interested reader is referred to [19-22] for details in the developments of the axiomatic foundation.

Following the fundamental suggestions made by Baras [7], we have developed similar, in principle, models for the data collected by agents in a distributed stochastic system.

Beginning from fundamental requirements on the data and propositions that appear in distributed sensor systems we first develop some algebraic structures. First, a *simple proposition or simple event* is a proposition that can admit a yes (usually assigned the binary value 1) or no (assigned the value 0) answer only, regarding their validity. Their validity can be verified (ascertained) by some combination of the data (measurements, experiments) performed by the various agents. We denote by E the set of simple events (or propositions).

It is important to note that "ambiguous" events (i.e., requiring probability assignments for their validity) are not simple events. There is a set of natural axioms we impose on E, supported by databases operating in multi-sensor, multi-agent distributed stochastic systems. We also have two important operations of *implication* (denoted by $\leq$) and *orthoemplementation* (denoted by $'$). We then have:

Theorem 1: *The set of simple propositions, in a distributed, asynchronous stochastic system is an orthomodular* σ *-orthoposet.*

This however is not a complete characterization. The reason is that the data bases of such a system cannot just be characterized by the *logic* (or logical structure) of the simple events that can be verified by the agents. The structure of the "logic" by itself is not sufficient to determine the mathematical formalism which should be employed. It is a fact that a mathematical theory of distributed detection (or estimation) (and more generally, of stochastic multi-agent systems) is used not so much to reproduce the logical properties of simple yes-no experiments performed and answerable by the agents, but rather to compute statistics of "system state" transitions and of outcomes of more complicated experiments (measurements). Therefore, we next *unify the probabilistic aspects with the logical aspects.*

Thus we are led to consider *event-state structures.* We think of states as the set of all possible (or just pertinent to the problem) configurations of the stochastic system. We want to emphasize that *we do not assume a memory interpretation for states.* This is done on purpose, as we do not want to have global causality due to the implication of existence of global time connected with it (i.e., synchrony), which in turn implies a centralized operation. Recall also that we want to incorporate or allow anticipatory agents. First, we consider global states and arrive at a more satisfactory model. This cannot be done without the introduction of probabilities. So we are led to consider the probability function $P : E \times S \rightarrow [0,1]$ with $p \epsilon E, a \epsilon S$:

$$P(p,a) = Pr\{p \text{ occurs when the global state is } a\}.$$

Thus we now have a triple (E, S, P). The details of our constructs are strongly motivated by certain quantum mechanical models [19-22], and are omitted. Under certain natural axioms we have established the following.

Theorem 2: *Given an event-state structure (E, S, P), satisfying certain axioms compatible with general principles of information flow and data bases in multi-agent systems, one can construct $(E, \leq, ')$, and S, such that*

(a) *$(E, \leq, ')$ is an orthomodular σ-orthoposet.*

(b) *$\hat{S}$ is a strongly order-determining σ-convex set of probability measures on $(E, \leq, ')$.*

(c) *$a \to \mu_a$ is a bijection of S onto $\hat{S}$.*

There is actually a converse, asserting that the above representation is "faithful", in the sense that it can generate the statistics on which it was based.

The set $\hat{S}$ is *dual* to the set S since it corresponds to a subset of the real valued functions on S. This is along well known mathematical methods of describing sets either by defining properties or by a set of functions on them. These constructions have quite interesting interpretations in the context of distributed systems. The states in $\hat{S}$ are basically *probability assignments* to simple events. Next, σ-convexity allows the construction of "mixture" states; for example, this shows how prior probabilities about states are captured in this formulation. The axiom about strong order-determinacy is a "minimality" assumption. In other words, we are using the smallest state set that can be supported by the observations (i.e., the outcomes of the simple experimentations to validate occurrence of events in E). It is interesting to observe how natural these models are. Indeed, states here (i.e., in the $\hat{S}$ picture) are supported by observed data (exactly as they should be on system theoretic grounds). They can be quite abstract constructs to represent the influence of events or hypotheses on observed data. It is then quite natural to demand certain minimality on the states. Next, observe that it is not difficult to see how this model can be extended to include local states. A reasonable departing point is as follows. A subset $\hat{\Xi}_i \subset \hat{S}$ is a set of local states for an agent A_i, if the states $\mu \epsilon \hat{\Xi}_i$ have support in the set of simple events that can be verified by the agent A_i alone, E_i. Several possibilities exist here since communications between agents will influence local state sets. Further research in this direction is needed.

The reason for this discussion is to indicate that such an event-state structure is a reasonable model for the probabilistic structure of a multi-agent stochastic control system. For example we can place Witsenhausen's formulation [13] in this model.

As a consequence of these results, we can now state our first result regarding appropriate probabilistic models in multi-agent stochastic systems.

Theorem 3: *The databases in a multi-agent stochastic control system can be used to construct an event-state structure (E, S, P) or $(E, \hat{S})$ as prescribed by Theorem 2. This representation is faithful in the sense that it can generate the satistics on which it was based.*

There are two generic examples of such event-state structures. They represent non-classical and classical probability models in multi-agent stochastic systems. In the first, we consider $Q(H)$, the set of all orthogonal projections on a separable complex Hilbert space H, and let $\leq$ be the usual order of projections. Let Q' be the orthogonal complement of Q. Then $(Q(H), \leq, ')$ is an orthomodular σ-orthoposet. Let S be the set of all positive, trace class, self-adjoint operators on H, with trace one. Let $\hat{S} = \{\mu_\rho(\cdot), \rho \epsilon S; \mu_\rho(Q) = Tr[\rho Q]\}$. Then $\hat{S}, Q(H)$, represent an event-state structure with the probability function being $P(Q, \rho) = Tr[\rho Q]$, which is, of course, von-Neumann's Hilbert space model.

The second example, consists of a σ-algebra E of subsets of a set X and $\hat{S}$ is a σ-convex, strongly order-determining set of probability measure on X. This is the classical Kolmogorov model of probability theory with several probability measures.

4 Communication Constraints, Incompatible Events and Event-State-Operation Structures

On intuitive grounds, we expect that in a multi-agent system we will have noncompatible events; that is, events whose occurrence cannot be simultaneously verified by two or more agents. Whether or not the verification procedure used is causal or non-causal is not essential. This can be seen as a manifestation of communication constraints in a networked control system, for example. It can also be interpreted on the grounds of expected interactions between agents in a multi-agent stochastic control problem.

Consider for a moment a distributed sensor network. It is clear that a specific sensor will be able to verify the occurrence or not of a restricted set of simple events. One may legitimately define this subset as the *domain of observation* or *sensor range* of the sensor. Similarly, an agent in a multi-agent stochastic control problem will be able to influence the occurrence of a subset of simple events. One may define this subset as the *domain of influence* or *control range* of the agent.

Having concluded that natural models for multi-agent stochastic systems must allow for incompatible events, a natural question to ask is: How can we build probabilistic models which have this incompatibility property built in automatically? Towards answering this question, let us remark that there is a fundamental difference between the two examples of Section 3: namely, the first allows for events which are not compatible (i.e., not simultaneously verifiable) while this cannot happen in the second. This indicates a fundamental, and highly nontrivial, limitation of classical Kolmogorov models for multiagent stochastic systems.

For a recent observation along similar lines we refer to [10]. To incorporate compatibility, we introduce a generalization of "conditioning" via certain constructs called *operations*. The notion of compatibility corresponds to a distinguished relation $\mathcal{C}$ on E [21,22]. We can assert that there exists at most one relation $\mathcal{C}$ on E, is determined (uniquely) by the following property: "for $p, q \epsilon E, p\mathcal{C}q$ iff there exists a subset $B \subset E$, with B Boolean, such that $p, q \epsilon B$." Boolean here means that $\wedge, \vee$ make B into a Boolean ring. Now clearly, if $p, q \epsilon E$ and $p\mathcal{C}q$, then $p \wedge q$ exists in E. In other words, if p, q are compatible events, their occurrence can be simultaneously verified in a multi-agent set-up. Now then, one deficiency of the non-commutative probability model, based on an event-state structure described in Section 3, is that the following question cannot be answered satisfactorily: If $p, q \epsilon E$, and $p\mathcal{C}q$ (i.e., p, q non-compatible), then does $p \wedge q$ exist in E? Note that this problem is related to the question whether or not E is a lattice.

The corresponding interpretation in a multi-agent stochastic system setting is apparently as follows. If data collection procedures or experiments or observations (or sensors) are described for two incompatible (i.e., non-simultaneously verifiable) events, how does one describe the data collection procedures or experiments or observations (or sensors) for the "and" (or conjunction) of these two events? The answer is provided by introducing a generalization of "conditional probability" and "conditional expectation".

This question is linked with the concept of "conditional probability" in an event-state structure. In the Kolmogorov model, the concept of conditional probability and the associated concept of Random-Nikodym derivative are of paramount significance for the analysis of classical stochastic control systems [8,15,16]. In the classical model, conditional probability is expressed as a mathematical object defined constructively in terms of the primitive entities of the theory. How can this be done in the generality of our discussion here? The hint comes from the interpretation of the operation introduced in [4,9,18] in order to handle joint statistics of repeated non-compatible measurements on the same system. That is, operations are a form of conditioning. Indeed, these set transformations are widely employed to represent the concept of conditional probability in von-Neumann's model. This leads to difficulties, however, and to the consideration of *event-state-operation structures* (Pool [21, 22]). Our constructs follow [21, 22]. An event-state-operation structure is a 4-tuple (E, S, P, T) where (E, S, P) is an event-state structure and T is a mapping, $T : E \rightarrow \sum =$ $\{$ set of all maps from S into S $\}$ which satisfies certain axioms. If $p \epsilon E, T_p$ is the *operation corresponding to the event* p. For $a \epsilon S$, we interpret $T_p a$ as the *new state* conditioned on occurrence of the event p and prior state a. The *domain* D_p of T_p are those states that can induce (or influence) the occurrence of p.

We let

$$\sum_T = \{T_{p_1} \circ T_{p_2} \circ \cdots, T_{p_n};\ p_1, p_s, \cdots, p_n \epsilon E\}$$

be the set of *operations*.

We assert the existence of an event q_x, which occurs with certainty in the state a if and only if $a\mathcal{C}D_x$. This event q_x can be used to design an experimental procedure or observation to determine whether or not a state belongs to the domain D_x of an operation.

$(\sum_T, \circ)$ is a multiplicative subsemigroup of $(\sum, \circ)$, where o is function composition. We introduce the $*$ operation on $\sum_T$, to mean reversal of application and show that it is an involution. Recall (Foulis [17]) that a *Baer*-semigroup* is an involution semigroup where the annihilator of each element is a principal left(right) ideal generated by a self-adjoint idempotent. We can now state our second fundamental result on non-commutative probability models for multi-agent stochastic control.

Theorem 4: *If (E, S, P, T) is an event-state-operation structure supported by the databases and conditioning of a multi-agent stochastic control system, then $(\sum_T, \circ, *, \sim)$ is a Baer*-semigroup.*

Here $\sim$ is the map $x \rightarrow T_{q_x}$. If one views the event-state structure of the previous section as a passive picture (in the sense that it considers only the probability of occurrence of events), then the introduction of the concept of operation provides an active picture. Now it is easily seen from Theorem 4 that for $p, q \epsilon E$, $p \leq q$ iff $T_p \circ T_q' = T_p$. This last property reminds us of the so-called "smoothing" property of conditional expectation.

We started this section with the problem of determining whether $p \wedge q$ exists for $p, q \epsilon E$. The natural question to ask now is: Can the greatest lower bound $p \wedge q$, for p, q be interpreted via the composition of the Baer*-semigroup, described? The answer is:

Theorem 5: *If (E, S, P, T) is an event-state-operation structure, then $(E, \leq, ')$ is an ortholattice. Moreover, if $p, q \epsilon E$, then $T_{p \wedge q} = (T_{p'} \tilde{\circ} T_q) \circ T_q$*

We also have:

Theorem 6: *Assumptions as in Theorem 5. Then for $p, q \epsilon E$ the following are equivalent:* (a) pCq, (b) $T_p \circ T_q = T_q \circ T_p$. *If pCq then $T_{p \wedge q} = T_p \circ T_q$.*

The important conclusion is that compatibility of events *has been interpreted as commutativity of the associated operations*, i.e., of the *associated conditioning* or *inferences*. This is a fundamental result which has significant implications in the analysis and control of multi-agent stochastic systems.

So in the setting of Baer*-semigroups, we can associate the compatible events with commutativity of the corresponding operations. Note [19-22] that operations and observables are quite different kinds of entities. It is accepted that the various constructs of Baer*-semigroups have nice physical interpretations in this model of quantum physics.

Returning to multi-agent stochastic control systems, we note that classical systems are characterized by a commutative Baer*-semigroup structure while non-classical information patterns correspond to non-commutative semigroups. Some natural questions that need to be answered are: How do specific structural properties of the system appear in the structure of the semigroup? Another important point is that an operation associates to every event a map of the state set into itself which can be interpreted as due to a control law. That is, we can think of an operation as a model for the combined operation of obtaining a measurement and applying a control law by an agent. This is the active interpretation of an operation which is significant for stochastic control. The passive can also be used to model the system's interaction to measurements.

We have also succeeded in interpreting $p \wedge q$ for non-compatible events in terms of the operations of the semigroup. Note, however, that we need both $\circ$ and $\sim$. This result can be interpreted as a "data fusion" or "agreement" among different agents. It is quite remarkable that it comes automatically out of the imposed structure.

We have also discovered that: the orthomodular orthoposet $(E, \leq, ')$ associated with an event-state structure is not an ortholattice necessarily; the introduction of operations with the axioms given forces $(E, \leq, ')$ to be an orthomolular ortholattice! This is a major gain in the structure theory.

Substantially more work is needed in the areas of representation theory for Baer*-semigroups, classification theory, coarse structure theory and fine structure theory, in order to take full advantage of our basic results. We are investigating these problems.

On the basis of these findings and results, we then can justify our second fundamental result in the hierarchy of models we have developed.

Theorem 7: *Databases and conditioning in a multi-agent stochastic control system can be used to construct an event-state-operation structure. This is a faithful representation.*

Regarding representation theory, there exists a large body of mathematical work on representation theory for Baer*-rings, Rickart*-rings. In particular, it is known that we can embed the Baer*-semigroup in a C^*-algebra structure. Then via the Gelfand-Naimark theorem we can ascertain that there exists a Hilbert space H such that a C^* algebra A is *-isomorphic to a closed *-subalgebra of $\mathcal{L}(H)$ = the space of bounded operators on H. These ideas can be pushed further toward a more detailed structure theory of Baer*-semigroups.

Let us give an example in that direction and of the results one can get. What one would like to do, of course, is to identify assumptions on the information pattern and flow that allow explicit representation of the Baer*-semigroup. Then one finds this structure. Subsequently, the value of these results, and the validity of the axioms, can be tested by examples in distributed stochastic control problems. We can for instance show the following.

Theorem 8: *If $(E, \leq)$ is atomic and the atoms $\hat{E}$ are mapped to pure states under T, then (E, S, P, T) can be represented as the lattice of projections on a Hilbert space.*

The assumption is typically valid in applications. Furthermore in applications the Hilbert space is often finite dimensional.

5 Non-commutative Probability Models: Convex Structures

The algebraic structures discovered in Sections 3 and 4 can be used effectively to understand several modeling problems in multi-agent stochastic control systems. However, they are not suitable for the formulation and analysis of optimization problems. For the latter, it is important to have such properties as convexity. Such more sophisticated models are developed here.

Let us consider here the particular Baer*-semigroup that pertains in the generic model described at the end of Section 4. We know that this model admits an operation map. It turns out that it is more convenient to work with the unnormalized version of operations [11]. So in the Hilbert space model, an operation is a positive linear map $T : \mathcal{T}_S(H) \to \mathcal{T}_S(H)$ which also satisfies

$$0 < Tr[T(\rho)] \leq Tr[\rho]$$

for all $\rho \in \mathcal{T}^+(H)$. We emphasize again the phenomenological interpretation of an operation: An operation describes the change of state associated with an observation. The probability of transmission of a state ρ by an operation S is taken to be $Tr[S(\rho)]$, while the output state conditioned upon transmission is taken to be

$$\rho_{out} = \frac{S(\rho)}{TrS(\rho)}$$

Associated with the operation S is its *effect*, defined as the unique operation A for which

$$Tr[S(\rho)] = Tr[\rho A]$$

for all $\rho \in \mathcal{T}_s(H)$. The interpretation of A is that it determines the probability of transmission but not the form of transmitted state. It is now seen that this unnormalized version of operations leads to a slightly different model for the propositional calculus which is actually more satisfactory. The set of all effects $\mathcal{EF}$ consists of all bounded operators A on H such that $0 \leq A \leq 1$. $\mathcal{EF}$ is a poset and has a least and greatest element. Furthermore, it has an orthocomplementation given by the map $A' = 1 - A$. It is easily seen that the set of orthogonal projections $\mathcal{P}(H)$ is the set of extreme points of $\mathcal{EF}$. Note, however, that $\mathcal{EF}$ is not a lattice, but on the other hand, it is a convex set in $\mathcal{L}(H)$. Furthermore, $\mathcal{P}(H)$ is dense in $\mathcal{EF}$ for the weak operator topology. This circle of ideas emphasizes $\mathcal{T}_S(H)$ as the state (or ensemble) set for the model, i.e., consider the normalization $Tr[\rho] = 1$ of secondary importance. Since we know that the density operator of a system is the analog of the probability density of a stochastic system, it is seen that the above argumentation is akin to considering the unnormalized probability density in classical formulations of filtering and control problems [6, 15]. This often turns some of the crucial equations to linear ones! See, for example, the unnormalized conditional density equation of classical nonlinear filtering of diffusion processes [6].

These ideas can be applied to the general Baer*-semigroup setting, and this was done by Ludwig [19]. Starting from an event-state structure (E, S, P) one embeds S into the real vector space V of functions on S defined by $X(p) = \sum_{i=1}^{n} c_i P(p, a_i)$, $a_i \in S, c_i$ real numbers, n arbitrary. Under this embedding $a \mapsto P(\cdot, a)$. By letting

$$\| X \| = \sup_{p \in E} | X(p) |$$

for $X \in V, V$ becomes a normed linear space, and we consider its completion which we also write as V. Then V is a real Banach space and let V^* be the dual. By considering $P(p, X) = X(p), P$ can be defined on the whole of V, and for p fixed this defines a linear functional on V, allowing us to identify E with a subset of V^*. Furthermore, one can introduce a partial order in V by a cone V^+. V is usually called the state space by Davies [11-13]. In V the norm is linear on V^+ and, therefore, can be uniquely extended to a positive linear functional $\tau : V \to R$, with $| \tau(x) | \leq \| x \|$ for all $x \in V$, and $\tau(x) = \| x \|$ for all $x \in V^+$. The states S are identified as the elements of $\{x \in V^+ : \tau(x) = 1\}$ and form a convex set. The set of effects $\mathcal{EF}$ is in this general setting identified with

$$\mathcal{EF} = \{\phi \in V^*;\ 0 \leq \phi \leq \tau\}.$$

$\mathcal{EF}$ is convex, weak* compact, partially ordered and has the orthocomplementation

$$\phi^\perp = \tau - \phi.$$

The events E are identified as the extreme points of $\mathcal{EF}$.

There are several advantages of doing this transformation from an event-state structure to a state space structure: (a) we embed a nonlinear structure into a linear richer structure, (b) the relationship with classical probability theory becomes more apparent, (c) the theory fits in nicely with the use of C^*-algebras [19, 20] in quantum statistical mechanics and quantum field theory.

We summarize these results in

Theorem 9: *To any event-state-operation structure (E, S, P, T), there corresponds a pair of Banach spaces V, V^*, with positive cones and a trace functional τ on V as above. S is identified with a convex subset of V. E as the set of extreme points of a convex subset of V^*. Operations in $\sum_T$ correspond to linear positive maps $T : V \to V$ such that $0 \leq \tau(T_x) \leq \tau(x)$ for all $x \in V^+$. To every operation T we can associate its effect φ_T via*

$$\tau(T_x) = \varphi(x), \quad \forall x \in V.$$

We turn now in a discussion of more complex (but more realistic) models of the observation process in a multi-agent system, utilizing the framework of a state space. It is important to realize that the concepts and constructions to be introduced can actually be worked out for the general setting of an abstract state space (or a Baer*-semigroup).

Thinking of the transformation performed on a state by an operation as one corresponding to a simple yes-no validation, it is easily seen that for a general continuous observation we need to consider operation valued measures OVM. Inspired by some similar constructs in quantum communication systems [4, 7, 9, 18], we introduce the following.

A *generalized sensor* (or observation) on a measurable space (U, B) is a map $M : B \to \mathcal{L}^+(V)$ such that

(i) $M(B) \geq M(\phi), \ \forall B \in \mathcal{B}$

(ii) for $B_i \cap B_j = 0, i \neq j, \ M(\coprod B_i) = \sum_{i=1}^{\infty} M(B_i)$

(iii) $\tau(M(U)\rho) = \tau(\rho)$ for all $\rho \in V$.

The interpretation is that a generalized sensor accepts a state, measures some properties and emits an output state conditioned on the value of the measurement (observation). Families of generalized sensors become important when we consider dynamic problems. For static problems we need only the *measurement* associated with a generalized sensor. Indeed for distributed detection or estimation for example there is no *action or control*.

The *measurement* K_M associated with the *generalized sensor* M is the unique V^*-valued measure such that

$$K_m(B)(p) = \tau[M(B)\rho]$$
$$\forall \rho \in V, \quad B \in \mathcal{B}.$$

We would like to note that the statistics of the observed (or collected) data, when the system "state" is ρ, by a generalized sensor M are given by the probability measure $K_M(B)(\rho), \forall B \in \mathcal{B}$. Consider, now, two generalized sensors M_1 on U_1 and M_2 on U_2 with values in $\mathcal{L}^+(V)$, then we can define their composition M_{12} as a sensor on $U_1 \times U_2$, which represents the operation of first M_1 and then M_2. The composition is uniquely determined from the equation

$$M_{12}(B_1 \times B_2)(\rho) = M_2(B_2)M_1(B_1)\rho$$

for all $\rho \in V$ and all B_1, B_2. This then leads naturally to a family of generalized sensors parametrized by time if we wish to describe repeated observations from the same system. In a series of papers [11-13], Davies introduced such families when the measurement outcomes form a marked point process and he termed them quantum stochastic processes. We introduced a generalization to allow for outcomes with continuous sample paths. Let U be a complete separable metric space, $\mathcal{B}$ the Borel σ-algebra on $U, \mathcal{Y}_t$ the set of all measurable functions form $[0, t]$ into U and $\mathcal{F}_t$ a σ-algebra on $\mathcal{Y}_t$. A *generalized stochastic process with outcomes adapted to* $\mathcal{F}_t$ is a family of generalized sensors M_t on $\mathcal{Y}_t$ such that:

(i) $\lim_{t \to 0} M_t(\mathcal{Y}_t)\rho = \rho$ for all $\rho \in \mathcal{T}_S(H)$,

(ii) $M_s(B)M_t(A)\rho = M_{t+s}(c(A \times B))\rho$ for all $A \in \mathcal{F}_t, B \in \mathcal{F}_s$, where c maps $\mathcal{Y}_t \times \mathcal{Y}_s$ onto $\mathcal{Y}_{t+s}$ via concatenation of sample paths.

Note that (ii) is the appropriate analog of the Chapman-Kolmogorov equation of classical probability. The physical interpretation is clear. So $M_t(A)\rho$ is the new state given the initial state was ρ and that the sample path of the outcome process was in $A \in \mathcal{F}_t$. The one parameter family $T_t = M_t(\mathcal{Y}_t)$ forms a semigroup of operators on $\mathcal{T}_S(H)$ describing the evolution of the state as perturbed by measurement. If we let z denote the empty sample path, then $S_t = M_t(\{z\})$ is also a semigroup on $\mathcal{T}_S(H)$ describing the evolution of the state unperturbed from measurements. For certain processes, Davies was able to characterize the differential version of the effects of measurement.

We close this section with an important result of Naimark [18] which has significant implications to questions of implementation of measurements, generalized sensors, etc. A natural question, is how are measurements interpreted? Briefly, consider a measurement K on the product measurable space $(U, \mathcal{B}) = (U_1, \mathcal{B}_1) \times (U_2, \mathcal{B}_2)$. If one could measure together two quantities which take values in U_1, U_2 then their composite measurement should be describable by such a K on $(U, \mathcal{B})$. Then we can define the "marginal measurements" M_1 on U_1, M_2 on U_2 via $M_1(B) = M(B \times \mathcal{B}_2)$, and similarly for M_2. It is a well known fact that if both M_1, M_2 are projection valued then M_{12} is projection valued and M_1, M_2, commute. There is a way to interpret M as a coordination of two noncompatible observers (sensors). This is provided by Naimark's theorem which asserts that given a POM M on $(U, \mathcal{B})$ with values in $\mathcal{L}(H)$, there exist a pure state ρ_e on a Hilbert space H_e and a PVM E_M on $H \times H_e$ such that $Tr[(\rho \times \rho_e)E(A)]$, for any ρ and any $A \in \mathcal{B}$.

The triple $\{H_e, \rho_e, E\}$ is called a *realization of* M and the physical interpretation is that M is statistically equivalent to the simultaneous measurements performed by compatible observers on an augmented system (the original augmented by the auxiliary system (H_e, ρ_e). Examples of such constructions appear in quantum communication problems [4,9,18].

The advantage of such representations is that they provide compact models for the statistics of fairly complicated observation processes. The disadvantage is that in the context of multi-agent stochastic systems, knowing $K_M(\cdot)$, it is not possible to recover the actual measurements and communications strategies or histories.

6 Distributed M-ary Detection

Let us now return to the distributed M-ary detection problem described in Section 2.2. Suppose that the N agents operate asynchronously over a time interval $[0, T]$. We collect all the local observation times $t_k^i, i = 1, \cdots, N$ and globally order them. Each agent $i, i = 1, \cdots, N$, at each local instant $t_k^i, i = 1, \cdots, N, k = 1, \cdots, L(i)$, has data $y^i(t_k^i)$ (his own) plus data $z^j(t_k^i), j \neq i, j = 1, \cdots, N$, communicated to him from other agents. z^j may be processed or unprocessed. We want to ask the following fundamental question. Given arbitrary communication, how can one represent the statistics of the collected data $(y^i(t_k^i), z^j(t_k^i)), i = 1, \cdots, N, j \neq i, j = 1, \cdots, N, k = 1., L(i)$? We have the following answer.

Theorem 10: *In the distributed, M-ary detection problem described above, any sequence of observations and communications between the agents can be represented by an appropriate measurement $K_M(\cdot)$ on some measurable space $(U, \mathcal{B})$.*

The proof is nonconstructive. This we consider as an important conceptual tool, particularly with respect to obtaining performance bounds. The latter is its greatest advantage. Its disadvantage is that it is not possible to recover from K_M the actual observation process *and* the communication strategy.

Another important point that was made earlier is that in this specific setting, Naimark's extension theorem provides a natural way of coordination between noncompatible observers. The fact that this comes out of the mathematical model automatically is a measure of success of the underlying models that we constructed.

To solve now the M-ary distributed detection problem in view of the representation result presented in Theorem 6, one proceeds as follows. Here, we concentrate on the Hilbert space model, but it should be clear by now how to extend the computation to more general models.

Given the M-hypotheses $H_i, \cdots, H_M$, one constructs risk operators $W_i \in \mathcal{T}_S(H), i = 1, \cdots, M$, based on assumed costs, states corresponding to the M hypotheses and prior probabilities. We allow, of course, randomized strategies, and we search for the optimal measurement $K_M(\cdot)$, subject to some information pattern constraint. Let

$$\Pi_i = \int_U \Pi_i(u) K_M(du), \; i = 1, \cdots, M$$

The problem becomes

$$\min \quad Tr \sum_{i=1}^{M} W_i \Pi_i$$

over all positive operator valued measures (POM) $\Pi_i, i = 1, \cdots, M$ such that

$$\{\Pi_i\}_{i=1}^M \in \mathcal{A}.$$

Here $\mathcal{A}$ is a convex set of POM's corresponding to some *information theoretic* constraint on information (communication) patterns such as capacity constraints, for example.

This problem is a convex linear programming problem and its duality theory is well understood (see, for example, [4, 18]). One can then, in this example, begin to understand how the *duality* between decisions and information patterns can be put in a firm framework. Further work is needed along this promising direction, however. For example, for the unconstrained problem, we have the following result [5, 18].

Theorem 10: *Suppose $\mathcal{A}$ above is the set of all POM's. Then a necessary and sufficient condition for the POM $\Pi_i^*, i = 1, \cdots M$, to be optimal is that*

(i) $\sum_{j=1}^M W_j \Pi_j^* \leq W_i, i = 1, \cdots M$

(ii) $\sum_{j=1}^M \Pi_j^* W_j \leq W_i, i = 1, \cdots, M.$

Furthermore, under any of the above conditions the operator

$$Y = \sum_{j=1}^{M} W_j \Pi_j^* = \sum_{j=1}^{M} \Pi_j^* W_j$$

is self-adjoint and is the unique solution of the dual problem.

It is easy to see that the above conditions are equivalent to Y being self-adjoint and

$$W_i \geq Y, \; i = 1, \cdots, M.$$

Then these imply

$$(W_i - Y)\Pi_i^* = \Pi_i^*(W_i - Y); \; i = 1, 2, \cdots, M$$

and that the minimum value is TrY.

We would like to close this section by mentioning that these results can be extended to include estimation problems. The major outstanding open problem is that of *implementation*. That is, if we find the optimal Π_i^* , how do we realize it by a communication pattern and a measurement process? It is also possible to interpret the Lagrange multipliers (here the Y), as sensitivities with respect to the information pattern constraints.

References

[1] L. Accardi and C. Cecchini, "Conditional Expectations in von Neumann Algebras and a Theorem of Takesaki", *J. Functional Anal.*, vol. 45, pp. 245-273, 1982.

[2] M. Andersland and D. Teneketzis, "Information Structures, Causality, and Nonsequential Stochastic Control II: Design-Dependent Properties", *SIAM J. on Control and Optimization*, 32, 6, pp. 1726-1751, 1994.

[3] D. Bansâl and T. Basar, "Solutions to a Class of Linear-Quadratic-Gaussian (LQG) Stochastic Team Problems With Nonclassical Information", *System and Control Letters*, 9, pp. 125-130, 1987.

[4] J.S. Baras, "An Optimization Problem from Linear Filtering with Quantum Measurements", *Journal of Applied Math. and Optim.*, Vol. 18, pp. 191-214, 1988.

[5] J.S. Baras, "Non-commutative Probability Models in Mutli-Agent Stochastic Control", to appear.

[6] J.S. Baras, G.L. Blankenship and W.E. Hopkins, Jr., "Existence, Uniqueness and Asymptotic Behavior of Solutions to a Class of Zakai Equations with Unbounded Coefficients", *IEEE Transactions on Automatic Control,* Vol. AC-28, pp. 203-214, February 1983

[7] J.S. Baras, "Noncommutative Probability Models in Quantun Communication and Multi-Agent Stochastic Control," in *Ricerche di Automation,* December 1979, Vol. 10, No. 2, pp. 217-264.

[8] J.S. Baras, R. Elliot and M. Kohlmann, "The Partially Observed Stochastic Minimum Principle", *SIAM Journal on Control and Optimization,* Vol. 27, No. 6, pp. 1279-1292, November 1989.

[9] J.S. Baras and R.O. Harger, "Quantum Mechanical Linear Filtering of Vector Signal Processes", *IEEE Trans. on Information Theory,* Vol. IT-23, No. 6, pp. 683-693, November 1977.

[10] A. Benveniste, E. Fabre, C. Jard, and S. Haar. "Diagnosis of Asynchronous Discrete Event Systems, a Net Unfolding Approach", to appear in *IEEE Trans. on Automatic Control*, 2003.

[11] E.B. Davies, *Quantum Theory of Open Systems*, Academic Press, London, 1976.

[12] E.B. Davies, "Generators of Dynamical Semigroups", *J. Funct. Anal.*, 34, pp. 421-432, 1979.

[13] E.B. Davies and J.T. Lewis, *Comm. Math. Phys.*, 17, p. 239, 1979.

[14] R.J. Elliott, *Viscosity Solutions and Optimal Control*, Pitman Research Notes in Mathematics, Longman, 1987.

[15] R.J. Elliott, L. Aggoun and J.B. Moore, *Hidden Markov Models: Estimation and Control,* Applications of Mathematics, Springer-Verlag, 1997.

[16] R. Elliott and M. Kohlmann, "On the Existence of Optimal Partially Observed Controls", J. of Aplied Mathematics and Optimization, 9, pp. 41-66, 1982.

[17] D.J. Foulis, *Baer* -Semigroups,* Proc. of Am. Math. Soc., 1960.

[18] A.S. Holevo, *Statistical Strucrure of Quantum Theory,* Lect. Notes Phys. m67, Springer, Berlin, 2001.

[19] G. Ludwig, *An Axiomatic Basis for Quantum Mechanics*, Springer-Verlag, 1987.

[20] D. Petz, "Conditional Expectation in Quantum Probability", in *Quantum Proba. and Appl.*, Vol III, Lect. Notes in Maths., Springer 1303, pp. 201-260, 1988.

[21] J.C.T. Pool, "Baer*-Semigroups and the Logic of Quantum Mechanics", *Comm. on Mathematical Physics*, 9, 118-141, 1968.

[22] J.C.T. Pool, "Semi-Modularity and the Logic of Quantum Mechanics", *Comm. on Mathematical Physics*, 9, pp. 212-228, 1968.

[23] A. Sahai, S. Tatikonda and S.K. Mitter, "Control of LQG Systems Under Communication Constraints", *Proc. 1999 American Control Conference*, Vol. 4, pp. 2778-2782, 1999.

[24] A. Sahai, *Anytime Information Theory,* Ph.D. Thesis, Dept. of Electrical Enginering and Computer Science, MIT, Febr. 2001.

[25] H.S. Witsenhausen, "A Counterexample in Stochastic Optimum Control", *SIAM J. of Control,* 6 (1), pp. 131-147, 1968.

[26] H.S. Witsenhausen, "Separation of Estimation and Control for Discrete Time Systems", *Proc. of the IEEE*, 59 (11), pp. 1557-1566, 1971.