

MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS-1963-A

12

January, 1983

LIDS-P-1268

AD A124940

OPTIMAL DECENTRALIZED CONTROL OF
FINITE NONDETERMINISTIC SYSTEMS¹

Robert R. Tenney
September 1982

DTIC FILE COPY

DTIC
SELECTED
FEB 25 1983
A

This document has been approved
for public release and sale; its
distribution is unlimited.

1. Supported by the Office of Naval Research, under contract No. N00014-77-0532 (N041-519) is gratefully acknowledged.
2. Laboratory for Information and Decision Systems, MIT, Rm. 35-213, Cambridge, MA 02139

83 02 03 009

ABSTRACT

A control problem with a system modeled as a nondeterministic finite state machine is considered. Several agents seek to optimize the behavior of the system under a minimax criterion, with each agent having different information about the system state. The nondeterministic model of uncertainty, combined with the minimax criterion, lead to equivalence relations on the past input/output histories of each agent which generate rather simple sufficient statistics for the optimal control laws. This sheds light on the basic nature of decentralized control and permits complete solution of a particular class of problems.



Accession For	
NTIS GRA&I	<input checked="" type="checkbox"/>
DTIC TAB	<input type="checkbox"/>
Unannounced	<input type="checkbox"/>
Justification	
<i>Put on file</i>	
Pr	
Distribution/	
Availability Codes	
Dist	Avail and/or Special
A	

OPTIMAL DECENTRALIZED CONTROL OF
FINITE NONDETERMINISTIC SYSTEMS

I. INTRODUCTION

It is almost a cliché that computational resources are becoming increasingly inexpensive. That this is not true of communications resources, particularly over long distances, is a prime motivation for the development of design methods for computationally distributed solutions to large scale control problems [1]. The basic structure of such a problem is that a system, subject to random influences, is to be controlled by a group of several powerful control nodes, each of which has different information about the events in the system due to communication constraints. The control algorithms in each node must be designed to produce good system behavior when they act in concert with the system.

Although problems of this nature have been considered for many years [2,3], progress towards a general solution has been surprisingly slow. Many problems with very special structure have been solved [4 - 9], but problems with a more general structure have defied tractible analysis [10,11]. Some understanding of the difficulties encountered has been gained, but there is no class or problems with a general, dynamic, decentralized structure for which exact solutions can be found and used to motivate further investigation.

The motivation for this work is to cure that deficiency. An exceedingly simple class of decentralized control problems is posed, and

solved completely. The simplicity stems from the assumed system model and performance objective - a simplicity which is deliberately introduced to illuminate the deep structure of decentralized problems, yet which preserves the essential features of more complex applications: stochastic dynamics, a general information structure with constrained or no communication, and a desire to achieve performance which is optimal in some precise sense.

The system to be controlled is assumed to be described as a finite state, nondeterministic automaton with inputs supplied by several control agents. Each agent receives an observation at each discrete time step which indicates a set in which the current state must lie. Each observation, coupled with knowledge of the system structure, can lead to inferences about the past system behavior, hence about the past observations of other agents and thus prediction of other agents' decisions. This simultaneous interleaving of inference by the agents, as each deduces the potential actions of others, and the deductions of others about itself, etc. leads to some of the complexity of the analytical process.

The notion of using nondeterministic, as opposed to Markov, models is not new. They are introduced in a centralized context in [12-13], and used in a decentralized problems in [14]. Their advantage is that the set theoretic operations which describe their operation are extremely simple in structure, and offer a rich opportunity for manipulation. Their disadvantage is that they tend to lead to "conservative" control laws, as decisions are made to avoid highly unlikely, expensive transitions which a classical probabilistic approach would discount.

Based on this model structure, an approach to addressing the problem of designing the optimal (in the appropriate worst-case sense) decision

rules has been developed. The approach is based on the identification of the set of sufficient statistics for each agent to use and the dynamic relations between them; these sufficient statistics are no more than the intertwined deductions of the agents about each other truncated at the point where they are no longer productive. The set of these statistics form an extended state space over which dynamic programming may be used to derive the optimal decision rules.

This approach is the major contribution of this work. It enables the study of a class of decentralized, stochastic, dynamic, optimal control problems which is sufficiently simple that exact solutions may be obtained. These solutions may serve to suggest heretofore unsuspected properties of optimal solutions in general. However, the possibility of direct application of this approach to engineering problems does exist: many protocols for distributed resource allocation in computer networks (e.g. data base management or asynchronous channel access) can be formulated in this framework (but have not yet been addressed).

In the sequel, notation and the problem formulation are introduced in section II. The complexity introduced by just the minimax, nondeterministic optimization approach in a static, decentralized setting is discussed in section III. The sufficient statistic, called the information relation, and its properties are then developed and these, coupled with the results of section III, solve the decentralized estimation problem [15] in this context. Finally, section VII solves the control problem - and quite unexpectedly, this involves only a rather straightforward extension of the estimation solution.

II. NOTATION AND PROBLEM FORMULATION

A. Notation

This work will use set valued functions

$$\underline{f}: X \rightarrow 2^Y \tag{2.1}$$

to model nondeterministic behavior, i.e. $y \in \underline{y} = \underline{f}(x)$ indicates that any element y of the set $\underline{y} \subseteq Y$ may arise as a result of applying \underline{f} to the point x . The extension of \underline{f} to a function on the power set of X

$$\underline{f}_e : 2^X \rightarrow 2^Y \tag{2.2}$$

where

$$\underline{f}_e(\underline{x}) = \bigcup_{x \in \underline{x}} \underline{f}(x) \tag{2.3}$$

will not be distinguished from \underline{f} itself. Abusing notation slightly, the preimage of y under f will be denoted by

$$\underline{x} = \underline{f}^{-1}(y) = \{x \mid y \in \underline{f}(x)\} \tag{2.4}$$

Abusing notation more than slightly, \underline{f}^{-1} will not be distinguished from its extension to

$$\begin{aligned} \underline{f}_e^{-1}(\underline{y}) &= \bigcup_{y \in \underline{y}} \underline{f}^{-1}(y) \\ &= \{x \mid \underline{y} \cap \underline{f}^{-1}(x) \neq \emptyset\} \end{aligned} \tag{2.5}$$

Finally, any cost functional

$$J : X \rightarrow \mathbb{R} \tag{2.6}$$

will have an extension

$$J_e(\underline{x}) = \max_{x \in \underline{x}} J(x) \tag{2.7}$$

which will also not be distinguished from its restriction, J .

Subscripts will indicate the agent associated with each variable.

Superscripts will indicate elements of a Cartesian product set; e.g.

$$x^t = (x(1), \dots, x(t)) \in X^t.$$

B. Dynamics

The system to be controlled will have a finite state space X with N elements and dynamics

$$x(t+1) \in \underline{f}(x(t), u_1(t), u_2(t)) \quad (2.8)$$

$u_1(t) \in U_1$, $u_2(t) \in U_2$ are inputs supplied by two decision agents (extension to several agents is straightforward). The sets U_1 and U_2 will generally be considered finite, although allowing them to be real intervals will aid the interpretation of some examples. The state and its dynamics may be taken to include any interagent communication mechanisms.

Observations

$$y_i(t) \in \underline{h}_i(x(t)) \quad (2.9)$$

are available to each agent at each time just before $u_i(t)$ is to be selected. These may include outputs of communications channels, representing messages introduced by other agents at earlier time steps.

The initial state is assumed fixed, known to all agents and is denoted x^0 . An imperfect initial information state can be modeled by letting $\underline{f}(x^0, u_1, u_2)$ be independent of u_1, u_2 and equal to what would otherwise be the set of initial states.

C. Information

Each agent is assumed to have perfect recall of all past observations and decisions. The information sequence I_i includes these via the dynamics¹

$$I_i(0) = () \quad (2.10)$$

$$I_i(t+1) = y_i(t+1) \cdot I_i(t)$$

¹ Since each agent has perfect recall, past decisions may be formally omitted from (2.10) as they may be recursively derived from previous observations only.

where () is the empty sequence, and \cdot denotes concatenation. The decision rule by which $u_i(t)$ is selected is restricted to being a causal function of the local information sequence

$$\gamma_i(\cdot; t) : I_i(t) \rightarrow U_i \quad (2.11)$$

Γ_i^T represents the entire sequence of decision rules for agent i .

D. Objective

A cost function

$$J: X \times U_1 \times U_2 \rightarrow \mathbb{R} \quad (2.12)$$

expresses the penalty incurred if a pair (u_1, u_2) are applied as input when the system occupies state x . (Note that when J is an indicator function on some subset of X , then the problem becomes one of maintaining $x(t)$ in a fixed target set in X , as treated in [14]).

Taken together, the dynamics (2.8 - 2.10) and control laws Γ_1^T, Γ_2^T recursively define a set of possible joint state, control, and information trajectories. Define

$$J(\Gamma_1^T, \Gamma_2^T) = \max_{t \in \{0, \dots, T\}} \max_{\substack{x(t) \\ u_1(t) \\ u_2(t)}} J(x(t), u_1(t), u_2(t)) \quad (2.13)$$

where the $x(t), u_1(t), u_2(t)$ range over all values jointly possible at time t as determined by this recursion. The overall objective is to

$$\text{minimize}_{\Gamma_1^T, \Gamma_2^T} J(\Gamma_1^T, \Gamma_2^T) \quad (2.14)$$

subject to (2.11).

III. THE STATIC CASE

Consider, as a start, the one-step decision case where x^0 gives rise to $x \in \underline{f}(x^0)$, each of two agents obtains an observation $y_i \in \underline{h}_i(x)$, and they seek to select $\gamma_i : y_i \rightarrow u_i$ to

$$\begin{array}{ll} \text{minimize} & \max \\ \gamma_1, \gamma_2 & x \\ & \gamma_1 \\ & \gamma_2 \end{array} J(x, u_1, u_2) \quad (3.1)$$

A. Necessary Conditions

The person-by-person-optimality (PBBO) conditions [4] are necessary conditions which the optimal decision rules γ_1^* and γ_2^* for (3.1) must satisfy. They are established as follows.

Assume γ_2^* is known. Then (3.1) becomes

$$\begin{array}{ll} \text{minimize} & \max \\ \gamma_1 & x \\ & y_1, y_2 \end{array} J(x, u_1, \gamma_2^*(y_2)) \quad (3.2)$$

Since γ_1 may depend on y_1 , it is sufficient to consider the values of (x, y_1, y_2) in (3.2) which are consistent with each observed value of y_1

$$\gamma_1^*(y_1) = \arg \min_{u_1} \max_{\substack{x \\ y_1 \\ y_2}} J(x, u_1, \gamma_2^*(y_2)) \quad (3.3)$$

} consistent

The values of x consistent with an observation y_1 (in the sense that both are possible simultaneously) are

$$\hat{x}(y_1) = \underline{f}(x^0) \cap \underline{h}_1^{-1}(y_1) \quad (3.4)$$

The \hat{x} defined in (3.4) will be called the conditional state set, as it plays the same role in nondeterministic problems as does the conditional state probability distribution in Markov problems. The values of y_2 can also be inferred to lie in $h_2(\hat{x}(y_1))$, so (3.3) becomes

$$\gamma_1^*(y_1) = \arg \min_{u_1} J(\hat{x}(y_1), u_1, \gamma_2^*(h_2(\hat{x}(y_1)))) \quad (3.5)$$

Symmetric arguments yield

$$\gamma_2^*(y_2) = \arg \min_{u_2} J(\hat{x}(y_2), \gamma_1^*(h_1(\hat{x}(y_2))), u_2) \quad (3.6)$$

and (3.5 - 3.6) are the PBPO conditions for the static problem.

These conditions have the same qualities as PBPO conditions for other static team problems [4,7,15]. First, the decision made in response to y_i depends on an inference of the other agents' decision, hence its observation - the so-called "second guessing" phenomenon. Second, determination of γ_1^* for a specific value of y_1 depends on the entire structure of γ_2^* , which in turn depends on the entire function γ_1^* . Thus these conditions do not allow γ_1^* and γ_2^* to be derived on a point-by-point basis as is the case with centralized problems; a joint solution for the both complete functions must be found. These characteristics are shared with classical team problems [2,4,7]; this helps substantiate the claim that the nondeterministic formulation is qualitative similar to the usual Markov one.

B. Example

To help interpret the PBPO conditions, consider a specific example. $x = \{1,2,3\}$. Agent 1 can distinguish 1 from $\{2,3\}$, and agent 2 may distinguish 2 from $\{1,3\}$, via observations (i.e. $h_1(1) \neq h_1(2) = h_1(3)$). Decisions are to be drawn from $U_1 = U = [0,1]$. The cost function is:

$$J(x, u_1, u_2) = \begin{cases} u_2 & x = 1 \\ 1 - u_1 & x = 2 \\ |u_2 - u_1| & x = 3 \end{cases} \quad (3.7)$$

This is sketched in Figure 1 according to conventions to be adhered to throughout this work. States which generate the same observation to agent 1 are depicted in the same row; those with the same image under h_2 are in the same column. J is plotted as a function of u_1 and u_2 for each state.

It is immediately apparent that $\gamma_1(1)$ and $\gamma_2(2)$ are inconsequential (where $\{1\} = h_1(\{1\})$, $\{2\} = h_2(\{2\})$). Thus γ_1 is completely characterized by $\gamma_1(2)$, and γ_2 by $\gamma_2(1)$.

Suppose $\gamma_1(2) = \beta_1$. Then (3.6) gives

$$\begin{aligned} \gamma_2(1) &= \operatorname{argmin}_{u_2} \max \{u_2, |u_2 - \beta_1|\} \\ &\triangleq \beta_2 \end{aligned} \quad (3.8)$$

(3.5) gives

$$\gamma_1(2) = \operatorname{argmin}_{u_1} \max \{1 - u_1, |\beta_2 - u_1|\} \quad (3.9)$$

These are not the usual relations encountered in optimization, but $\beta_1 = 2/3$, $\beta_2 = 1/3$ can be found as the unique solution.

This example demonstrates the character of decentralized non-deterministic static optimization, and underscores the comments at the end of the previous section. In general, these conditions are quite difficult to solve computationally; it is critical to the practical solution of the dynamic problem that the domains of the γ_i be as small as possible.




Figure 1: Cost and Information Pattern

C. Structure

There is some simplification possible. Begin with

Lemma 1: If $\hat{x}(y_i) = \hat{x}(y_i')$, then $\gamma_i^*(y_i) = \gamma_i^*(y_i')$.

Proof: (3.5) and (3.6). □

This states that the conditional state set is a sufficient statistic for the static team problem, as is the conditional state distribution in Markov problems [15]. However, the set theoretic nature of the problem allows another simplification:

Lemma 2: If $\hat{x}(y_i) \subseteq \hat{x}(y_i')$, then we may set

$$\gamma_i^*(y_i) \triangleq \gamma_i^*(y_i')$$

without additional cost.

Proof: (3.5) and (3.6), along with the fact that

$$\hat{x} \subseteq \hat{x}' \implies J(\hat{x}, u_1, u_2) \leq J(\hat{x}', u_1, u_2) \quad \forall u_1, u_2$$
 □

This has no analogy in Markov problems. It is a result of the worst case objective, since if \hat{x}' is a worse case than \hat{x} , in terms of set containment, then there is no harm in replacing the decision corresponding to \hat{x} with that of \hat{x}' .

There are two effects of Lemma 2. The first is beneficial - the complexity of γ_i may be reduced. The second is that for certain events, namely, if y_i is observed rather than y_i' , the actual cost will be greater if the substitution is performed than if it were not. This is a result of the fact that optimal strategies for the problem as posed may not be unique, and that consideration of events other than the worst case may

allow one to determine that one strategy is intuitively superior to another even though their worst case performances are identical. The absence of such secondary considerations are partially responsible for the conservatism associated with minimax, nondeterministic problems; we will be sensitive to this issue in the sequel, but not dwell on it.

D. Interpretation

The static, nondeterministic decentralized problem brings out three points concerning more complex problems:

1. The PBPO conditions are not particularly helpful analytically. The fact that there is a tight dependence between all points of the decision rules suggests that the solution will remain tractable only if the number of points for which each γ_i must be specified remains small.

2. While many problems (e.g. estimation with no feedback of controls to dynamics) may be reduced to a sequence of static problems, this usually leads to a sufficient statistic of high dimension in decentralized settings (e.g. sets of state trajectories [15]). In view of 1, this is not sufficiently simple to lead to computable solutions.

3. There are available two mechanisms for reducing information sets, and thus the complexity of γ_i . Lemma 1 is the usual equivalence between information sets based on the property of state; Lemma 2 exploits the set theoretic structure of the problem.

The next section will introduce a construction to which both mechanisms can apply, and develop some of its algebraic properties. Section V returns to the decentralized estimation problem, showing how its decomposition to a sequence of static problems can take advantage of the new structure. Section VI achieves similar results for the general control problem.

IV. THE INFORMATION RELATION

A. Definition

Consider the system (2.8 - 2.10). In this and the next section, we will be concerned only with the autonomous case, where

$$x(t+1) \in \underline{f}(x(t)) \quad (4.1)$$

Definition: The global conditional state set at time t , denoted $\hat{x}_1(t) \subseteq X$, is the set of all possible states which the system may occupy at time t and which may be reached along a trajectory x^t which is possible given both information sequences $I_1(t)$ and $I_2(t)$. \square

This global conditional state set is analogous to the Markov conditional state distribution; it may be computed recursively.

Lemma 3: The global conditional state set may be computed from

$$\hat{x}(0) = \{x^0\} \quad (4.2)$$

$$\hat{x}(t+1) = \underline{f}(\hat{x}(t)) \cap \underline{h}_1^{-1}(y_1(t+1)) \cap \underline{h}_2^{-1}(y_2(t+1))$$

Proof: Set manipulations. (4.2) first generates all $x(t+1)$ which may arise from states in $\hat{x}(t)$, then removes those which could not generate the observed values of $y_1(t+1)$ and $y_2(t+1)$. \square

In the autonomous case, one may also generate local conditional state sets $\hat{x}_i(t)$ based on $I_i(t)$ alone:

$$\hat{x}_i(0) = \{x^0\}$$

$$\hat{x}_i(t+1) = \underline{f}(\hat{x}_i(t)) \cap \underline{h}_i^{-1}(y_i(t+1)) \quad (4.3)$$

and the various conditional state sets are related by

Corollary: $\hat{x}(t) \subseteq \hat{x}_1(t) \cap \hat{x}_2(t)$ (4.4)

Proof: By induction. $\hat{x}(0) = \hat{x}_1(0) \cap \hat{x}_2(0) = \{x^0\}$.

If

$$\hat{x}(t) \subseteq \hat{x}_1(t) \cap \hat{x}_2(t) \quad (4.5)$$

then

$$\begin{aligned} \underline{f}(\hat{x}(t)) &\subseteq \underline{f}(\hat{x}_1(t) \cap \hat{x}_2(t)) \\ &\subseteq \underline{f}(\hat{x}_1(t)) \cap \underline{f}(\hat{x}_2(t)) \end{aligned} \quad (4.6)$$

$$\begin{aligned} &[\underline{f}(\hat{x}(t)) \cap \underline{h}_1^{-1}(y_1(t+1)) \cap \underline{h}_1^{-1}(y_2(t+1))] \subseteq \\ &[\underline{f}(\hat{x}_1(t)) \cap \underline{h}_1^{-1}(y_1(t+1))] \cap [\underline{f}(\hat{x}_2(t)) \cap \underline{h}_2^{-1}(y_2(t+1))] \end{aligned} \quad (4.7)$$


□

Note that the second containment in (4.6) may not be replaced with equality in general; figure 2 shows an example where (4.4) must be a containment. Thus there is a more subtle relationship between local information sequences and the global state set than that provided by the local state sets.

Definition: An information relation R for a two agent problem is a function from two sets $Z_1(t)$ and $Z_2(t)$ to the power set of another set X . For autonomous systems, R is a primitive information relation if $Z_1(t) = Y_1^t$, $Z_2(t) = Y_2^t$, X is the state space, and

$$R(y_1^t, y_2^t) = \hat{x}(t) \quad \square \quad (4.8)$$

This definition will be extended to the case of feedback control in section VI. For autonomous systems, the primitive information relation contains all of the logical relationships between global state sets which are necessary for solving decision problems.



state	output	
	$y' \in$	$y'' \in$
1	{1, 2}	{1, 2}
2	{1, 2}	{2}
3	{2}	{1, 2}
4	{2}	{2}

$$f(x^0) = \{1, 2, 3, 4\}$$

a) Dynamics

b) Observations

t	$y'(t)$	$y''(t)$	$\hat{x}_1(t)$	$\hat{x}_2(t)$	$\underline{x}(t)$
1	1	1	{1, 2}	{1, 3}	{1}
2	2	2	{1, 2, 3, 4}	{1, 2, 3, 4}	{1, 2, 3}

c) Sample path

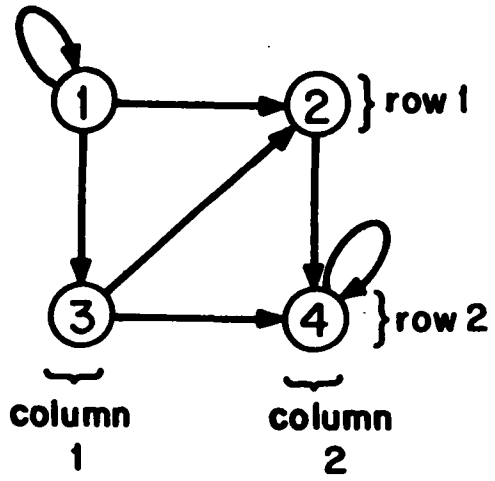
Figure 2: Example of Relation Between Conditional State Sets

Before developing the general properties of information relations, consider a specific example of a primitive information relation, taken from system (a) of figure 3. For simplicity, all examples in figure 3 will be assumed to have an observation structure whereby agent 1 can observe the row of the diagram which is occupied by the state, and agent 2 the column. (Thus observations are deterministic functions of the state). All examples have $x^0 = 1$.


At $t = 0$, $I_1 = I_2 = ()$, and $x^0 = 1$; the information relation is as shown in figure 4a. At time 1, $Y_1^1 = Y_2^1 = \{1,2\}$, and the primitive relation is shown in figure 4b. Note that there is no x consistent with observing $y_1(1) = y_2(1) = 2$, as only state 4 could cause this and it cannot be reached from state 1 in one step. This is indicated by ϕ here; subsequently inconsistent cases are left blank. In every other case, the global conditional state set is a singleton as the system is perfectly observable globally (as is the case with all of these examples).

The information relation captures the second guessing structure - the deductions each agent can make about others. The R depicted in figure 4b summarizes the fact that agent 1, receiving $y_1(1) = 2$, knows the state is 3 and that agent 2 observed $y_2(1) = 1$. However, in this case agent 2 cannot distinguish state 1 from 3 and thus cannot determine agent 1's observation; it can conclude that if agent 1 also observed 1, agent 1 in turn could not distinguish between states 1 and 2. This logic becomes quite byzantine even for this simple case; the information relation captures it succinctly.


At time 2, the primitive information relation is shown in figure 4c. The dashed lines indicate that it is comprised of four components, each the primitive information relation which would have resulted had the system




a) Simple Reducing



b) Joint Reducing




c) Infinite




d) Splitting


Figure 3. Example Systems





a) R(0)



b) R(1)



c) R(2)



d) R(3)

Figure 4. Primitive Information Relations for Example (a)

started at the prior step in each possible state. Thus just as $x^0 = 1$ expanded to $R(1)$, so the corner of $R(1)$ corresponding to $x(1) = 1$ expands to the same structure within $R(2)$. The information relations thus preserve a nesting structure generated by the nesting of the observation sets (in the sense that $y_i(1)$ is nested in $(y_i(2), y_i(1))$).

If the order of the elements of Y_i^2 are rearranged, $R(2)$ can be "unfolded" and rewritten as the alternate form in figure 4c. Labels for the rows and columns, as well as impossible cases, have been dropped for clarity. This alternate tableau form clarifies the second-guessing inference structure.

Finally, figure 4d shows $R(3)$ as generated from the alternate representation in 4c by replacing each cell occupied by a state with the primitive information relation generated by starting in that state. An alternate, unfolded form is included on the right.

Figure 4d suggests the utility of the information relation. Agent 1 can receive three distinct observation sequences from which it may infer that $x(3) = 4$. For one of these, agent 2 may not be able to distinguish whether the true state is 2 or 4; for the other two, agent 2 also knows that $x = 4$. In fact, the two cells at the upper right of the tableau are completely isomorphic in their relations to the rest of the structure, and it will be shown that the optimal decisions for agent 1 to make in these two cases are identical. Thus agent 1 need not distinguish between the two observation sequences leading to these two identical rows in the tableau; the information set Y^3 may be collapsed to a smaller sufficient statistic by identifying the two cases.

Thus the motivation for studying the information relations is to reduce a primitive relation to one of smaller dimension which still serves

for the generation of optimal decision rules. The generic sets Z_i which comprise any relation will be aggregations of the primitive information sets $I_i(t)$. The remainder of this section establishes the algebraic structure of information relations.

B. Homomorphisms.

The global conditional state sets have a lattice structure superimposed on them by the set containment relation. Information relations have a similar structure which, while not a lattice, is (almost) a partial order.

Definition: A homomorphism ϕ from R , defined from Z_1 and Z_2 to 2^X , to R' defined from Z_1', Z_2' , to 2^X , is a pair of functions $\phi = (\phi_1, \phi_2)$

$$\begin{aligned} \phi_1: Z_1 &\rightarrow Z_1' \\ \phi_2: Z_2 &\rightarrow Z_2' \end{aligned} \tag{4.9}$$

satisfying

$$\begin{aligned} R(z_1, \tilde{z}_2) &\subseteq R'(\phi_1(z_1), \phi_2(\tilde{z}_2)) & \forall \tilde{z}_2 \in Z_2 \\ R(\tilde{z}_1, z_2) &\subseteq R'(\phi_1(\tilde{z}_1), \phi_2(z_2)) & \forall \tilde{z}_1 \in Z_1 \end{aligned} \tag{4.10}$$

In the tableau representation, a homomorphism from R to R' is an association of the rows and columns of R to those of R' such that the state set in each cell of R is contained in the state set of the corresponding cell in the image.

Definition: An information relation R is contained in another R' ($R \subseteq R'$) if there is a homomorphism from R to R' . \square

This containment relation is a generalization of set containment. When $Z_2 = Z_2' = \{\phi\}$, and $R(z_1, z_2)$ is at most a singleton, this is exactly

set containment where Z_1 and Z'_1 index elements of subsets of X , (e.g. figure 5a). When $R(z_1, z_2)$ may be any set, this defines containment between two classes of subsets of X (figure 5b); note that one class may be contained in another although the second has fewer elements. Containment between information relations may be neither antisymmetric nor related to the size of the Z_i and Z'_i (figure 5c).

The containment relation is clearly reflexive (if ϕ_i is the identity map) and transitive (through composition of ϕ_i 's). It thus imposes a structure on the set of information relations which is a partial order on equivalence classes, where:

Definition: R is equivalent to R' ($R \sim R'$)

if $R \subseteq R'$ and $R \supseteq R'$. \square

Both the partial order and equivalence relation have useful interpretations in the decentralized problem.

C. Automorphisms

This section focusses on the equivalence \sim on information relations.

Definition: An automorphism ϕ on an information relation R is a homomorphism from R to R . If both ϕ_1 and ϕ_2 are 1:1, then it is an isomorphism. Otherwise, it is a reducing automorphism. \square

Clearly isomorphisms on an information relation correspond to permuting rows and columns of a tableau, and are of little conceptual interest. There are automorphisms which are not isomorphisms, and these are of considerable interest.


Definition: An information relation R is irreducible if all automorphisms on R are isomorphisms. Otherwise it is reducible. \square



a) Set containment



b) Class containment



c) Information relation containment

Figure 5. Containment Examples

As the terminology suggests, a reducing automorphism will be used to simplify, in some appropriate sense, an information relation. Provided we can show that using the simplified information relation in decision rule design yields results no worse than the original, then there will exist a mechanism for generating sufficient statistics which are simpler than the $I_i(t)$. The reducing automorphism provides this mechanism.

Definition: The reduction of an information relation $R : Z_1 \times Z_2 \rightarrow Z^X$ by a reducing automorphism $\phi = (\phi_1, \phi_2)$ is an information relation $R' : \phi_1(Z_1) \times \phi_2(Z_2) \rightarrow Z^X$ where

$$R'(\phi_1(z_1), \phi_2(z_2)) = R(\phi_1(z_1), \phi_2(z_2)) \quad (4.11)$$


□

For example, figure 6a shows successive reductions of the primitive information relation at time 3 for the system in figure 3a (as was derived in figure 4) by the reducing automorphisms indicated by arrows. The result is an irreducible relation. In this case, the irreducible relation can be constructed by a sequence of automorphisms which only reduce Z_1 or Z_2 alone. One might conjecture that this is always the case, but figure 6b shows how the primitive information relation for system 3b requires a joint reduction of both Z_1 and Z_2 at the second step in order to produce an irreducible relation.


However, the fact that every information relation can be reduced to an irreducible one by suitable compositions of reducing automorphisms (thus by some single reducing automorphism) is important.

D. The Core

This construct is the most essential part of this work.



a) Example: $R(2)$ of system 3a



b) Example: $R(3)$ of system 3b

Figure 6: Reducing Information Relations

Definition: A core of an information relation R, denoted core (R) or R*, is an irreducible information relation obtained from R by some automorphism. □

The core has a number of interesting properties. The most basic rely on the following lemma.

Lemma 4: If two irreducible information relations R and R' are equivalent under \subseteq , then they are equal¹.

Proof: If $R \subseteq R'$, then there is a homomorphism ϕ from R to R', and another, ϕ' , from R' to R.

Consider the composed homomorphism $\tilde{\phi}$ from R to R:

$$\tilde{\phi}_1(z_1) = \phi_1'(\phi_1(z_1)) \quad (4.12)$$

$$\tilde{\phi}_2(z_2) = \phi_2'(\phi_2(z_2))$$

That $\tilde{\phi}$ is indeed a homomorphism from R to R is shown by

$$R(z_1, \tilde{z}_2) \subseteq R'(\phi_1(z_1), \phi_2(\tilde{z}_2)) \quad \forall \tilde{z}_2 \in z_2 \quad (4.13)$$

since ϕ is a homomorphism from R to R', and in turn

$$R'(\phi_1(z_1), \phi_2(\tilde{z}_2)) \subseteq R(\phi_1'(\phi_1(z_1)), \phi_2'(\phi_2(\tilde{z}_2))) \quad \forall \tilde{z}_2 \in z_2 \quad (4.14)$$

by ϕ' being a homomorphism. Thus ϕ is an automorphism on R, and in fact must be an isomorphism since R is irreducible. However, the composition of ϕ and ϕ' can be an isomorphism if and only if ϕ and ϕ' are isomorphisms; hence R and R' are equal. □

¹ Up to isomorphism. This qualifier will be left implicit in the sequel.

This lemma immediately gives:

Theorem 1: The core of an information relation is unique. Moreover, if

$$R \subseteq R', \text{ then } \text{core } (R) = \text{core } (R').$$

Proof: Any core $(R) \subseteq R$, since a homomorphism from R to a core (R) exists by definition, and a homomorphism from core (R) to R exists by construction: core (R) is a (perhaps relabeled) restriction of R to subsets of Z_1 and Z_2 ; construct the homomorphism from elements of these subsets back into their original values in Z_1 and Z_2 . If two cores, core (R) and core $'(R)$ exist, then

$$\text{core } (R) \subseteq R \subseteq \text{core}'(R) \tag{4.15}$$

Both are irreducible; by transitivity of \subseteq and lemma 4 they must be equal, hence unique.

Moreover,

$$\text{core } (R) \subseteq R \subseteq R' \subseteq \text{core } (R') \tag{4.16}$$

similarly implies core $(R) = \text{core } (R')$. \square

Thus all information relations in an equivalence class of \subseteq share the same core. Restricted to irreducible information relations, the containment relation \subseteq is antisymmetric ($R \subseteq R'$ and $R \supseteq R' \Rightarrow R = R'$), transitive, and reflexive, thus inducing a partial order on the irreducible information patterns. These are the two important structural properties of information relations for any time t .

E. Dynamics

Section A informally showed how primitive information relations for the same autonomous system at different times related to one another.

Section D derived a way to reduce primitive relations to their core, thereby

creating simpler relations. This section formalizes the process by which any relation, primitive or reduced, can be extended over time and presents the basic property inherent in that process.

At any point in time $Z_i(t)$ in $R(t)$ will represent a reduced, perhaps trivially, version of Y_1^t . There is a natural way to extend this aggregation to its counterpart at time $t+1$.

Definition: The expansion of an information relation $R(t)$ to another relation $R(t+1)$ is denoted

$$R(t+1) = F(R(t)) \quad (4.17)$$

and is constructed by setting

$$Z_i(t+1) = Z_i(t) \times Y_i \quad (4.18)$$

and

$$\begin{aligned} R(t+1)[z_1(t+1), z_2(t+1)] \\ &= R(t+1)[(z_1(t), y_1), (z_2(t), y_2)] \\ &\stackrel{\Delta}{=} \underline{f}(R(t)[z_1(t), z_2(t)]) \cap \underline{h}_1^{-1}(y_1) \cap \underline{h}_2^{-1}(y_2) \quad \square \end{aligned} \quad (4.19)$$

The structure of this expansion is captured in

Lemma 5: If $R(t)$ is the primitive relation at time t , then $F(R(t))$ is the primitive relation for time $t+1$. □

Proof: By definition, $R(t) : Y_1^t \times Y_2^t \rightarrow 2^X$, and $R(t)(y_1^t, y_2^t) = \hat{x}(t)$, the global conditional state set based on (y_1^t, y_2^t) . Identifying y_i in (4.19) with the observation $y_i(t+1)$, yields

$$\begin{aligned} R(t+1)(y_1^{t+1}, y_2^{t+1}) &= \underline{f}(\hat{x}(t)) \cap \underline{h}_1^{-1}(y_1(t+1)) \cap \underline{h}_2^{-1}(y_2(t+1)) \\ &= \hat{x}(t+1) \end{aligned} \quad (4.20)$$

by (4.2). Thus $R(t+1)$ is a primitive information relation. □

In general, the expansion process augments the $Z_i(t)$ with $y_i(t+1)$ and computes the conditional state set for $z_i(t)$, $z_2(t)$, $y_1(t)$, and $y_2(t)$ based on the system dynamics \underline{f} and the representation in $R(t)$ of a conditional state set for $z_1(t)$ and $z_2(t)$. Section B introduced containment between information relations to convey the notion that one was related to a substructure of another. Expansion preserves containment.

Lemma 6: If $R \subseteq R'$, then

$$F(R) \subseteq F(R') \quad (4.21)$$

Proof: Since $R \subseteq R'$, there exist functions $\phi_i: Z_i \rightarrow Z'_i$ satisfying

$$(4.10). \text{ Construct functions } \phi_i^+: Z_i \times Y_i \rightarrow Z'_i \times Y_i$$

where

$$\phi_i^+(z_i, y_i) = (\phi_i(z_i), y_i) \quad (4.22)$$

Then for all \tilde{z}_2 and \tilde{y}_2 ,

$$\begin{aligned} F(R((z_1, y_1), (\tilde{z}_2, \tilde{y}_2))) \\ = \underline{f}(R(z_1, \tilde{z}_2)) \cap \underline{h}_1^{-1}(y_1) \cap \underline{h}_2^{-1}(\tilde{y}_2) \end{aligned} \quad (4.23)$$

$$= \underline{f}(R'(\phi_1(z_1), \phi_2(\tilde{z}_2))) \cap \underline{h}_1^{-1}(y_1) \cap \underline{h}_2^{-1}(\tilde{y}_2) \quad (4.24)$$

by (4.10) and set inequalities, and this

$$= F(R')(\phi_1^+(z_1, y_1), \phi_2^+(\tilde{z}_2, \tilde{y}_2)) \quad (4.25)$$

This establishes the first half of (4.10) for ϕ^+ ; the other half is shown by a symmetric argument. Thus ϕ^+ is a homomorphism from $F(R)$ to $F(R')$; hence $F(R) \subseteq F(R')$.

□

This sets up the second major result:

Theorem 2: Let $R(t)$ and $R(t+1)$ be primitive information relations at successive times. Then

$$\text{core}(R(t+1)) = \text{core}(F(\text{core}(R(t)))) \quad (4.26)$$

Proof: By Lemma 5,

$$R(t+1) = F(R(t)) \quad (4.27)$$

From Theorem 1

$$\text{core}(R(t)) \subseteq R(t) \quad (4.28)$$

Hence applying lemma 6 to each containment in \subseteq .

$$F(\text{core}(R(t))) \subseteq F(R(t)) \quad (4.29)$$

Then by Theorem 1 and (4.27)

$$\text{core}(F(\text{core}(R(t)))) = \text{core}(R(t+1)) \quad (4.30)$$

□

Thus the cores of the primitive information relations can be computed recursively, rather than just directly. This is the analog of the fact that conditional state probability distributions for Markov processes may be computed recursively, rather than requiring reference to $I(t)$ at each time, and has the same structural implications. The size of the primitive information relations $R(t)$ generally increases exponentially with t , since $Z_i = Y_i^t$. However, the cores of the primitive information relations, $R^*(t)$, may remain much more manageable in size.

Figure 7 shows the recursive computation of the core information relations for the example of figure 3b. F denotes expansion; $*$ denotes reduction to an irreducible relation. First, notice that the reduction of $R(z)$ led to an expansion of $R^*(z)$ which was a bit simpler than the primitive $R(3)$ displayed in figure 6, so reductions are indeed cumulative over time.

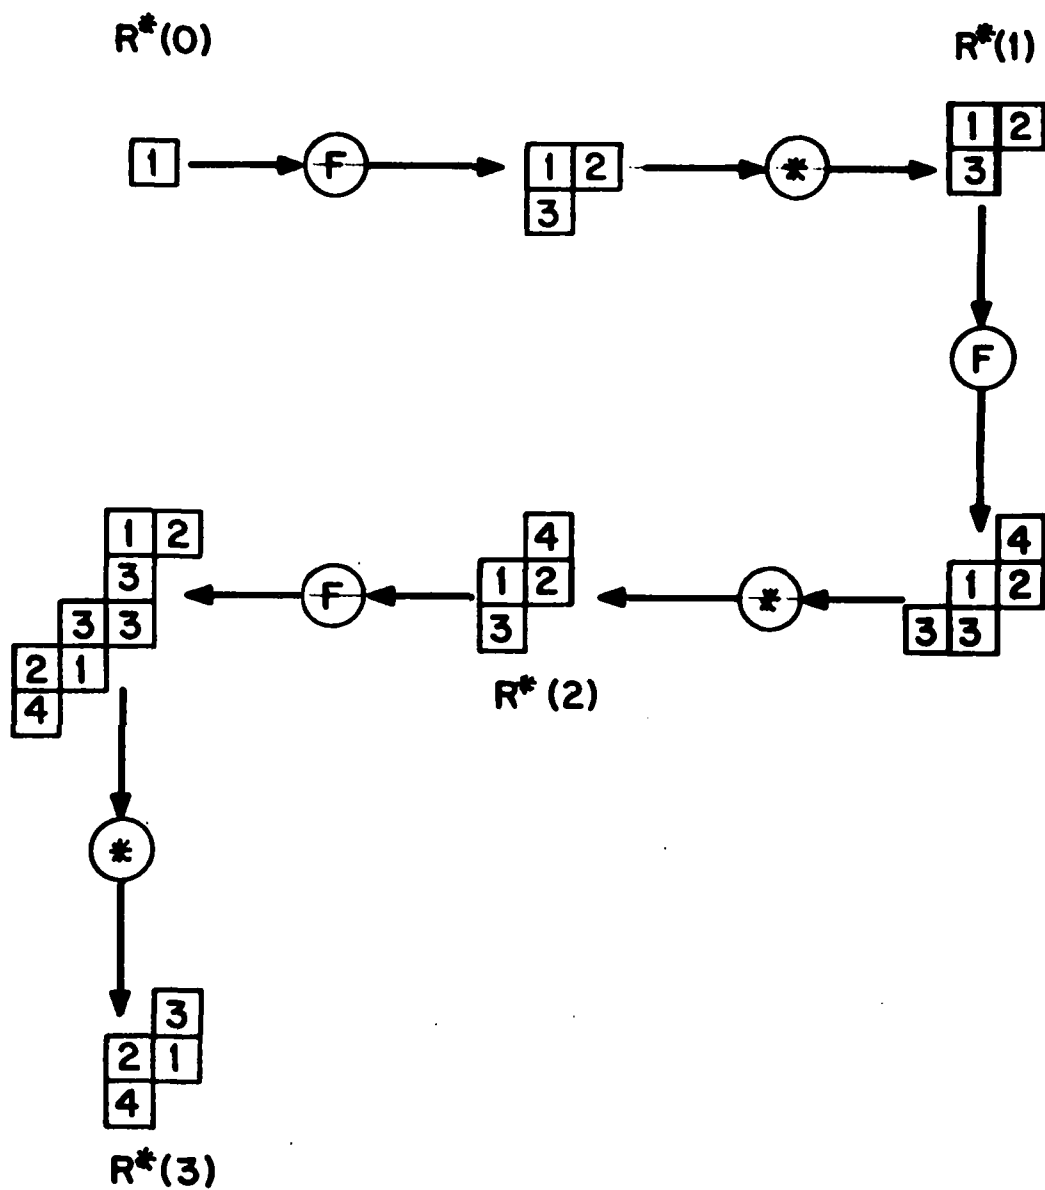


Figure 7. Recursive Derivation of Core

More importantly, notice that $R^*(3) = R^*(2)$ (isomorphically), and thus the core information relation for this system is fixed for all future time at a finite size!

Definition: The steady state core information relation R^* , if it exists, is the core of some primitive information relation $R(t)$ and satisfies

$$R^* = \text{core}(F(R^*)) \quad (4.31)$$

□

If a steady state core can be found for a system, then a great deal of the work required to solve the system is complete. However, not all systems have a steady state core; the system in figure 3c generates cores of the form shown in figure 8. (In fact, these are identical to the primitive information relations). While an irreducible information relation can be found for this system which satisfies (4.31), and which contains (but is not contained in) $R^*(t)$ for all time, it is not yet clear what the relation to the notion of a core for an infinite horizon problem is. This question is revisited in Appendix A.

A final property of the core dynamics $R^*(t+1) = \text{core}(F(R^*(t)))$ is illustrated by the system of figure 3d. $R^*(2)$ is shown in figure 9; notice that it separates into two pieces: the solitary 2 state shares no rows or column with any other entry. What is the physical reason for this?

The solitary 2 is generated by a transition from 4 to 2. This is an unusual transition in that both agents observe it unambiguously - for agent 1, it is the only transition from row 2 to row 1; for agent 2 it is the only transition within column 2. Thus its occurrence leads each agent to know exactly what happened, and that the other agent knows exactly what

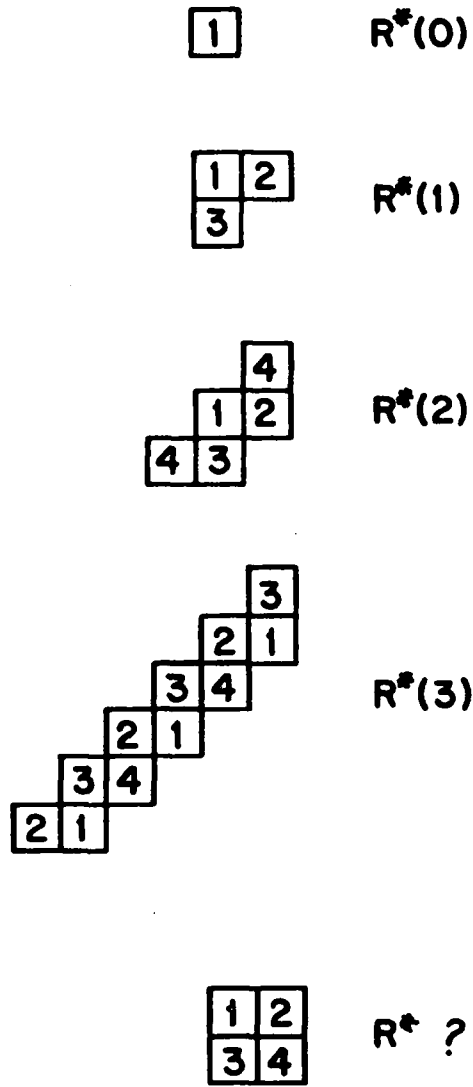


Figure 8. Cores for Example 3.c

happened, etc. This is an event which is simultaneously and unambiguously observable to both agents, and thus decouples it from any other events and second guessing logic.

Definition: An information relation R splits if the sets Z_i can be partitioned into nonempty disjoint subsets Z_i^I, Z_i^{II} with


$$\begin{aligned} R(z_1^I, z_2^{II}) &= \emptyset & \forall z_1^I \in Z_1^I, \quad \forall z_2^I \in Z_2^I \\ R(z_1^{II}, z_2^I) &= \emptyset & \forall z_1^{II} \in Z_1^{II}, \quad \forall z_2^I \in Z_2^I \end{aligned} \tag{4.32}$$

If a relation splits, then the expansion of each part of the relation can be computed separately (i.e. if R splits, then $F(R)$ splits in the compatible way). In the context of core dynamics, this means that the two pieces evolve separately. This suggests a representation for the core dynamics as shown in figure 9b - $R^*(1)$ expands and splits into two parts - one of which reduces back to $R^*(1)$, the other of which evolves separately. In this case, however, the steady state core is still that illustrated in figure 9a.


F. Local Observers

Thus far, attention has centered on off-line calculations - the reduction of primitive information relations to core relations, and the core dynamics. The ostensible purpose of this is to aid the design of on-line decision rules. The connection between the two is the dynamics of the elements of $Z_i(t)$ which are implied by the core dynamics.

The core dynamics alternately expand $Z_i(t)$ by appending Y_i , a new observation, and then reduce it via an equivalence relation implied by an automorphism ϕ_i . This equivalence relation combines



a) $R^*(2)$ for example 3d



b) Core dynamics

Figure 9: Splitting Cores

those elements of $Z_i \times Y_i$ which need not be distinguished in the future as far as the core dynamics are concerned. The entire sequence of these equivalence relations map every observation sequence y_i^t into some element of $Z_i(t)$, and thus dictate the structure of a finite state machine with Y_i as an input set and $Z_i(t)$ as the states.


Definition: The local core observer for agent i is the system with state set $Z_i(t)$ and dynamics

$$\begin{aligned} z_i(t+1) &= \hat{f}_i(z_i(t), y_i(t+1)) \\ &\triangleq \phi_i(z_i(t), y_i(t+1)) \end{aligned} \tag{4.33}$$

where the $Z_i(t)$ are the components on which $R^*(t)$ is defined, and ϕ_i is a component of the reducing automorphism used to reduce $F(R^*(t))$ to $R^*(t+1)$.


For example, consider the system of figure 10a. Its steady state core is shown in figure 10b, with elements of Z_1 and Z_2 labeled arbitrarily. Figure 10c shows the expansion $F(R^*)$, with ϕ_1 and ϕ_2 as required to reduce it back to R^* . Figure 10d is the resulting structure for the local observer of agent 1 - agent 2 has an isomorphic structure by the symmetry of the problem. Arcs are labeled with the values of y_1 causing transitions in the observer; no arc labeled 1 leaves z_{11} as this represents a case where agent 1 knows $x(t) = 2$, and 2 is followed only by 4 which gives the observation 2.

The dynamics of these observers are intertwined despite the fact that they operate on the basis of different observations. This fact shows up in the information relation on which they are based. $R^*(z_1, z_2)$ is the subset of states which the system may occupy when each agent i 's observer is in state z_i . When $R^*(z_1, z_2) = \emptyset$, it is impossible that z_1 and z_2 be occupied simultaneously - thus the observers are at least




a) System


$z_{11} z_{22} z_{23} z_{24}$



b) Steady state core



c) Expansion and reduction



d) Agent 1 observer

Figure 10: Local Steady State Core Observers

partially synchronized.

Other insight as to the meaning of the local observer states can be gained from noting that $R^*(z_1, z_2)$ is the complete set of states which the system may occupy if agent 1's observer is in state z_1 . It is the union of the sets of states which the system could occupy given each y_1^t which leads to z_1 . Thus the state sets corresponding to each z_1 are related to the local conditional state sets defined in (4.3). However, the relation is rather subtle, as both z_{12} and z_{14} indicate that $x \in \{3,4\}$; they must remain distinct due to the existence of agent 2 with different information.

Thus the core dynamics define some automata for processing local observations in a way which maintains the relationships described in R^* .

G. Special Cases

There are certain nondeterministic systems which have more structure than the general case discussed thus far. For example, a system yields the equivalent of centralized information if $y_1(t) = y_2(t)$ at all times. This special case appears in $R(t)$ in that it is diagonal.

$$R(t)(y_1^t, y_2^t) = \begin{cases} \emptyset & \text{if } y_1^t \neq y_2^t \\ \hat{x}(t) & \text{else} \end{cases} \quad (4.34)$$

The local core observers, in this case, are perfectly synchronized and each implements a reduced order realization of the conditional state set computation (4.2). The realization is reduced order as any reachable

conditional state set which is contained in another is treated as equivalent to that other. Thus only the maximal (in the sense of set containment) reachable conditional state sets are tracked.

A slightly more general case is where:

Definition: The observations of agent 1 are included in those of agent 2 if $y_1(t) = h_{21}(y_2(t))$, for some h_{21} . □

Thus if $y_2(t)$ is a vector, $y_1(t)$ may be one component of that vector.

In this case, the information relations have the property that for any z_2 , there is exactly one z_1 for which $R(z_1, z_2) \neq \emptyset$. This is true for the primitive information relations since $z_2 = y_2^t$ uniquely determines the y_1^t that will arise. This property is preserved when $R(t)$ is reduced by any automorphism; proof of this fact is deferred. Figure 11 illustrates such an information relation.

In this case, the local core observer for agent 1 turns out to just be a reduced order realization of the local conditional state set computation. The local core observer for agent 2 is a reduced order realization of the conditional state set for the combined system and local core observer for agent 1. This is the structure demonstrated in the example in [15].

A final special case has one agent, say 1, operating open loop without observations. In this case, set

$$y_1(t) = 0 = h_{12}(y_2(t)) \tag{4.35}$$

to see that it is a variation of the above case. However, now Z_1 always has exactly one element; $z_1(t)$ indicates the steady state conditional state set (generally X unless part of the state space is not reachable), and $z_2(t)$ realizes a reduced order computation of $\hat{x}(t)$.

	z_{21}	z_{22}	z_{23}	z_{24}	z_{25}	z_{26}
z_{11}	{1}	{2,4}				
z_{12}			{1,3}	{1,4}	{2}	
z_{13}						{3}

Figure 11. Example Information Relation with Included Observations

H. Summary

This section has developed the notion of an information relation, and some of the properties thereof. The properties all stem from the definition of information relation containment via homomorphisms (4.10). The basic premise, as yet unjustified, has been that this definition is useful in providing the structure needed to solve decentralized nondeterministic decision problems. The structure resulting from this definition, particularly the core, is now known; the next section shows what it can do in a subset of the problems of interest.

V. DECENTRALIZED ESTIMATION

This section continues the assumption of the previous one that the system dynamics are autonomous, but introduces the objective function in order to derive decision rules. Since the cost function J compares agents' decisions with the true state, these will be referred to as decentralized estimation problems. They can be solved by (a) reducing them to a sequence of static problems of the type addressed in section III, (b) relating the static problems to the information relation and its structure, and (c) combining the resulting decision rules with the observer dynamics necessary to compute sufficient statistics on-line. Some special cases will be of particular interest.

A. Reduction to Static Problems

The general decentralized estimation problem is to find the

$$\min_{\Gamma_1^T} \max_{\Gamma_2^T} \max_{t \in \{1, \dots, T\}} \max_{x(t) \text{ possible}} J(x(t), u_1(t), u_2(t)) \quad (5.1)$$

where $x(t)$ is possible if $x(t) \in \underline{f}^t(x^0)$, (the t -fold composition of \underline{f} with itself). The information restriction

$$u_i(t) = \gamma_i(t) (y_i^t) \quad (5.2)$$

still applies.

Define

$$\tilde{J}(\gamma_1(t), \gamma_2(t)) = \max_{\left. \begin{array}{l} x(t) \\ y_1^t \\ y_2^t \end{array} \right\} \text{possible}} J(x(t), \gamma_1(x_1^t), \gamma_2(y_2^t)) \quad (5.3)$$

so (5.1) becomes

$$\begin{aligned} & \min_{\Gamma_1^T, \Gamma_1^T} \max_t \tilde{J}(\gamma_1(t), \gamma_2(t)) \\ & = \max_t \min_{\gamma_1(t), \gamma_2(t)} \tilde{J}(\gamma_1(t), \gamma_2(t)) \end{aligned} \quad (5.4)$$

It applies in this since each component of Γ_i^T may be chosen separately. case since decisions do not affect dynamics or costs other than that a single time; \tilde{J} depends only on decision rules at one time. Thus

Lemma 7: The solution to the decentralized estimation problem may be found by solving the sequence:

$$\text{minimize}_{\gamma_1(t), \gamma_2(t)} \tilde{J}(\gamma_1(t), \gamma_2(t)) \quad (5.5)$$

Proof: (5.4)

B. Use of the Information Relation

The important structural quantity in \tilde{J} is the characterization of the set of possible triples $(x(t), y_1^t, y_2^t)$. This is a role which the information relation can play.

Definition: The optimal value of an information relation R in a non-deterministic decentralized estimation problem, denoted $J^*(R)$, is

$$\min_{\gamma_1, \gamma_2} \max_{\substack{z_1 \in Z_1 \\ z_2 \in Z_2 \\ x \in R(z_1, z_2)}} J(x, \gamma_1(z_1), \gamma_2(z_2)) \quad (5.6)$$

with Z_1 and Z_2 the sets on which R is defined, and the restriction

$$\gamma_i : Z_i \rightarrow U_i \quad \square \quad (5.7)$$

Thus if R is a primitive information relation, $J^*(R)$ is the optimal value of (5.5), since $Z_i = Y_i^t$ and R specifies the set of $x(t)$ possible (perhaps empty) for each (y_1^t, y_2^t) . Naturally we are interested in simpler information relations; the following lemma establishes the connection between information relations and their optimal values.

Lemma 8: If R and R' are information relations, then

$$R \subseteq R' \Rightarrow J^*(R) \leq J^*(R') \quad (5.8)$$

Proof: If $R \subseteq R'$, then a homomorphism $\phi = (\phi_1, \phi_2)$ exists from R to R'.

Consider any strategy $\gamma' = (\gamma'_1, \gamma'_2)$ with

$$\gamma'_i : Z'_i \rightarrow U_i \quad (5.9)$$

Build a strategy $\gamma = (\gamma_1, \gamma_2)$ with

$$\gamma_i(z_i) = \gamma'_i(\phi_i(z_i)) \quad (5.10)$$

Then

$$\begin{aligned} & \max_{\substack{z_1 \\ z_2 \\ x \in R(z_1, z_2)}} J(x, \gamma_1(z_1), \gamma_2(z_2)) & (5.11) \\ & = \max_{\substack{z_1 \\ z_2 \\ x \in R(z_1, z_2)}} J(x, \gamma_1'(\phi_1(z_1)), \gamma_2'(\phi_2(z_2))) \end{aligned}$$

$$\begin{aligned} & \leq \max_{\substack{z_1' \\ z_2' \\ x \in R'(z_1', z_2')}} J(x, \gamma_1'(z_1'), \gamma_2'(z_2')) & (5.12) \end{aligned}$$

where the inequality stems from the fact that

$$\phi_i(z_i) \subseteq z_i' \quad (5.13)$$

$$R(z_1, z_2) \subseteq R'(\phi_1(z_1), \phi_2(z_2))$$

from the definition of a homomorphism. Since (5.12) holds for all γ' , evaluate it at a γ' which achieves $J^*(R')$. The corresponding γ defined in (5.10), must achieve a value of (5.11) no larger than $J^*(R')$, and the best γ is at least as good as this one, so the conclusion holds. This immediately suggests the third major result:

Theorem 3: Let $R(t)$ be a primitive information relation and $R^*(t)$ its core.

Then

$$J^*(R(t)) = J^*(R^*(t)) \quad (5.14)$$

Proof: From theorem 1, $R(t) \subseteq R$. Applying Lemma 8 in both directions, the conclusion follows. □

This theorem implies that the pair of decision rules γ_1, γ_2 which achieve the optimum value of (5.6) when R is a primitive relation perform exactly as well as the rules optimizing (5.6) for the core of that relation. Since the former are functions on Y_i^t , they are more cumbersome than the latter, which are functions of only the Z_i of the core. Since each agent can compute the value of $z_i(t)$ representing the situation given an observation sequence y_i^t by using its local core observer, the solution to (5.6) with $R^*(t)$ can be implemented on line. Thus the complexity of the problem can be significantly reduced from that of (5.5).

C. General Solution

The procedure for solving the nondeterministic optimal decentralized estimation problem is thus:

1. Compute the core dynamics (4.18 - 4.19, 4.11)

$$\begin{aligned} R^*(t+1) &= \text{core } (F(R^*(t))) \\ R^*(0) &= \{R(\{z_1^0\}, \{z_2^0\}) = \{x^0\}\} \end{aligned} \tag{5.15}$$

2. Derive from these the local core observers (4.33)

$$z_i(t+1) = \hat{f}_i(z_i(t), y_i(t+1)) \tag{5.16}$$

3. Find the core decision rules $\gamma_i^* : Z_i \rightarrow U_i$ minimizing

$$\begin{aligned} \max_{\substack{z_1 \\ z_2 \\ x \in R^*(t)(z_1, z_2)}} J(x, \gamma_1(z_1), \gamma_2(z_2)) \end{aligned} \tag{5.17}$$

4. Implement the local core observers and decision rules

$$u_1(t) = \gamma_1^*(z_1(t)) \tag{5.18}$$

Thus a recursive solution can be found based on the structure of the core information relation. The solution of (5.17) can be obtained as described in section III by exploiting the PBPO necessary conditions and finding the best pair (γ_1^*, γ_2^*) which satisfy them. If a steady state core exists, then the dynamics of the local core observers, and the structure of the decision rules, become stationary.

Figure 12 illustrates the resulting solution structure. Note the decomposition of each agent's processing into an estimator with memory preceding a memoryless decision rule. The estimators may be designed independently of the cost function J , but must be derived jointly. Similarly, the decision rules are derived jointly, but only in response to those dynamics retained in the core.

D. Special Cases

Three special cases of information relations were identified in the last section. Their impact on the estimation rules is summarized here.

1. **Splitting Cores:** If R^* splits, then the design of the decision rules for each piece can be done separately. If $Z_1^i, Z_1^ii, Z_2^i,$ and Z_2^{ii} are the partitions on Z_1 and Z_2 which support the split, and R^i and R^{ii} are the restrictions of R to the appropriate subsets, then

$$J^*(R) = \max \{J^*(R^i), J^*(R^{ii})\} \quad (5.19)$$

and the optimization over R decomposes.

2. **Centralized information:** When $y_1(t) = y_2(t)$, R^* becomes diagonal as $R(y_1^t, y_2^t) = \emptyset$ unless $y_1^t = y_2^t$. This is an extreme case of splitting as the value of the decision rule for each case of $z_1 = z_2$ can be derived separately.

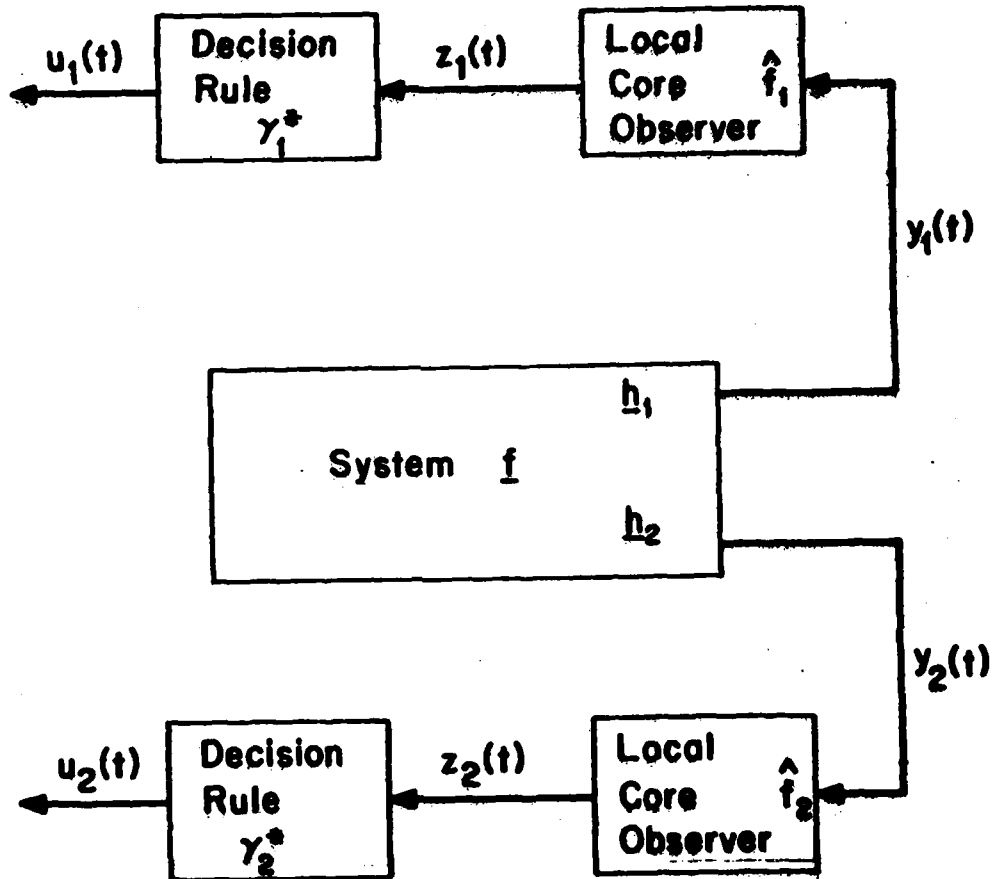


Figure 12. Optimal Decentralized Estimation Structure

3. Included information: When $y_1(t) = h_{12}(y_2(t))$, so agent 2 knows agent 1's information exactly, then R^* has a structure where, for every z_2 , there is exactly one z_1 for which $R^*(z_1, z_2) \neq \emptyset$. Again this is a case of splitting; decision rules can be found via:

- a) For each z_1 and u_1 , find

$$\gamma_2(z_2 | z_1, u_1) = \operatorname{argmin}_{u_2} \max_{x \in R(z_1, z_2)} J(x, u_1, u_2) \quad (5.20)$$

and define the value of this minimum as $V_2(z_1, u_1)$

- b) For each z_1 , set

$$\gamma_1^*(z_1) = \operatorname{argmin}_{u_1} V_2(z_1, u_1) \quad (5.21)$$

- c) Set

$$\gamma_2^*(z_2) = \gamma_2(z_2 | z_1, \gamma_1^*(z_1)) \quad (5.22)$$

This is possible because each z_2 determines the z_1 required in (5.22) uniquely.

This is a recursive procedure taking advantage of the nesting of the information of agent 1 with that of agent 2. Interpreted as a Stackelberg solution to a team problem, this implies that agent 1, with the reduced information, is a leader and agent 2, with more information, is a follower. This rather curious connection between Stackelberg solutions and included observations arises naturally.

VI. THE CONTROL PROBLEM

This section finally relates the information relation and its properties to the full control problem defined in section II. Additional complexity is generated by the fact that information relations can only be expanded in the context of specific decision rules. This sets up some deterministic (unless splitting occurs) core dynamics which are controlled by selections of decision rules. Viewing cores as states and decision rules as inputs, a dynamic programming approach leads to the complete solution.

A. Information Dynamics

The definitions of information relation, homomorphisms, automorphisms, containment, and cores established in section IV carry over to the control problem without change. The expansion process is the only place where new information relations are generated, so the influence of decisions on dynamics requires a modification there.

Definition: The expansion of an information relation R by decision

rules $\gamma_i : Z_i \rightarrow U_i$ is an information relation R' , where

$$R' = F(R, \gamma_1, \gamma_2) \quad (6.1)$$

$$R' : Z_1' \times Z_2' \rightarrow 2^X \quad (6.2)$$

$$Z_i' = Z_i \times Y_i \quad (6.3)$$

$$R'((z_1, y_1), (z_2, y_2)) = \underline{f}(R(z_1, z_2), \gamma_1(z_1), \gamma_2(z_2)) \cap \underline{h}_1^{-1}(y_1) \cap \underline{h}_2^{-1}(y_2) \quad (6.4)$$

□

This is interpreted as follows. If R is information relation at a system at time t , then $F(R, \gamma_1, \gamma_2)$ is information relation between two agents observing the system via y_1 and y_2 while it autonomously evolves for one time space step under the decisions generated by γ_1 and γ_2 . (6.4) is the update for the conditional state set given $R(z_1, z_2)$ as the prior conditional state set, decisions generated by γ_1 and γ_2 and new observations y_1 and y_2 . Thus R' will describe the information relation between the agents if γ_1 and γ_2 are implemented as decision rules.

Only one new concept is needed.

Definition: Let ϕ be a homomorphism from R to R' . Then a decision rule $\gamma_i : Z_i \rightarrow U_i$ is contained in $\gamma_i' = Z_i' \rightarrow U_i'$, denoted $\gamma_i \subseteq \gamma_i'$, if

$$\gamma_i(z_i) = \gamma_i'(\phi_i(z_i)) \quad \forall z_i \in Z_i \quad (6.5)$$

□

With this notion, the results of section IV generalize to the control case as:

Theorem 4: Let R and R' be information relations, and γ_i, γ_i' decision rules of the appropriate structure.

a) If R is a primitive information relation, then so is

$$F(R, \gamma_1, \gamma_2)$$

b) If $R \subseteq R'$, $\gamma_1 \subseteq \gamma_1'$, and $\gamma_2 \subseteq \gamma_2'$, all by the same homomorphism, then

$$F(R, \gamma_1, \gamma_2) \subseteq F(R', \gamma_1', \gamma_2') \quad (6.6)$$

c) If R is a primitive information relation, then

$$\text{core}(F(R, \gamma_1, \gamma_2)) = \text{core}(F(\text{core}(R), \gamma_1, \gamma_2)) \quad (6.7)$$

Proofs: All proofs are direct extensions of

- a) Lemma 5
- b) Lemma 6
- c) Theorem 2, using (a) and (b). \square

As in the autonomous case, (c) is the critical fact. The core information relations can be viewed as the state of a new process with dynamics (6.7) and control (γ_1, γ_2) . (If (6.7) produces a new core which splits, then the dynamics must be viewed as nondeterministic, as different sample paths in the original process may lie in either of the two pieces). These dynamics establish half of a new problem structure.

B. Costs

The other half involves the objective. Properties developed in Section V for the estimation problem also generalize to the control problem.

Definition: The cost of an information relation R with compatible decision rules γ_1, γ_2 is

$$J(R, \gamma_1, \gamma_2) = \max_{\substack{z_1 \\ z_2 \\ x \in R(z_1, z_2)}} J(x, \gamma_1(z_1), \gamma_2(z_2)) \quad (6.8)$$

\square

The overall problem objective (2.13) can be written as

$$J(\Gamma_1, \Gamma_2) = \max_{t \in \{1, \dots, T\}} \left. \begin{array}{l} \max_{x(t)} \\ \max_{y_1^t} \\ \max_{y_2^t} \end{array} \right\} \text{possible } J(x(t), \gamma_1(y_1^t), \gamma_2(y_2^t)) \quad (6.9)$$

Given y_1^t and y_2^t , the only $x(t)$'s possible are given by $R(t)(y_1^t, y_2^t)$, the primitive information relation for time t constructed recursively from (6.4) (Theorem 4a guarantees such an $R(t)$ is primitive). (6.9) then becomes

$$J(\Gamma_1, \Gamma_2) = \max_{t \in \{1, \dots, T\}} \max_{\substack{y_1^t \\ y_2^t \\ x \in R(y_1^t, y_2^t)}} J(x(t), \gamma_1(y_1^t), \gamma_2(y_2^t)) \quad (6.10)$$

$$= \max_{t \in \{1, \dots, T\}} J(R(t), \gamma_1(t), \gamma_2(t)) \quad (6.11)$$

where the dependence of $R(t)$ on prior decision rules is left implicit. Thus the overall objective can be written in terms of costs of primitive relations with decision rules.

The final result needed is:

Theorem 5: Let R and R' be information relations, with γ_1, γ_2 and γ_1', γ_2' compatible decision rules. Then

a) If $R \subseteq R'$, $\gamma_1 \subseteq \gamma_1'$, and $\gamma_2 \subseteq \gamma_2'$ by the same homomorphism ϕ , then

$$J(R, \gamma_1, \gamma_2) \leq J(R', \gamma_1', \gamma_2') \quad (6.12)$$

- b) If $R = \text{core}(R')$, $\gamma_1 \subseteq \gamma_1'$, and $\gamma_2 \subseteq \gamma_2'$ by the same homomorphism ϕ , then

$$J(R, \gamma_1, \gamma_2) = J(R', \gamma_1', \gamma_2') \quad (6.13)$$

Proof: Also direct extensions of previous results.

- a) Lemma 8
 b) (a) with theorem 1. □

Applying (6.13) to (6.11), the objective becomes

$$J(\Gamma_1, \Gamma_2) = \max_{t \in \{1, \dots, T\}} J(\text{core}(R(t)), \gamma_1'(t), \gamma_2'(t)) \quad (6.14)$$

where γ_1' and γ_2' are the restrictions of γ_1 and γ_2 to $\text{core}(R(t))$:

C. General Solution

Let $R^*(t) = \text{core}(R(t))$, where $R(t)$ is the primitive information relation created by decision rules prior to time t . From theorem 4c,

$$\begin{aligned} R^*(t+1) &= \text{core}(F(R^*(t), \gamma_1(t), \gamma_2(t))) \\ &\stackrel{\Delta}{=} F^*(R^*(t), \gamma_1(t), \gamma_2(t)) \end{aligned} \quad (6.15)$$

and the overall objective is

$$J(\Gamma_1, \Gamma_2) = \max_{t \in \{1, \dots, T\}} J(R^*(t), \gamma_1(t), \gamma_2(t)) \quad (6.16)$$

where Γ_1 and Γ_2 are now sequences of decision rules whose arguments are in $Z_1(t)$ and $Z_2(t)$, respectively. Solution of this problem would involve straightforward minimax dynamic programming if the set in which the $R^*(t)$ could lie were determinable ahead of time. Unfortunately, it is not known how to do this at this point; however, one may construct the set of all cores

reachable under all Γ_1, Γ_2 from $R^*(0)$ and take this to be the requisite set.

This generates the following procedure for solving the decentralized optimal control problem for finite nondeterministic automata:

1. Using (6.15) for all possible $\gamma_i : Z_i(t) \rightarrow U_i$, construct the set \mathbb{R} of reachable cores.

2. Dynamic program:

a) Set

$$V(R, T) = \min_{\gamma_1, \gamma_2} J(R, \gamma_1, \gamma_2) \quad (6.17)$$

for each $R \in \mathbb{R}$.

b) Given $V(R, t+1)$ for every $R \in \mathbb{R}$, compute

$$V(R, t) = \min_{\gamma_1, \gamma_2} \max \{ J(R, \gamma_1, \gamma_2), \quad (6.18)$$

$$V(F^*(R, \gamma_1, \gamma_2), t+1) \}$$

(if $F^*(R, \gamma_1, \gamma_2)$ splits, include all pieces separately in (6.18)).

c) After the recursion (b) ends, determine the actual cores $R^*(t)$

and optimal decision rules $\gamma_1^*(t), \gamma_2^*(t)$ reached from $R^*(0)$,

using (6.15) and deriving γ_1^*, γ_2^* as the arguments achieving

a minimum in (6.18)

3. Implement:

a) Construct the local observers from the dynamics of $Z_1(t)$ and $Z_2(t)$ along the optimal trajectory of cores,

b) Implement $\gamma_1^*(t)$ and $\gamma_2^*(t)$ as memoryless feedback maps from $Z_1(t)$ and $Z_2(t)$ to decision sets.

There are several structural points to note about this solution. First, the implementation is of the same structure as that of figure 12, except both decisions are now fed back into the system. Thus each agent consists of an estimator followed by a memoryless decision rule. As before, the local estimators must be designed jointly; however, their structure is highly dependent on the decision rules of both agents. There is no separation principle in this class of decentralized control problems: the structure of the estimator for agent 1 depends on the decision rules which agent 2 uses, as knowledge of these impacts on how agent one "second guesses" the activities of agent 2 and their impact on the system. Finally, the least well understood step is step 1 - for many problems R may be reasonably small; however, problems for which one set of strategies generates dynamics such as those of figure 3c, increasing over time without limit, may yield unmanageably large sets of R after a few steps. The importance of this issue is thus highly dependent upon the problem.

As a final note, it is easy to verify by induction that: If $R \subseteq R'$, then

$$V(R,t) \leq V(R',t) \quad \forall t \in \{1, \dots, T\}$$

VII. CONCLUSIONS

Having achieved the solution of the general control problem, several lessons can be learned from this exercise.

A. Summary

The three basic constructs introduced were the information relation, the notation of containment of one relation in another, and the process of reducing one into another. These are the decentralized generalizations

of the constructs required for a centralized treatment of this problem: sets, set containment, and the elimination of redundant elements in set unions (in the computation of the conditional state set). The dynamics of the information relation were derived. The facts that the containment relation is preserved by these dynamics, and that it leads to a compatible inequality on the cost of an information relation, are what establish the role of irreducible information relations (the core) in the problem. Reformulation of the control problem in terms of the core yields a standard dynamic programming problem.

The core exists as a simple description of one agent's information set to both the system state and to other agents' information. The fact that it is simpler than the primitive information relation is strongly dependent upon the nondeterministic system model and the minimax cost objective. This is most evident in the decentralized estimation case. At each point, the static problem which each agent must solve is originally posed in terms of a primitive information relation, then simplified to one on the core. Consider the relation of figure 13. z_{21} can be mapped onto z_{11} by a reducing automorphism. This implies that we may set $\gamma_1(z_{21}) = \gamma_1(z_{11})$ without loss of optimality. Let γ_2 be any decision rule on $\{z_{21}, z_{22}\}$. The above equality assures that the event $(x=1, z_{11}, z_{21})$ will be penalized identically as would $(x=1, z_{12}, z_{21})$, and this cannot increase the value of the worst case. However, it may well be that

$$\min_{\gamma_1(z_{11})} \max_{x \in \{1,2\}} J(x, \gamma_1(z_{11}), \gamma_2(z_{21})) > \min_{\gamma_1(z_{12})} \max_{x \in \{1\}} J(x, \gamma_1(z_{12}), \gamma_2(z_{11}))$$

(7.1)

These two could not be treated as equal with other objective structures,

	z_{21}	z_{22}
z_{11}	1	2
z_{12}	1	

Figure 13. A Simple, Reducible Information Relation

and thus one could not consider them equivalent and the information relation could not be reduced.

Computationally, the information relation simplifies some problems, often significantly. However, two difficulties remain: the core may always grow in size without bound, and the number of possible decision rules which may be defined on a large information relation, and which defines the size of the control space in the reformulated problem, may be quite large. These difficulties may be resolved through deeper understanding of the role of the state core or through branch and bound techniques, but these are open questions.

B. Conclusions

Whether these results may suggest important new approaches to practical decentralized control design, or whether they exist only as an interesting special case, is only a matter of speculation now. However, they allow two principle conclusions to be drawn. First, the information relation is a richer structure than a set, and essential to this problem. Extension of these results to probabilistic models, with which we are more familiar, will require a similar generalization of probabilistic structures to the decentralized case. However, behind every probability measure lies a σ -algebra, so part of the work may be done.

Second, there are some problems which can be cast in the model and objective structure assumed here, and they can now be addressed directly. Events which can be expected to happen sometimes, but not always, can be modeled by nondeterministic structures which allow them to occur, say,

at least three, but no more than five, times out of any eight sequential time steps. By this type of remodeling, the apparent limitations of the formulation can be partially overcome.

For example, consider the problem of designing protocols for decentralized management of a multiaccess channel [18]. A source of messages may produce them no more than N_M out of any N_T consecutive time steps. They may appear at any one of M transmitting sites. There they are stored in a queue, and the agent at each site must decide whether to attempt to transmit one message or to wait, at each time step. It has knowledge of its queue length and past activity on the channel. Depending on the number of agents transmitting simultaneously, the channel may be clear, busy (one attempt, successful) or jammed (several attempts, unsuccessful). A successful attempt reduces a queue length by one. The objective function penalizes the maximum length of the queue at any site, at any time, over all possible message production patterns. This suggests a relatively simple nondeterministic, finite state model with a minimax objective - and is thus suitable for solution by the techniques given here. Variation of the structure of the solution as N_M , N_T , and M are changed would be quite interesting.

References

- [1] N.R. Sandell, Jr., P. Varaiya, M. Athans, and M.G. Safonov, "A Survey of Decentralized Control Methods for Large Scale Systems", IEEE Trans.A.C. Vol. AC 23, No. 2, April 1978.
- [2] R. Radner, "Team Decision Problems", Ann. Math. Stat., Vol. 23, No.3, pp. 857-881, 1962.
- [3] H. W. Kuhn, "Extensive Games and the Problem of Information", Contributions to the Theory of Games II, Princeton University Press, Princeton, N.J., 1953.
- [4] Y.C. Ho and K.C. Chu, "Team Decision Theory and Information Structures in Optimal Control Problems, Part I", IEEE Trans. A.C., Vol. AC-17, p.p. 15-22, Feb. 1972.
- [5] T. Basar, "Equilibrium Strategies in Dynamic Games with Multilevels of Hierarchy", Automatica, Vol. 17, No. 5, 1981.
- [6] C.Y. Chong and M. Athans, "On the Periodic Coordination of Linear Stochastic Systems", Automatica, Vol. 12, July 1976.
- [7] R.R. Tenney and N.R. Sandell, Jr., "Detection with Distributed Sensors", IEEE Trans. AES, Vol. AES-17, No. 4, July 1981.
- [8] L.K. Ekchian and R.R. Tenney, "Detection Networks", Proc. 21 C.D.C., Orlando, FLA, Dec. 1982.
- [9] D. Teneketsis, "The Decentralized Quickest Detection Problem", Proc. 21 C.D.C., Orlando, FLA, Dec. 1982.
- [10] H. Witsenhausen, "A Counterexample in Stochastic Optimal Control", SIAM J. Control, Vol. 6, No. 1, 1968.
- [11] M. Athans and N.R. Sandell, Jr., "Solution of Some Nonclassical LQG Stochastic Decision Problems", IEEE Trans. A.C., Vol. AC-19, pp.108-116, April 1974.
- [12] F.C. Schweppe, Uncertain Dynamical Systems, Prentice-Hall, Englewood Cliffs, N.J. 1973.
- [13] D.P. Bertsekas, "Control of Uncertain Systems with a Set-Membership Description of the Uncertainty", Ph.D. thesis, MIT, Cambridge, MA., 1971.
- [14] H. Witsenhausen, "Minimax Control of Uncertain Systems", Ph.D. Thesis MIT Dept. EECS, 1966.
- [15] H. Sira-Ramirez, "On the Robust Control of Systems to Set-Valued Objectives", to appear IEEE Trans. A.C.
- [16] R.R. Tenney, "Sufficient Statistics for Decentralized Estimation", to appear, IEEE Trans.A.C.

- [17] D.P. Bertsekas, Dynamic Programming and Stochastic Optimal Control, Academic Press.
- [18] L.S. Bobrow and M.A. Arbib, Discrete Mathematics, Saunders, Phil. Pa.1974.
- [19] J. Capetekanasis, "Tree Algorithms for Packet Broadcast Channels," IEEE Trans. I.T., Vol. 25, No. 5, Sept. 1979.

APPENDIX A

Steady State Cores: Nonexistence

This appendix demonstrates a system which does not have a steady state core for a decentralized estimation problem. The dynamics and observation structure are those of figure 3c, so the primitive information relations are themselves irreducible (figure 8). Let the decision sets be $U_1 = U_2 = [0,1]$. Construct a cost functional

$$J(u_1, u_2) = \begin{cases} \bar{\delta}(u_1, u_2) & x \in \{1,2,3\} \\ \bar{\delta}(u_1, [u_2 - \epsilon]_{\text{mod } 1}) & x = 4 \end{cases} \quad (\text{A-1})$$

where $\bar{\delta}(x,y) = 0$ iff $x = y$, and where ϵ is any irrational number.

First, we can construct decision rules for this problem which achieve a worst-case cost of 0. Clearly no rules exist which give better worst-case performance. Consider any information relation of the form illustrated in Figure 8, say $R(3)$. If z_1 is the local observer state of agent 1 corresponding to the uppermost row, set $\gamma_1(z_1) = \alpha \in [0,1]$ arbitrarily. Now proceed recursively: assume some cell has one decision rule corresponding to it already defined (i.e. for the cell at (z_1, z_2) , either $\gamma_1(z_1)$ or $\gamma_2(z_2)$ is specified). Construct the other by setting

$$\gamma_1(z_1) = \begin{cases} \gamma_2(z_2) & \text{if } R(z_1, z_2) \in \{1,2,3\} \\ [\gamma_2(z_2) - \epsilon]_{\text{mod } 1} & \text{if } R(z_1, z_2) = 4 \end{cases} \quad (\text{A-2})$$

if $\gamma_2(z_2)$ is defined, and

$$\gamma_2(z_2) = \begin{cases} \gamma_1(z_1) & \text{if } R(z_1, z_2) \in \{1, 2, 3\} \\ [\gamma_1(z_1) + \epsilon]_{\text{mod } 1} & \text{if } R(z_1, z_2) = 4 \end{cases} \quad (\text{A-3})$$

if it is $\gamma_1(z_1)$ that is defined. This introduces a new cell with only one component defined, or terminates the recursion. Note that aside from the choice of α , these decision rules are unique in achieving $J^* = 0$. Note also that the values of the decision rules are of the general form

$$\gamma_i(z_i) = [\alpha \pm n \epsilon]_{\text{mod } 1} \quad n \text{ integer} \quad (\text{A-4})$$

Now assume that a reducing homomorphism ϕ exists for an information relation of this type which produces a (strictly simplified) core. Then there must be a pair (z_i, z_i') such that $z_i \neq z_i' = \phi_i(z_i)$; otherwise ϕ would be an isomorphism. By theorem 3, the value of $R(3)$ must equal that of its reduction. From above, that value is zero. Also from above, the decision rules which achieve that value must be of the form

$$\begin{aligned} \gamma_i(z_i) &= [\alpha \pm n_i \epsilon]_{\text{mod } 1} \\ \gamma_i(z_i') &= [\alpha \pm n_i' \epsilon]_{\text{mod } 1} \end{aligned} \quad (\text{A-5})$$

Moreover, since $z_i \neq z_i'$, the recursive procedure for generating these rules must have at least one step of the form

$$\gamma_j(\cdot) = [\gamma_j(\cdot) \pm \epsilon]_{\text{mod } 1} \quad (\text{A-6})$$

between the derivation of $\gamma_i(z_i)$ and $\gamma_i(z_i')$. Thus $n_i \neq n_i'$. From the proof of Lemma 8, we must be able to set

$$\gamma_i(z_i) = \gamma_i(z_i') \quad (\text{A-7})$$

and preserve optimality. This occurs iff

$$[(n_i - n_i')\epsilon]_{\text{mod } 1} = 0 \quad (\text{A-8})$$

i.e. iff the quantity in brackets is an integer. Since ϵ may be irrational, we have a contradiction. All primitive information relations for the system of Figure 3c must be irreducible; hence the core of the system is unbounded in size; hence no finite steady state core exists.

3-8

DT