

Decentralized Stochastic Decision Problems and Polynomial Optimization

Randy Cogill¹

Sanjay Lall²

Abstract

In this paper we consider the problem of computing decentralized control policies in a discrete stochastic decision problem. For the problem we consider, computation of optimal decentralized policies is NP-hard. We present a relaxation method for this problem which computes suboptimal decentralized policies as well as bounds on the optimal achievable value. We then show that policies computed from this relaxation are guaranteed to be within a fixed bound of optimal. The relaxation is derived from an equivalent formulation of this decentralized decision problem as a polynomial optimization problem. The method is illustrated by an example of decentralized detection.

1 Introduction

Decentralized decision problems are optimization problems in which a collection of decisions are made in response to a set of observations with the goal of maximizing some cost. The complicating factor is that decisions can only be made to depend on some specified subset of the observations. That is, the complete set of observations can be thought of as the state of the environment. Each decision is made on the basis of an incomplete observation of the state, although the cost incurred depends on the entire state and set of decisions. Such problems are common in engineering and economics. Much of the early work on team decision problems was motivated by economic problems [4]. In certain engineering problems, such as the design of distributed detection schemes and distributed data transmission protocols, the key difficulty lies in the design of good rules for interacting decision makers to follow.

Here we consider a fairly general discrete version of this problem, where the sets of possible observations and decisions are finite. The problem considered is a *static* decision problem, where a single set of decisions is

made in response to a single set of observations. Given the probabilities of all sets of observations, the goal is to choose decentralized decision rules which maximize the expected cost. This problem is shown in [9] to be NP-hard, even for the case of two decision makers. Therefore, the goal of this paper is to determine effective methods of computing good *suboptimal* solutions to this problem. Here we show that this problem can be equivalently formulated as a maximization of a polynomial subject to linear constraints. Relaxations of this polynomial optimization problem can then be efficiently solved. From these relaxations, we obtain upper bounds on the maximum achievable value for the original problem, as well as suboptimal decision rules. We also show that policies computed from this relaxation are guaranteed to be within a fixed bound of optimal.

2 Previous work

Much of the previous work on decentralized decision problems can be roughly categorized as complexity results, tractable special cases, and applications. Some of the earliest work on decentralized decision problems is the work of Radner and Marschak [4, 6]. Along with introducing the team decision problem, they have shown that for concave differential costs, person-by-person optimality is sufficient for global optimality. A nice survey of the early work in the field of decentralized decision problems, including extensions to dynamic problems, can be found in [1]. In [9], it is shown that the general static decentralized decision problem with finite state and action spaces is NP-hard.

A great deal of work on the team decision problem has been done for the application of decentralized detection. The decentralized detection problem was introduced in [7], where it was shown that under certain independence assumptions, optimal decentralized detection rules take the form of likelihood ratio tests. However, it is shown in [9] that the problem of decentralized detection, a special case of the team decision problem, is also NP-hard. Therefore, most approaches to the problem of decentralized detection focus on determining person-by-person optimal detection rules. Surveys of the field of decentralized detection can be found in [8] and the book [10].

¹Department of Electrical Engineering, Stanford University, Stanford, CA 94305, U.S.A.
Email: rcogill@stanford.edu

²Department of Aeronautics and Astronautics, Stanford University, Stanford CA 94305-4035, U.S.A.
Email lall@stanford.edu

^{1,2}Partially supported by the Stanford URI *Architectures for Secure and Robust Distributed Infrastructures*, AFOSR DoD award number 49620-01-1-0365.

3 Motivating example

In this section we motivate the study of the general problems discussed in this paper by a specific application. The problem of *decentralized detection* is an example of a decentralized stochastic decision problem. Here we present a very brief overview of this subject. Detailed surveys can be found in [8] and [10].

In a detection problem, we have several hypotheses on the underlying state of our environment, and we would like use measurements of our environment to decide which hypothesis is true.




Figure 1: The correct hypothesis $H \in \{h_1, \dots, h_M\}$ is to be detected. In this figure, N independent detectors produce decisions u_i based on their measurements y_i .

Classical detection methods assume all measurements are available to a single detector, which estimates the true hypothesis based on all measurements. Such a detection scheme is called *centralized*. When minimizing the probability of error, optimal decision rules in centralized schemes are given by the well-known MAP (maximum a-posteriori probability) detector. In a *decentralized* detection scheme, each sensor is responsible for making a decision based only on its own measurement. The goal is to choose decision rules for all sensors which are optimal with respect to some system-wide cost function.

For example, suppose we have a collection of sensors each monitoring various elements of an industrial process. We would like the sensors to sound an alarm when some part of the process is malfunctioning. In this case we may wish to maximize the probability that the alarm sounds when there is a malfunction and does not sound when there is no malfunction. One option is to transmit all sensor measurements to a central location, where a decision to sound the alarm is made on the basis of all measurements. An alternative is to equip each sensor with its own decision rule and the ability to sound the alarm. When the loss of performance associated with employing the second alternative is small, such a scheme is preferable due to the reduced implementation complexity associated with the elimination

of the communication requirements.

One might initially assume that good decentralized decision rules can be obtained by allowing each sensor to use a MAP detection rule. While this is true in some special cases, it is not true in general. Unlike the centralized case, the general problem of computing optimal decentralized detection rules is NP-hard [9]. Also, decentralized decision rules can appear considerably more complex than their centralized counterparts. For example, optimal decentralized decision rules typically involve hedging among the sensors, a strategic element which is not present when simply using MAP rules at each detector.

Due to the complexity of this problem, most existing methods for computing decentralized detection rules produce locally optimal equilibrium policies. Such policies are said to be *person-by-person optimal*; for a set of such decision rules, no improvement can be obtained by adjusting the decision rule for any given sensor while leaving the others fixed. In general, a single problem instance may have many equilibrium policies. The globally optimal policy is clearly an equilibrium policy. However, for any given equilibrium policy, we have no way of knowing how this policy relates to the globally optimal policy. In particular, we have no way of knowing how much improvement we could obtain by using the globally optimal policy. In the next section we will show by a simple example that an equilibrium policy can perform arbitrarily poorly compared to the optimal policy. The methods that we present in this paper are *relaxations*. In addition to generating an equilibrium policy, they return an upper bound on the maximum achievable cost by any decentralized policy. When the bound is exact, we have a proof that our computed policy is globally optimal. Even when the bound is not exact, we still have a measure of the suboptimality of the computed policy.

4 Main results

4.1 Formulation and complexity

In this section we consider a general *static* decentralized decision problem, also commonly referred to as a *team decision problem* [1, 6]. For notational simplicity, we only discuss problems involving two decision makers. Extensions of the results to the general case of N decision makers is straightforward.

The distributed detection problem of the previous section is an example of the class of problems to be considered in this paper. The objective of the detector can be specified by an objective function to be minimized or maximized over the set of decentralized policies. For example, consider a problem with two detectors where we would like to determine the decentralized policy which maximizes the probability that all detec-

tors are correct. This is carried out by maximizing the objective function

$$J(\gamma_1, \gamma_2) = \sum_{y_1, y_2} \sum_i r_i(\gamma_1(y_1), \gamma_2(y_2), h_i) p(y_1, y_2 | h_i) p(h_i),$$

over all decentralized policies, where $p(h_i)$ is the a-priori probability of truth of each hypothesis, $p(y_1, y_2 | h_i)$ is the conditional probability of observing (y_1, y_2) given that h_i is the correct hypothesis, and

$$r_i(u_1, u_2, h_i) = \begin{cases} 1 & \text{if } u_1 = u_2 = h_i \\ 0 & \text{otherwise.} \end{cases}$$

The objective function above is often referred to as the *Bayes risk function*.

In its most general form, the specific problem considered in this paper is the following:

TEAM DECISION PROBLEM: *Given finite sets Y_1, Y_2, U_1, U_2 and a cost function $c : Y_1 \times Y_2 \times U_1 \times U_2 \rightarrow \mathbb{R}_+$, find policies $\gamma_i : Y_i \rightarrow U_i, i = 1, 2$ which maximize the cost*

$$J(\gamma_1, \gamma_2) = \sum_{y_1, y_2} c(y_1, y_2, \gamma_1(y_1), \gamma_2(y_2))$$

This problem is presented in [9]. There, the problem considered was that of maximizing expected cost given a cost function $\hat{c} : Y_1 \times Y_2 \times U_1 \times U_2 \rightarrow \mathbb{R}$ and probability mass function $p : Y_1 \times Y_2 \rightarrow \mathbb{R}$. It is easily shown that these two formulations are equivalent by taking our cost function to be $c(y_1, y_2, u_1, u_2) = \hat{c}(y_1, y_2, u_1, u_2) p(y_1, y_2)$. It was shown in [9] that this problem is NP-hard. Unless $P=NP$, we cannot hope to find an efficient algorithm capable of always producing globally optimal policies. In fact, even the restricted class of problems where $|U_1| = |U_2| = 2$ is still NP-hard. Efficient methods for computing policies must aim to find good suboptimal solutions. This problem is particularly interesting since the centralized version is trivial. For the centralized problem, the optimal policy simply chooses the decisions u_1 and u_2 which maximize $c(y_1, y_2, u_1, u_2)$ for each (y_1, y_2) .

One way of formulating the team decision problem involves expressing policy i as a $|Y_i| \times |U_i|$ Boolean matrix for each i :

$$K_{y_i u_i}^i = \begin{cases} 1 & \text{if } \gamma_i(y_i) = u_i \\ 0 & \text{otherwise} \end{cases}$$

Similarly, we can express the system cost as a $|Y_1| |Y_2| \times |U_1| |U_2|$ matrix $C_{yu} = c(y_1, y_2, u_1, u_2)$. Here, the matrix is indexed according to an order on the pairs

$y = (y_1, y_2)$ and $u = (u_1, u_2)$. The static decentralized decision problem can be equivalently formulated as

$$\begin{aligned} \text{maximize:} & \sum_{y, u} C_{yu} K_{yu} \\ \text{subject to:} & K_{yu} = K_{y_1 u_1}^1 K_{y_2 u_2}^2 \\ & K^i \geq 0 \quad i = 1, 2 \\ & K^i \mathbf{1} = \mathbf{1} \quad i = 1, 2 \\ & K_{yu} \in \{0, 1\} \quad \text{for all } y, u \end{aligned} \quad (1)$$

This problem is clearly a nonconvex optimization problem due to the Boolean constraints and the bilinear constraint. However, we can eliminate the Boolean constraints and show that the resulting problem is equivalent to (1):

Theorem 1. *The optimization problem*

$$\begin{aligned} \text{maximize:} & \sum_{y, u} C_{yu} K_{yu} \\ \text{subject to:} & K_{yu} = K_{y_1 u_1}^1 K_{y_2 u_2}^2 \\ & K^i \geq 0 \quad i = 1, 2 \\ & K^i \mathbf{1} = \mathbf{1} \quad i = 1, 2 \end{aligned} \quad (2)$$

always has an optimal solution satisfying $K_{yu} \in \{0, 1\}$ for all y, u .

Proof: Suppose \overline{K}^1 and \overline{K}^2 are optimal for (2). Note that \overline{K}^1 and \overline{K}^2 may have non-integer entries. The problem

$$\begin{aligned} \text{maximize:} & \sum_{y_1, u_1} \left(\sum_{y_2, u_2} C_{yu} \overline{K}_{y_2, u_2}^2 \right) K_{y_1, u_1}^1 \\ \text{subject to:} & K^1 \geq 0 \\ & K^1 \mathbf{1} = \mathbf{1} \end{aligned}$$

is a linear program in the variable K^1 . An optimal solution \widehat{K}^1 to this LP satisfies

$$\sum_{y, u} C_{yu} \widehat{K}_{y_1 u_1}^1 \overline{K}_{y_2 u_2}^2 \leq \sum_{y, u} C_{yu} \overline{K}_{y_1 u_1}^1 \overline{K}_{y_2 u_2}^2$$

Also, it is clear that \widehat{K}^1 can be chosen to have 0-1 entries. Now consider the linear program

$$\begin{aligned} \text{maximize:} & \sum_{y_2, u_2} \left(\sum_{y_1, u_1} C_{yu} \widehat{K}_{y_1, u_1}^1 \right) K_{y_2, u_2}^2 \\ \text{subject to:} & K^2 \geq 0 \\ & K^2 \mathbf{1} = \mathbf{1} \end{aligned}$$

Again, an optimal solution \widehat{K}^2 satisfies

$$\sum_{y, u} C_{yu} \widehat{K}_{y_1 u_1}^1 \widehat{K}_{y_2 u_2}^2 \leq \sum_{y, u} C_{yu} \widehat{K}_{y_1 u_1}^1 \overline{K}_{y_2 u_2}^2$$

and can be chosen to have 0-1 entries. Therefore, \widehat{K}^1 , \widehat{K}^2 , and $\widehat{K} = \widehat{K}^1 \otimes \widehat{K}^2$ constitute an optimal 0-1 solution to (2). ■

When the optimal solution in (2) is not unique, there may be a mixed optimal solution. However, the above theorem shows that there is always a Boolean solution which achieves the same objective value.

Although we were able to eliminate the Boolean constraints, finding a globally optimal solution to (2) is still a difficult problem. The most common approach for handling this problem is to employ an iterative scheme for finding a *person-by-person* optimal solution [1, 6, 10]. This type of scheme starts by initially choosing an arbitrary pair of policies. Policies are then modified by alternately optimizing each policy while leaving the other policy fixed (as in the proof of Theorem 1). Since there are a finite number of policies, and each step never results in a decrease in the objective, this method leads to an equilibrium solution in a finite number of steps. The problem with such methods is that problems may have many equilibria, and it not clear if any given equilibrium solution is necessarily a good one. In fact, we can show by a simple example that an equilibrium policy can be arbitrarily poor compared to the optimal policy.

Consider a trivial case where $|Y_1| = |Y_2| = 1$, $|U_1| = |U_2| = 2$. In this case, there is only one state, and each decision maker can choose between two decisions. Consider the cost function

$$c(y_1, y_2, u_1, u_2) = \begin{cases} \rho & \text{for } u_1 = u_2 = 1 \\ 2 & \text{for } u_1 = u_2 = 2 \\ 1 & \text{for } u_1 \neq u_2 \end{cases}$$

Consider the decentralized policy where $\gamma_1(y_1) = 2$ and $\gamma_2(y_2) = 2$. This policy achieves a cost of $J(\gamma_1, \gamma_2) = 2$. Leaving one decision rule fixed while changing the other always achieves a cost less than 2. Therefore, this policy is person-by-person optimal. However, the globally optimal decentralized policy in this case achieves a cost of $J(\gamma_1^*, \gamma_2^*) = \rho$. Since ρ is arbitrary, we can choose its value so that a suboptimal person-by-person optimal policy achieves a cost arbitrarily worse than the globally optimal cost.

In this paper, we consider an alternate approach to searching for equilibrium policies. We treat this problem as a polynomial optimization problem, and apply lifting methods to obtain convex relaxations. Such methods either produce a globally optimal solution, or produce a suboptimal solution along with a bound on its suboptimality.

4.2 A relaxation for the static problem

Consider the problem (2) with additional valid constraints added by taking products of the original linear

constraints:

$$\begin{aligned} \text{maximize:} & \quad \sum_{y,u} C_{yu} K_{yu} \\ \text{subject to:} & \quad K_{yu} = K_{y_1 u_1}^1 K_{y_2 u_2}^2 \\ & \quad \sum_{u_1} K_{yu} = K_{y_2 u_2}^2 \quad \text{for all } y_1, y_2, u_2 \\ & \quad \sum_{u_2} K_{yu} = K_{y_1 u_1}^1 \quad \text{for all } y_1, y_2, u_1 \\ & \quad K^i \mathbf{1} = \mathbf{1} \quad i = 1, 2 \\ & \quad K \geq 0 \end{aligned}$$

The additional constraints $K^i \geq 0$ and $K\mathbf{1} = \mathbf{1}$ are implied by the *linear* constraints, so they are left out for brevity. By dropping the bilinear constraint, we obtain the linear programming relaxation:

$$\begin{aligned} \text{maximize:} & \quad \sum_{y,u} C_{yu} K_{yu} \\ \text{subject to:} & \quad \sum_{u_1} K_{yu} = K_{y_2 u_2}^2 \quad \text{for all } y_1, y_2, u_2 \\ & \quad \sum_{u_2} K_{yu} = K_{y_1 u_1}^1 \quad \text{for all } y_1, y_2, u_1 \quad (3) \\ & \quad K^i \mathbf{1} = \mathbf{1} \quad i = 1, 2 \\ & \quad K \geq 0 \end{aligned}$$

Solving this linear program produces an upper bound on the maximum value achievable by a decentralized policy, as well as suboptimal policies described by K^1 and K^2 . When the relaxation is not exact, the policies K^1 and K^2 may not be person-by-person optimal. However, we can apply the iterative scheme described at the end of the previous subsection using these policies as a starting point to obtain improved person-by-person optimal policies. In the next section, we will show that applying such a scheme leads to policies which are guaranteed to be within a fixed bound of optimal.

4.3 Guaranteed suboptimality bounds

In the previous section, it was shown that by solving a relaxation, we produce suboptimal policies and bounds on the optimal achievable value for instances of TEAM DECISION PROBLEM. In this section, we show that by solving the relaxation and applying a person-by-person optimization scheme using the policies obtained from the relaxation as a starting point, we obtain policies which are guaranteed to be within a *fixed bound of optimal*. This is in contrast to the person-by-person optimization scheme with an arbitrary initial policy, where we saw in Section 4.1 that computed policies may achieve a value arbitrarily far from the optimal value.

Here we consider the following algorithm.

Algorithm 1.

1. Solve the relaxation (3) to obtain K^2 .

2. Choose \widehat{K}^1 with entries

$$\widehat{K}^1_{y_1 u_1} = \begin{cases} 1 & \text{if } u_1 = \operatorname{argmax} \left\{ \sum_{y_2, u_2} C_{yu} K_{y_2 u_2}^2 \right\} \\ 0 & \text{otherwise.} \end{cases}$$

3. Choose \widehat{K}^2 with entries

$$\widehat{K}^2_{y_2 u_2} = \begin{cases} 1 & \text{if } u_2 = \operatorname{argmax} \left\{ \sum_{y_1, u_1} C_{yu} \widehat{K}^1_{y_1 u_1} \right\} \\ 0 & \text{otherwise.} \end{cases}$$

The policies described by \widehat{K}^1 and \widehat{K}^2 have the following property:

Theorem 2. *Let the optimal decentralized policy be described by \overline{K}^1 and \overline{K}^2 . The policies produced by Algorithm 1 satisfy*

$$\sum_{y,u} C_{yu} \widehat{K}^1_{y_1 u_1} \widehat{K}^2_{y_2 u_2} \geq \frac{1}{|U_1|} \sum_{y,u} C_{yu} \overline{K}^1_{y_1 u_1} \overline{K}^2_{y_2 u_2},$$

where it is assumed, without loss of generality, that $|U_2| \geq |U_1|$.

Proof: Let K , K^1 , and K^2 be the optimal solutions to the relaxation (3). Clearly

$$\sum_{y,u} C_{yu} K_{yu} \geq \sum_{y,u} C_{yu} \overline{K}^1_{y_1 u_1} \overline{K}^2_{y_2 u_2}.$$

From the constraints of (3), $K^2_{y_2 u_2} = \sum_{u_1} K_{yu}$ for all y_1, y_2 , and u_2 . Since the elements of K are nonnegative, $K^2_{y_2 u_2} \geq K_{y_1 y_2 u_1 u_2}$ for all y_1, y_2, u_1 , and u_2 . Therefore,

$$\sum_{y,u} C_{yu} K^2_{y_2 u_2} \geq \sum_{y,u} C_{yu} \overline{K}^1_{y_1 u_1} \overline{K}^2_{y_2 u_2}.$$

For each y_1 ,

$$\max_{u_1} \left\{ \sum_{y_2, u_2} C_{yu} K^2_{y_2 u_2} \right\} \geq \frac{1}{|U_1|} \sum_{y_2, u_2} C_{yu} K^2_{y_2 u_2}.$$

Therefore,

$$\begin{aligned} \sum_{y,u} C_{yu} \widehat{K}^1_{y_1 u_1} \widehat{K}^2_{y_2 u_2} &\geq \sum_{y,u} C_{yu} \widehat{K}^1_{y_1 u_1} K^2_{y_2 u_2} \\ &\geq \frac{1}{|U_1|} \sum_{y,u} C_{yu} K^2_{y_2 u_2} \\ &\geq \frac{1}{|U_1|} \sum_{y,u} C_{yu} \overline{K}^1_{y_1 u_1} \overline{K}^2_{y_2 u_2}. \end{aligned}$$

■

Theorem 2 tells us that Algorithm 1 always produces policies that achieve an objective value within a factor of $1/|U_1|$ of optimal. For the special case where $|U_1| = |U_2| = 2$, which is still NP-hard, computed policies achieve a value within a factor of $\frac{1}{2}$ of the optimal value.

5 Numerical example

Here we illustrate some of the concepts discussed in this paper with a numerical example. Consider a decentralized detection problem with four hypotheses and two detectors. Let H denote the current hypothesis. The a-priori probabilities for each hypothesis are given by:

$$\operatorname{Prob}\{H = h_i\} = \begin{cases} 0.39 & \text{for } i = 1 \\ 0.31 & \text{for } i = 2 \\ 0.16 & \text{for } i = 3 \\ 0.14 & \text{for } i = 4 \end{cases}$$

The measurements M_1 and M_2 are made by each detector are each quantized to one of ten measurements. The conditional probabilities of each possible pair of measurements given each hypothesis are illustrated by the figure below.

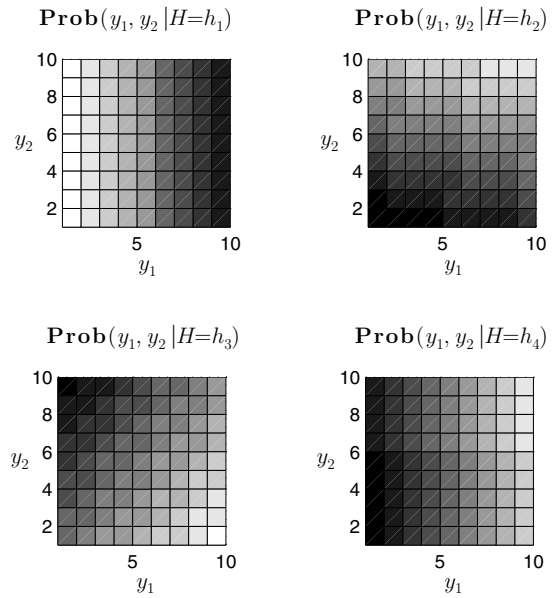


Figure 2: Conditional probabilities of each pair of measurements given each hypothesis. Dark areas on the plots represent low probabilities.

Each detector will estimate the hypothesis based only on its own observation. We would like to find decentralized detection rules which maximize the probability that *at least one detector is correct*.

We can formulate a relaxation of this problem as the linear program (3). In this case, the costs are

$$C_{yu} = \mathbf{Prob}\{(H \neq u_1) \cap (H \neq u_2) \cap (M_1 = y_1) \cap (M_2 = y_2)\}$$

Solving the relaxation, we obtain the globally optimal detection rules:

$$\gamma_1^*(y_1) = \begin{cases} 1 & \text{for } y_2 \leq 7 \\ 4 & \text{otherwise} \end{cases}$$

$$\gamma_2^*(y_2) = \begin{cases} 3 & \text{for } y_1 \leq 3 \\ 2 & \text{otherwise} \end{cases}$$

For this particular example, the optimal strategy achieves $\mathbf{Prob}\{\text{at least one correct}\} = 0.77$.

It is interesting to compare the optimal strategy to the one obtained by using a maximum a-posteriori detection rule for each detector. The MAP strategy is identical for both detectors and is given by

$$\gamma_i^{\text{MAP}}(y_i) = \begin{cases} 1 & \text{for } y_i \leq 6 \\ 2 & \text{otherwise} \end{cases}$$

In this example, the MAP strategy achieves $\mathbf{Prob}\{\text{at least one correct}\} = 0.62$. The key difference between the optimal strategy and the MAP strategy is the element of *hedging* employed by the optimal strategy. That is, the first hypothesis is the most likely, and it is most reliably detected by the first detector. In the optimal strategy, the second detector never guesses the first hypothesis. This is done to maximize the probability of guessing the correct hypothesis when the first detector guesses incorrectly. In the MAP strategy, both detectors are often both guessing the first hypothesis. When one is incorrect, the other is likely to be incorrect as well.

6 Conclusions

In this paper, we considered the problem of determining optimal decentralized decision rules in stochastic decision problems. It was shown that a general discrete decision problem has an equivalent formulation as a polynomial optimization problem. We obtain a relaxation of this polynomial optimization problem which can be used to compute suboptimal policies as well as bounds on the optimal achievable value. We also show that policies computed from this relaxation are guaranteed to be within a fixed bound of optimal.

In this paper, we restricted our focus to problems involving two decision makers. Of course, most practical problems where decentralization is a required property of the decision strategies will consist of many decision makers. We only consider problems with two decision makers for several reasons. First, it is known

that problems with two decision makers are computationally intractable. Therefore, these problems are the simplest cases which are still affected by the inherent difficulties of decentralized decision making. All of the results presented here will extend in a straightforward way to cases involving many decision makers, but we present all results in terms of the two decision maker case for simplicity. Secondly, for the general problem formulation considered in this paper, the input complexity required to specify a problem instance scales exponentially with the number of decision makers. In other words, specifying a cost function for a particular instance requires a number of parameters which is exponential in the number of decision makers. To discuss algorithms which are efficient in some meaningful way for problems involving many decision makers, we must restrict our attention to cost functions which admit a representation which scales gracefully with the number of decision makers. Similar issues are faced in the theory of multi-player games. There, it has been observed that cost functions in many real problems have special structure, such as symmetry or graphical structure [2, 3]. We believe that the results of this paper can be extended in a way which exploits similar structures in the cost function, although we have not pursued this issue here.

References

- [1] Y.C. Ho. Team decision theory and information structures. *Proceedings of the IEEE*, 68(6):644–654, 1980.
- [2] M. Kearns, M. Littman, and S. Singh. Graphical models for game theory. *Proceedings of the Conference on Uncertainty in Artificial Intelligence*, pages 253–260, 2001.
- [3] D. Koller and B. Milch. Multi-agent influence diagrams for representing and solving games. *Games and Economic Behavior*, 45:181–221, 2003.
- [4] J. Marschak and R. Radner. *The economic theory of teams*. Yale University Press, New Haven, CT., 1972.
- [5] C. Papadimitriou and J. Tsitsiklis. The complexity of Markov decision processes. *Mathematics of operations research*, 12(3):441–450, 1987.
- [6] R. Radner. Team decision problems. *Annals of Mathematical Statistics*, 33(3):857–881, 1962.
- [7] R. Tenney and N. Sandell. Detection with distributed sensors. *IEEE Trans. on Aerospace and Elec. Sys.*, 17(4):501–510, 1981.
- [8] J. Tsitsiklis. *Decentralized Detection*, volume 2, pages 297–344. 1993. H.V. Poor and J.B. Thomas, editors, JAI Press.
- [9] J. Tsitsiklis and M. Athans. On the complexity of decentralized decision making and detection problems. *IEEE Trans. Automatic Control*, 30(5):440–446, 1985.
- [10] P. Varshney. *Distributed detection and data fusion*. Springer, 1997.