Automatica 61 (2015) 178–191

Contents lists available at ScienceDirect

Automatica

journal homepage: www.elsevier.com/locate/automatica

An algorithmic approach to identify irrelevant information in sequential teams[☆]

Aditya Mahajan[a], Sekhar Tatikonda[b]

[a] Department of Electrical and Computer Engineering, McGill University, Montreal, Canada

[b] Department of Electrical Engineering, Yale University, New Haven, USA

ARTICLE INFO

Article history:
Received 8 August 2013
Received in revised form
16 June 2015
Accepted 28 July 2015
Available online 31 August 2015

Keywords:
Sequential teams
Decentralized stochastic control
Markov decision theory
Graphical models
Bayesian networks

ABSTRACT

An algorithmic framework that identifies irrelevant data (i.e., data that may be ignored without any loss of optimality) at agents of a sequential team is presented. This framework relies on capturing the properties of a sequential team that do not depend on the specifics of state spaces, the probability law, the system dynamics, or the cost functions. To capture these properties the notion of a team form is developed. A team form is then modeled as a directed acyclic graph and irrelevant data is identified using D-separation properties of specific subsets of nodes in the graph. This framework provides an algorithmic procedure for identifying and ignoring irrelevant data at agents, and thereby simplifying the form of control laws that need to be implemented.

1. Introduction

Teams refer to multi-agent stochastic control systems in which all agents have a common objective. Teams arise in many modern technologies including networked control systems, communication networks, sensor and surveillance networks, environmental remote sensing, and smart grids. Dynamic programming, which is the main solution concept for optimal design of centralized stochastic control, only works for specific sub-classes of team problems (Nayyar, Mahajan, & Teneketzis, 2013). To apply the dynamic programming principle to general team problems, one needs to identify *the structure of optimal control laws*. Such structural results are of two type: (i) remove irrelevant information at the controller; (ii) identify a sufficient statistic of the data available at the controller. In this paper, we present an algorithmic approach to identify the first type of structural result. As an example of such structural results, consider the problem of real-time communication studied in Witsenhausen (1979).

Example 1. Consider a real-time communication system consisting of a source, an encoder, and a decoder. The source is a first-order Markov process $\{S_t\}_{t=1}^{\infty}$. The encoder observes the source output and generates quantized symbols Q_t, causally and in real-time, as follows

$$Q_t = e_t(S_{1:t}, Q_{1:t-1})$$

where $S_{1:t}$ is a short hand notation for $(S_1, \dots, S_t)$ and $Q_{1:t-1}$ has a similar interpretation. The decoder is a finite state machine. M_t denotes the state of the machine at time t. The decoder generates an estimate $\hat{S}_t$ of the source as follows

$$\hat{S}_t = d_t(Q_t, M_{t-1})$$

and updates the contents of its memory as follows

$$M_t = g_t(Q_t, M_{t-1}).$$

At each time a distortion $c_t(S_t, \hat{S}_t)$ is incurred. The objective is to choose an encoding policy $\mathbf{e} := (e_1, e_2, \dots, e_T)$, a decoding policy $\mathbf{d} := (d_1, d_2, \dots, d_T)$, and a memory update policy $\mathbf{g} := (g_1, g_2, \dots, g_T)$ to minimize

$$\mathbb{E}_{(\mathbf{e},\mathbf{d},\mathbf{g})}\Big[\sum_{t=1}^{T} c_t(S_t, \hat{S}_t)\Big].$$

☆ A. Mahajan was supported by the Natural Science and Engineering Research Council of Canada (NSERC) Grant NSERC-RGPIN 402753-11 and S. Tatikonda was partially supported by National Science Foundation (NSF) grant CCF-1017744. The material in this paper was partially presented at the 2009 ConCom (Control over communication channels) workshop, June 27, 2009, Seoul, South Korea (Mahajan & Tatikonda, 2009a) and at the 47th Allerton conference on communication, control, and computing, September 30–October 2, 2009, Monticello, IL, USA (Mahajan & Tatikonda, 2009b). This paper was recommended for publication in revised form by Associate Editor Valery Ugrinovskii under the direction of Editor Ian R. Petersen.

E-mail addresses: aditya.mahajan@mcgill.ca (A. Mahajan), sekhar.tatikonda@yale.edu (S. Tatikonda).

http://dx.doi.org/10.1016/j.automatica.2015.08.002
0005-1098/

The following two structural results hold for Example 1.

(1) For any decoding and memory update strategies $(\mathbf{d}, \mathbf{g})$, there is no loss of optimality in restricting attention to an encoding strategy of the form

$$Q_t = e_t(S_t, M_{t-1}).$$

(2) When $M_{t-1} = Q_{1:t-1}$, there is no loss of optimality in restricting attention to encoding and decoding strategies of the form

$$Q_t = e_t(S_t, \Pi_t) \quad \text{and} \quad \hat{S}_t = d_t(Q_t, \Pi_t)$$

where $\Pi_t(s) = \mathbb{P}(S_t = s \mid Q_{1:t-1})$.

The first structural result shows that $S_{1:t-1}$ and $Q_{1:t-1}$ are *irrelevant* at the encoder. The second structural result shows that Π_t is a sufficient statistic for $Q_{1:t-1}$. The first structural result was proved in Witsenhausen (1979) and the second in Walrand and Varaiya (1983). Based on the second structural result, a dynamic programming decomposition was presented in Walrand and Varaiya (1983).

In this paper, we develop an algorithmic framework to identify irrelevant data at agents in a team. Removing such irrelevant data is usually the first step in deriving a dynamic programming decomposition for teams. For example, in the above example, the derivation of the second structural result (and therefore of the dynamic program) relies on the first structural result.

Structural results that remove irrelevant data are robust to various modeling assumptions: the specifics of the state spaces, the underlying probability measure, and the specifics of the plant dynamics and the cost functions. All that matters is the *form* of the system. We model dynamical systems using directed acyclic graph in such a manner that captures the *form* of a team. Removing irrelevant data is equivalent to removing edges from the corresponding directed acyclic graph. To identify the irrelevant data, we use graphical modeling algorithms to iteratively apply Blackwell's principle of irrelevant information (Blackwell, 1964), which we state below for completeness.

Theorem 1 (*Blackwell's Principle of Irrelevant Information*). *For any Borel spaces $\mathbb{X}$, $\mathbb{Y}$, and $\mathbb{U}$, let P be a probability measure on $\mathbb{X} \times \mathbb{Y}$ and $c : \mathbb{X} \times \mathbb{U} \to \mathbb{R}$ be a bounded Borel-measurable function. Then for any Borel-measurable function $g : \mathbb{X} \times \mathbb{Y} \to \mathbb{U}$, there exists another Borel measurable function $h : \mathbb{X} \to \mathbb{U}$ such that*

$$\mathbb{E}[c(x, h(x))] \leq \mathbb{E}[c(x, g(x, y))]$$

where the expectation is with respect to P.

A consequence of Theorem 1 is the following. Consider the optimization problem of choosing a control law $g : \mathbb{X} \times \mathbb{Y} \to \mathbb{U}$ to minimize $\mathbb{E}[c(x, g(x, y))]$. Then, there is no loss of optimality in restricting attention to control laws of the form $h : \mathbb{X} \to \mathbb{U}$. Equivalently, the observation y is irrelevant for optimal control. In this paper, we present algorithms that recursively apply Blackwell's principle at each agents and groups of agents to identify irrelevant data in teams.

1.1. Literature overview

Team problems were introduced in the economics literature in the 1950s (Marschak & Radner, 1972; Radner, 1962) and have been extensively analyzed in the control literature since the 1970s (Ho, 1980; Sandell, Varaiya, Athans, & Safonov, 1978; Witsenhausen, 1971). Motivated by applications in networked control systems, there has been tremendous activity in the study of team problems in the last decade. We refer the reader to Mahajan, Martins, Rotkowitz, and Yüksel (2012) and references therein for a detailed literature overview.

Broadly speaking, team problems are modeled either in state space using information structures or in input–output formulation using sparsity constraints. We follow the former modeling paradigm in this paper. Such models are analyzed either for the LQG setup (linear dynamics, quadratic cost, and Gaussian disturbances) or general (non-linear) Markovian setup. In this paper, we follow the latter setup and develop an algorithmic procedure to identify and remove irrelevant data at each agent.

We model teams using a directed acyclic graph (DAG) and use algorithms from graphical models to remove edges that correspond to irrelevant data. A DAG is a natural structure to model the causality and partial order relationship between the system variables of a sequential team. Other researchers have also used DAGs to model sequential teams (Gattami, 2007; Ho & Chu, 1972; Witsenhausen, 1971; Yoshikawa, 1978) but, to the best of our knowledge, the idea of using graphical modeling algorithms on the DAG representation to identify and remove redundant information has not been used before.

1.2. Contributions

Our main contribution is to present a graphical model for sequential team. This model captures the information structure of the system and the conditional independence relations between all system variables.

Using this graphical mode, we develop graphical modeling algorithms that identify irrelevant data at each agent. An agent can ignore this data without any loss of optimality. Two such algorithms are presented. The first algorithm sequentially identifies irrelevant data at each agent in the system. Preliminary versions of this algorithm were presented in Mahajan & Tatikonda, 2009a,b. The second algorithm sequentially identifies irrelevant data at all possible subsets of agents in the system. These algorithms do not depend on the type of system dynamics or the cost function.

The rest of the paper is organized as follows. In Section 2 we define team form and team type and formulate the problem of simplification of a team form. In Section 3 we present background material on graphical models and in Section 4 we describe how to represent a team form using a DAG (directed acyclic graph). Simplification of a team form may be viewed as removing edges from this DAG. Algorithms that perform this simplification are presented in Section 5 (for a single agent) and Section 6 (for a group of agents). Examples illustrating this approach are presented in Section 7 and we conclude in Section 8.

1.3. Notation

We use the following notation in the paper.

- For a set A, $|A|$ denotes the cardinality of A.
- For two sets A and B, $A \times B$ denotes their Cartesian product.
- For two measurable spaces $(\mathbb{X}, \mathscr{F})$ and $(\mathbb{Y}, \mathscr{G})$, $\mathscr{F} \otimes \mathscr{G}$ denotes the product σ-field on $\mathbb{X} \times \mathbb{Y}$.
- For two probability measures μ on $(\mathbb{X}, \mathscr{F})$ and ν on $(\mathbb{Y}, \mathscr{G})$, $\mu \otimes \nu$ denotes the product probability measure on $\mathscr{F} \otimes \mathscr{G}$.
- $X_{1:t}$ is a short hand for the sequence $(X_1, X_2, \dots, X_t)$.
- For a set N and a sequence of random variables $\{X_n\}$, X_N is a short-hand for $(X_n : n \in N)$.
- For a set N and a sequence of state spaces $\{\mathbb{X}_n\}$, $\mathbb{X}_N$ is a short-hand for $\prod_{n \in N} \mathbb{X}_n$.
- For a set N and a sequence of σ-fields $\{\mathscr{F}_n\}$, $\mathscr{F}_N$ is a short-hand for $\bigotimes_{n \in N} \mathscr{F}_n$.

2. Modeling sequential team using team form and team type

A sequential team is a decentralized control system consisting of multiple agents (also called controllers or decision makers),