16 APPLICATIONS OF MARKOV DECISION PROCESSES IN COMMUNICATION NETWORKS

Eitan Altman

Abstract: We present in this chapter a survey on applications of MDPs to communication networks. We survey both the different application areas in communication networks as well as the theoretical tools that have been developed to model and to solve the resulting control problems.

16.1 INTRODUCTION

Various traditional communication networks have long coexisted providing disjoint specific services: telephony, data networks and cable TV. Their operation has involved decision making that can be modeled within the stochastic control framework. Their decisions include the choice of routes (for example, if a direct route is not available then a decision has to be taken which alternative route can be taken) and call admission control; if a direct route is not available, it might be wise at some situations not to admit a call even if some alternative route exists.

In contrast to these traditional networks, dedicated to a single application, today's networks are designed to integrate heterogeneous traffic types (voice, video, data) into one single network. As a result, new challenging control problems arise, such as congestion and flow control and dynamic bandwidth allocation. Moreover, control problems that had already appeared in traditional networks reappear here with a higher complexity. For example, calls corresponding to different applications require typically different amount of network resources (e.g. bandwidth) and different performance bounds (delays, loss probabilities, throughputs). Admission control then becomes much more complex than it was in telephony, in which all calls required the same performance characteristics and the same type of resources (same throughput, bounds on loss rates and on delay variation).

We do not aim at a complete survey of the area, since several other surveys [90, 94, 149, 119, 150, 152, 251, 248, 250, 253] on related issues already exist, see also [221]. Other references that focus on the methodology that allows one to apply MDPs to communication networks can be found in the following books [4, 54, 111, 239, 271, 283] as well as in the survey paper [192].

We have two goals in this survey. First, we wish to present to researchers who specialize in MDPs a central application area which provides a vast field of challenging problems. We would like to familiarize these researchers with special complex features in control problems that arise in communications: complex information structure, problems with multiobjective and multiagents. A second objective is to familiarize researchers in communications with tools that have been developed for modeling and for solving control problems in networks.

Problems that are described in this survey are MDPs in a general sense: they are described as Markov chains whose transition probabilities are controlled. This control can be done by a single or several controllers, having the same or having different objectives. It is often tools other than the standard dynamic programming that are used to solve these control problems. For completeness, we occasionally mention approaches that have been used for control of communication systems which are not based on MDPs, and then relate or compare these to MDPs.

16.2 THE CONTROL THEORETICAL FRAMEWORK

The most popular telecommunication network architectures today are the Internet and the ATM (Asynchronous Transfer Mode) networks. The Internet offers today a "best effort" type service, i.e. the resources in the network are shared among all users, and when the number of users increases, the quality of service (in terms of delay, throughput, losses) per user decreases. In contrast, ATM networks provide mostly guaranteed services: if a session needs a given Quality Of Services (QOS) it may establish a contract with the network that guarantees that all along the duration of the session, some required bounds would hold on given performance measures (loss probabilities, delays, delay variation, throughput), as long as the source respects the terms of the contract (in terms of the average and the maximum throughput it sends, as well as some constraints on its bursty behavior). Two guaranteed service classes are defined in ATM: the CBR (Constant Bit Rate) and VBR (Variable Bit Rate) service classes. ATM contains also two best effort type services: the Available Bit Rate (ABR) service and the Unspecified Bit Rate (UBR) service. In the ABR service, the network determines the allowed transmission rate of each source by sending to them periodically appropriate control signals. As long as a source abides to those commands, the network can guarantee some given bounds on its loss rates. In the UBR service no guarantees are given on performance measures.