CCIT Report #519 February 2005

Robustness of policies in Constrained Markov Decision Processes

Alexander Zadorojniy and Adam Shwartz, *Senior Member, IEEE*

Available at www.ee.technion.ac.il/~adam/PAPERS/robustCMDP05R.pdf

Abstract

We consider the optimization of finite-state, finite-action Markov Decision processes, under constraints. Cost and constraints are of the discounted type. We introduce a new method for investigating the continuity and robustness of the optimal cost and the optimal policy under changes in the constraints. This method is also applicable for other cost criteria such as finite horizon and infinite horizon average cost.

Index Terms

Markov Decision Processes, Constrained MDP, Discounted Cost, Sensitivity, Robustness.

I. Introduction

Consider the standard model of a Markov Decision Process (MDP) with finite state and action spaces. A natural generalization of the optimization problem is to include cost constraints. Such models arise in relation to resource-sharing systems. For example, in telecommunication networks which are designed to enable simultaneous transmission of different types of traffic: voice, file transfer, interactive messages, video, etc. Typical performance measures are transmission delay, power consumption, throughput, etc. [1]. A trade-off exists, for example, between minimizing delay and reducing power consumption: to minimize delay we should transmit with the highest possible power, since this increases the probability of successful transmission. Such problems are formulated as constrained MDP [2], where we wish to minimize the costs related to the delay subject to constraints on the average and peak power.

Work of Alexander Zadorojniy was performed while he was with the Faculty of Electrical Engineering, Technion, Israel Institute of Technology.

Adam Shwartz is with the Faculty of Electrical Engineering, Technion, Israel Institute of Technology. Work of this author was supported in part by the fund for promotion of research at the Technion.