European Journal of Operational Research 218 (2012) 698–707

Contents lists available at SciVerse ScienceDirect

European Journal of Operational Research

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

Stochastics and Statistics

Approximate dynamic programming with Bézier Curves/Surfaces for Top-percentile Traffic Routing

Andreas Grothey, Xinan Yang *

School of Mathematics, College of Science and Engineering, The University of Edinburgh, United Kingdom

ARTICLE INFO

Article history:
Received 16 November 2010
Accepted 28 November 2011
Available online 8 December 2011

Keywords:
Top-percentile pricing
Multi-homing
Stochastic
Routing policy
Approximate dynamic programming
Bézier Curves/Surfaces

ABSTRACT

Multi-homing is used by Internet Service Providers (ISPs) to connect to the Internet via different network providers. This study develops a routing strategy under multi-homing in the case where network providers charge ISPs according to top-percentile pricing (i.e. based on the θth highest volume of traffic shipped). We call this problem the Top-percentile Traffic Routing Problem (TpTRP).

Solution approaches based on Stochastic Dynamic Programming require discretization in state space, which introduces a large number of state variables. This is known as the curse of dimensionality in state space. To overcome this, in previous work we have suggested to use approximate dynamic programming (ADP) to construct value function approximations, which allow us to work in continuous state space. The resulting ADP model provides well performing routing policies for medium sized instances of the TpTRP. In this work we extend the ADP model, by using Bézier Curves/Surfaces to obtain continuous-time approximations of the time-dependent ADP parameters. This modification reduces the number of regression parameters to estimate, and thus accelerates the efficiency of parameter training in the solution of the ADP model, which makes realistically sized TpTRP instances tractable. We argue that our routing strategy is near optimal by giving bounds.

1. Introduction

Internet Service Providers (ISPs) do not generally have their own network infra-structure to route the incoming traffic of their customers, but instead use external network providers. Multi-homing is used by ISPs to connect to the Internet via more than one network provider. This technique is currently widely adopted to provide fault tolerance and traffic engineering capabilities [1].

Traditionally network providers charge ISPs based on a combination of fixed cost and per usage pricing. Top-percentile pricing is a relatively new and increasingly popular pricing regime (although it usually appears as part of a mixed pricing strategy), that is quickly becoming established [9]. In this scheme, the network provider divides the charge period, say a month, into several time intervals with equal, fixed length. Then, it measures and evaluates the amount of data (traffic) sent in these time intervals. At the end of the charge period, the network provider selects the traffic volume of the top q-percentile interval as the basis for computing the cost. For example, if the charge period (i.e. 30 days) is divided into 4320 time intervals with a length of 10 min, and if top 5-percentile pricing is used, the cost computed by top-percentile pricing is based on the traffic volume of the top 216th interval.

It has been discussed (e.g. in [9] what the optimal multi-homing routing strategies look like under traditional pricing regimes and whether they are economically viable. In contrast, very little work has been done on network operation under top-percentile pricing. The deterministic problem (in which we assume that we know all the traffic volumes in advance) has been analysed in [2], where the authors investigate the Traffic Routing Problem under a combined pricing policy – top-percentile pricing and fixed cost pricing. In the stochastic case, Levy et al. [8] developed a probabilistic model and provided analysis of the expected costs, thus demonstrated that multi-homing could be economically efficient under top-percentile pricing though they did not investigate the optimal routing policy. On the other hand, Goldenberg et al. [5] focused on the development of smart routing algorithms for optimising both cost and performance for multi-homing users under top-percentile pricing. However, in the case where traffic volumes are not available in advance (stochastic case), they apply the deterministic algorithm directly on the prediction of one later time interval's traffic and then use heuristics to accommodate the prediction error. As a conclusion, to the best of our knowledge there is no result dealing with the optimal multi-homing routing policy under top-percentile pricing in the stochastic case.

The purpose of our study is to find an optimal or near optimal routing strategy in order to allow the ISP to make full use of the underlying networks with minimum cost, when all network

* Corresponding author.
E-mail address: s0677435@sms.ed.ac.uk (X. Yang).

0377-2217/$ - see front matter © 2011 Elsevier B.V. All rights reserved.
doi:10.1016/j.ejor.2011.11.041

providers operate independently in parallel and charge the ISP based on the volume of the top q-percentile time interval's traffic (pure top-percentile pricing). In the following parts of this paper we call this problem, the Top-percentile Traffic Routing Problem (TpTRP). The TpTRP is a stochastic problem, where the ISP cannot predict the volume of future time intervals' traffic. Instead, we assume that the ISP knows the probabilistic distributions for all time intervals' traffic ahead of time.

In [6], the authors have shown that solving the TpTRP as an Stochastic Mixed-Integer Programming (SMIP) problem is intractable for all but extremely small instances. The reason is that modelling of the top-percentile cost introduces integer variables in the final stage subproblem, which makes the stochastic programming model non-convex and hard to solve. On the other hand, the authors suggested a Stochastic Dynamic Programming (SDP) model based on a discretization of the state space. The solution of the SDP model gives routing policies that outperform all available naive routing policies for small sized instances (with up to 3 network providers, each dividing the charge period into 10 time intervals and charges based on the traffic volume in the 3rd highest time interval). Nevertheless, due to the huge number of state variables arising from discretization, an effect well known as the curse of dimensionality prevents the use of the SDP model on larger problem instances. As a modification, in [7] the authors applied approximate dynamic programming (ADP) to solve the TpTRP. ADP works on a continuous state space thus overcomes the curse of dimensionality introduced by discretization. With the ADP model, medium sized TpTRP instances (with up to 3 network providers, each dividing the charge period into 86 time intervals and charges based on the traffic volume in the 5th highest time interval) can be solved within reasonable time.

This work follows the study in [7]. We intend to develop an ADP model based aggregation algorithm to make realistically sized TpTRP instances (with up to 3 network providers, each dividing the charge period into 4320 time intervals and charges based on the traffic volume in the 216th highest time interval) tractable. The focus of this work is on the investigation of the numerical structure of the time-dependent parameters in the original discrete ADP model. Based on this investigation, a Bézier Curves approximation of the parameter structure is introduced and integrated into the ADP model. This modification reduces the number of regression parameters to estimate (e.g. by a factor of 20 in a 86 time interval instance), thus accelerates the efficiency of parameter training in the solution of the ADP model. TpTRP instances with 3 network providers and 216 time intervals can be trained within 20 min with the ADP-Bézier-Curve model. To accelerate the process further, we extend the Bézier Curves approximation into a two-dimensional, Bézier Surfaces approximation. The resulting ADP-Bézier-Surface model, as expected, takes around 10 min to solve instances with 432 time intervals (with 2 network providers). Although it is not directly applicable to realistically sized instances due to the long running time, with a simplification of the decision process, realistically sized TpTRP instances become tractable. Note that the discretization in state space is only part of the reason for the curse of dimensionality, indeed the number of realisations of the state also depends exponentially on the number of network providers. However, at least for low number of network providers the computational complexity of the ADP method effectively increases linearly with the number of network providers. This is discussed in Section 3.2. Test examples with 2 or 3 network providers show that the solution difficulty increases linearly with the number of network providers, which justifies that the method should be applicable to up to 5 network providers (which is a realistic upper bound in practice).

The remainder of this paper is structured as follows. We firstly present the TpTRP problem and its basic SDP modelling elements in Section 2. In Section 3, we give a brief introduction to the ADP technique and build the ADP model. Then we analyse the complexity of the ADP model and show how to exploit its structure to improve it with Bézier Curves in Section 4. Section 5 gives the numerical results and provides a stronger – Bézier Surfaces aggregation model, which makes realistically sized instances tractable. Finally we give our conclusions in Section 6.

2. The Top-percentile Traffic Routing Problem

This section gives a formal description of the TpTRP parameters and defines the main modelling elements in the dynamic programming model formulation.

2.1. Problem parameters

- $\mathcal{I},\ |\mathcal{I}| = n$: The set of network providers.
- $\mathcal{T},\ |\mathcal{T}| = \Gamma$: The set of time intervals.
- q: The percentile parameter.
- $\theta = \lfloor \Gamma q \rfloor$: The index of the top-percentile time interval.
- $c_i,\ i \in \mathcal{I}$: The per unit cost charged by network provider i on the top-percentile traffic.

In this work, we assume that all network providers divide the charge period equally into Γ time intervals. Network providers use pure top-percentile pricing with parameter q, namely the cost charged on the ISP depends solely on the θth highest volume of traffic that has been sent to network provider i:

$$C_i = c_i T_i^{\theta},$$

where T_i^{θ} indicates the θth highest volume of traffic that has been sent to network provider i. Thus under pure top-percentile pricing, if a given network provider only carries traffic in $\theta - 1$ time intervals, then the network provider charges nothing to the ISP.

We assume that there is no upper bound on the volume of traffic that can be shipped to each network provider, and that no failure occurs in any network during the planning period.

- $T^{\tau},\ \tau \in \mathcal{T}$: The volume of traffic in time interval τ.

Note that in practice traffic arrives at the ISP continuously. In order to route each packet of traffic without delay, a routing decision for period τ needs to be in place before the amount of traffic is known. Thus in this work, we assume that before the routing decision for period τ is made, $T^{\tau}(\omega^{\tau})$ is a random variable depending on the random event ω^{τ}. When the random event $\hat{\omega}^{\tau}$ becomes known, we use $\widehat{T}^{\tau} = T^{\tau}(\hat{\omega}^{\tau})$ to represent the realisation of T^{τ}.

- $x^{\tau},\ \tau \in \mathcal{T}$: The routing decision (also referred as an action in Dynamic Programming terminology) for time interval τ. The detailed description of the feasible decision set is presented later in Section 2.3.

In this work, the decision for period τ should be made before knowing the realization of T^{τ} and all the remaining time intervals' traffic. Thus the TpTRP is a stochastic problem in which decisions are made under uncertainty. Note that as we use dynamic programming based approaches to solve the problem, the routing policies considered in this work only depend on the current state rather than the whole traffic history.

2.2. State variables and value function

At the beginning of time interval τ, all the previous realisations of traffic volumes $\widehat{T}^{t},\ t = 1, \ldots, \tau - 1$ and routing decisions x^{t},