Automatica 72 (2016) 158–165

Contents lists available at ScienceDirect

Automatica

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

Brief paper

Optimization-based, QoS-aware distributed traffic control laws for networks with time-varying link capacities[☆]

Wenjing Su[a], Constantino M. Lagoa[a], Hao Che[b]

[a] The Pennsylvania State University, University Park, PA, USA
[b] University of Texas at Arlington, TX, USA

ARTICLE INFO

Article history:
Received 17 October 2014
Received in revised form
22 February 2016
Accepted 16 May 2016

Keywords:
Overlay networks
Traffic engineering
Optimization
Sliding modes control
Quality of service

ABSTRACT

It is a challenge to design optimal, distributed traffic control mechanisms in a network where the link capacities and Class of Service (CoS) requirements may vary with time, such as a virtual network, an overlay network, or a wireless network. In this paper, we develop a family of optimization-based distributed traffic control laws to meet this challenge. This family of control laws enables Quality of Service (QoS), Traffic Engineering (TE), and Failure Recovery (FR) features simultaneously in a network where the link capacities and CoS requirement may vary with time. The approach taken relies on the concept of Sliding Modes to solve the resulting time-varying optimization problem. Running at the edge of a network, a set of control laws selected from this family enables class-of-service-based multi-path load balancing and/or rate adaptation to respond to network congestion, CoS requirement variation and link failures. The only nonlocal information needed as input to the control laws is the number of congested links along a forwarding path. This family of control laws is particularly viable to be implemented in a software-defined network (SDN) where the available underlying resources may not be accurately predictable and vary with time, due to, e.g., network virtualization.

1. Introduction

As communication networks become large, e.g., the Internet with global coverage, user traffic control must be distributed by design. Namely, the flow rates for the flows sending between various source–destination host pairs must be controlled by individual pairs independent of one another. The problem then arises as to how to control individual flows so that the network will converge to a stable state that achieves certain design objective. Mathematically, this problem is generally stated as follows: Find flow control laws that maximize the sum of individual user utilities as functions of flow rates, under network resource constraints. Significant research effort has been made in developing such traffic control laws. The aim is to provide theoretical underpinnings for the development of distributed traffic control protocols with provable properties, including scalability, stability, and optimality. This line of research has proven to be successful and has resulted in the finding of large families of optimization-based, distributed control laws (e.g., Kelly, Maulloo, & Tan, 1998, Kelly & Voice, 2005, Lagoa, Che, & Movsichoff, 2004, Low & Lapsley, 1999, Wang, Palaniswami, & Low, 2003 and Movsichoff, Lagoa, & Che, 2007). This development is highly encouraging because by using these control laws, distributed traffic control protocols can now be developed based on *theory*, rather than ad hoc/empirical design.

However, all the existing control laws work under the assumption that the link bandwidths (or equivalently, link capacities) are fixed, which largely limits the applicability of such control laws. This is simply because the assumption may not hold true for some segments of the Internet, e.g., wireless networks, overlay networks, or virtual networks.

In a wireless network, the effective radio channel capacity is both time and space varying (Li, Che, & Li, 2007), as a result of several factors, including signal power attenuation, inter-channel interference, thermal noise, Doppler frequency shift, and multipath channel fading. An overlay network particularly enabled at the application layer, provides a convenient way to quickly deploy services with minimum interruption to the IP layer infrastructure, such as a P2P network (Wang & Li, 2003) and a

---

[☆] This work was partially supported by the National Science Foundation under grants CNS-1329422 and ECCS-1201973. The material in this paper was partially presented at the 46th IEEE Conference on Decision and Control, December 12–14, 2007, New Orleans, LA, USA. This paper was recommended for publication in revised form by Associate Editor Peng Shi under the direction of Editor Toshiharu Sugie.

E-mail addresses: wenjing_su@hotmail.com (W. Su), lagoa@engr.psu.edu (C.M. Lagoa), hche@cse.uta.edu (H. Che).

http://dx.doi.org/10.1016/j.automatica.2016.06.003

VoIP network (Singh & Chauhan, 2014). In such networks, the underlying network topology and resources are hidden. Moreover, the network resources must be shared between the overlay traffic and the underlying network layer traffic. As a result, an overlay link capacity is time varying and the network layer path that forms an overlay link may change from time to time.

Network virtualization (Scroggins, 2013) is a natural extension of server virtualization to allow fully virtualized distributed computing; e.g., in a data center. Virtualization aims at emulating a hardware platform, such as a server, storage device or network resource, in software. Network virtualization has come into focus recently, with the true decoupling of the control and forwarding data planes, as advocated by software-defined networking (SDN) (https, 0000a) and network functions virtualization (NFV) (https, 0000b). On one hand, the decoupling of the virtual networks from its physical counterparts allows the use of any general off-the-shelf physical platforms and multiple virtual networks sharing the same physical platform. On the other hand, the decoupling of the control and data planes makes it difficult to accurately predict the resource availability in virtual networks, which may change from time to time.

The main contribution of this paper is the development of a family of optimization-based, distributed, multipath enabled traffic control laws, which can function properly in the above mentioned networks with time-varying link bandwidths, hidden and/or changing topologies. These distributed traffic control laws enable rich QoS, TE, and FR features through QoS-aware multipath load balancing and rate adaptation, and require only a single-bit binary congestion feedback between two nodes interconnected through a link with time-varying link capacity. For example, in the case of an virtual/overlay network, it requires only a single-bit binary congestion feedback between two virtual/overlay nodes, which can be inferred by the virtual/overlay nodes themselves. This allows an overlay/virtual link to be truly virtual, in the sense that the network layer path that forms this virtual/overlay link is completely hidden from the overlay and may change from time to time.

Our approach has its basis in the theory of Sliding Mode Control. The results in Utkin (1992) indicate that Sliding Modes theory can be a powerful tool for optimizing convex functions subject to a large number of convex constraints. Motivated by this result, we provide a family of optimal distributed traffic allocation algorithms for the time varying problem described above.

Finally, we note that our control laws are not meant to be a comprehensive solution, nor concerned with the protocol development and implementation. Instead, it is a theoretical underpinning upon which highly scalable and versatile traffic control mechanisms can be developed. It is quite similar to the distributed Bellman–Ford algorithm, which underpins any distance-vector routing protocols, rather than the protocols themselves. With mathematical rigor, it simply states that with any given resources including given multipaths, given set of next hops and given workload such as x number of best-effort (BE) flows and y number of assured service (AS) flows, the family of control laws will drive the traffic allocation to track the time-varying link bandwidths and topology changes so that a given global utility function is maximized. Moreover, this paper is not concerned with the implementation issues, e.g., how to design the utility function, how to translate a control law associated with a given utility function into a window-based traffic control protocol, how to engineer the paths and allocate the link bandwidths, and so on. Interested readers are referred to Ye, Wang, Che, and Lagoa (2011) for the implementation of an end-to-end traffic control protocol based on the distributed controllers found in Lagoa et al. (2004).

The remainder of this paper is organized as follows. In Section 2, we briefly discuss related work. In Section 3, we introduce the notation that is used throughout the paper and provide an exact problem statement. In Section 4, the family of optimal control laws is introduced. Examples of application of the proposed algorithm are presented in Section 5. Conclusions and directions for future research are given in Section 6. Finally, a sketch proof of the main results is provided in the Appendix.

2. Related work

In this section, we focus exclusively on surveying literature relevant to the work presented in this paper, i.e., optimization-based, multipath-enabled, distributed traffic control schemes. But note that none of the existing solutions can deal with the case with time-varying link capacities.

The general approach taken to address the optimization-based, distributed traffic control problem is to formulate it, based on a fluid-flow model, as a nonlinear programming problem, taking into account of link bandwidth constraints. The aim is to find distributed traffic control laws with the property that by working independent of one another, they will drive the network to an operation point where a given utility function of flow rates is maximized. Since different flows share the link resources which are constrained, the key challenge in the design of distributed control laws is the fact that there is a high degree of interaction between different flows. This challenge has been tackled in three ways.

The first approach (e.g. Elwalid, Jin, Low, & Widjaja, 2001 and Golestani & Bhattacharyya, 1998) is to incorporate link congestion costs into the overall utility function to convert a constrained problem into a non-constrained problem. The optimization problem was then solved using a gradient type of algorithm. Iterative algorithms were proposed where individual sources periodically adjust/balance their flow sending rates to multiple paths based on the congestion cost information fed back from the (congested) links along each path. It is generally used for a network where each path is pinned, such as a label switched path (Elwalid et al., 2001), and, hence, not suitable for an overlay network where each physical path forming an overlay link is hidden.

The second approach (e.g., Han, Shakkottai, Hollot, Srikant, & Towsley, 2003, Kelly et al., 1998 and Kelly & Voice, 2005) is to solve a relaxation of the original problem, by incorporating a price function into the overall utility function. Distributed control laws were found, which were proven to be locally stable in the presence of variable feedback delays. Working independently at a source, a control law adjusts/balances its flow sending rates to multiple paths based on periodic, cumulative price feedbacks from the destination node. Each component price is collected from the intermediate links along the forwarding path. The involvement of core nodes in conveying price information for each flow makes it difficult, if ever possible, to apply this scheme directly to an overlay environment where the underlying network topology is hidden.

The third approach is to solve the original problem directly (e.g., Lagoa & Che, 2000, Lagoa et al., 2004, Movsichoff, Lagoa, & Che, 2005, Movsichoff et al., 2007, Su et al., 2015 and Wang et al., 2003). Using a duality model, an algorithm was provided in Wang et al. (2003). In Lagoa and Che (2000), Lagoa et al. (2004), Movsichoff et al. (2005) and Movsichoff et al. (2007), this problem was tackled by using a technique based on the Theory of Sliding Modes. Both end-to-end (Lagoa & Che, 2000; Lagoa et al., 2004; Movsichoff et al., 2007) and hop-by-hop (Movsichoff et al., 2005) optimal control laws were found. Moreover, distributed multidomain optimal control laws are also found in Su et al. (2015). These algorithms allow multipath forwarding and enable multiple Classes of Service (CoSes), and require minimum information feedback for control. For an overview of recent results in sliding mode theory and its applications see, for example, Bartolini, Fridman, Pisano, and Usai