Computer Networks 52 (2008) 3373–3391

Contents lists available at ScienceDirect

Computer Networks

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

Pricing the quality of differentiated services for media-oriented real-time applications: A multi-attribute negotiation approach

Marco Comuzzi, Chiara Francalanci, Paolo Giacomazzi *

Politecnico di Milano, Dipartimento di Elettronica e Informazione, Piazza Leonardo da Vinci, 32, I-20133 Milano, Italy

ARTICLE INFO

Article history:
Received 21 November 2007
Received in revised form 12 August 2008
Accepted 5 September 2008
Available online 18 September 2008

Responsible Editor: A. Pitsillides

Keywords:
Negotiation
Differentiated service
Real-time applications
Multimedia applications

ABSTRACT

The paper proposes a new algorithm for negotiating the price of quality of service in IP differentiated services networks. We focus on real-time media-oriented applications such as video conferencing, online gaming, broadcast TV and live events, streaming video, and audio on demand. The performance objectives of end-users are expressed as the end-to-end delay thresholds to be exceeded with a given maximum probability. The target of negotiation is a multi-attribute description of traffic profile and quality of service, rather than the simple raw-bandwidth attribute. The service class chosen for each traffic flow is the result of negotiation and depends on the user's and supplier's utility and quality functions and on their conceding versus selfish negotiation attitude. We model non-linear utility and quality functions in such a way to represent the user's and supplier's perception of quality of service parameters. This represents a fundamental contribution of this paper with respect to current approaches accounting for simple linear utility functions of the raw-bandwidth attribute. We analyze the utilization of network resources as well as the customer's and supplier's utility through simulation by comparing our algorithm with previous algorithms negotiating raw-bandwidth instead of end-to-end quality of service.

1. Introduction

Most of current research in the field of pricing Internet services follows the game-theoretical approach by implementing auctioning systems, because of their significant incentive properties. Marbach [24] makes pricing decisions based on the priorities of customers in obtaining a min–max fair allocation of bandwidth. In [19], MacKie–Mason and Varian use a per-packet auction to solve congestion problems through pricing. Reichl et al. [28] adopt Delta auctions, a second-price mechanism useful with multiple service providers. Semret and Lazar [30] propose progressive second-price auctions to price bandwidth. Prouskas et al. [35] propose an auction-based protocol for improving load control capabilities of an Intelligent Network. In [23], Maillé and Tuffin propose a multibid auction pricing mechanism, consisting in a one-shot auction to share network capacity among several users. The one-shot auction approach is particularly interesting as it minimizes the communication overhead required to perform the multilateral bargaining.

Altman et al. [1] are the first to tackle the pricing problem in the context of IP differentiated services [4,12]. They consider the packet dropping rate as the relevant quality of service (QoS) metric. In this way, the authors position their methodology within the assured forwarding per hop behavior [13], where the flexibly customizable active queue management techniques can be used to finely adjust the packet dropping rate of service categories that can be differentiated, and priced, accordingly. A further development in the context of IP differentiated services is presented in [34] by Zachariadis and Barria, using effective bandwidth as the QoS metric (effective bandwidth is a metric closer to the actual QoS than the raw link capacity), without considering packet loss and without selecting a specific per hop behavior for their analysis.

* Corresponding author. Tel.: +39 02 2399 3596; fax: +39 02 2399 3413.
E-mail addresses: comuzzi@elet.polimi.it (M. Comuzzi), francala@elet.polimi.it (C. Francalanci), giacomaz@elet.polimi.it (P. Giacomazzi).

1389-1286/$ - see front matter © 2008 Elsevier B.V. All rights reserved.
doi:10.1016/j.comnet.2008.09.003

A common feature of this previous research is the selection of a single attribute to be negotiated. Raw-bandwidth is adopted in [20–22,28,30], the packet dropping rate is used in [1], the effective bandwidth is selected in [34], while [36] uses the packet delay. Semret et al. [31] propose a single-attribute approach to the pricing of edge-allocated bandwidth.

In this paper, we propose a novel algorithm for the negotiation of the quality of service of multimedia traffic flows sharing a bottleneck link. We refer to the IP differentiated services architecture [4] and focus on media-oriented real-time applications with tight QoS requirements on end-to-end delay and packet loss rate. The first innovative contribution is that we implement a multi-attribute QoS negotiation algorithm, as opposed to the usual single-attribute approach. Our algorithm dynamically sets the end-to-end QoS characteristics of traffic flows. This is in line with the most advanced offers of IP differentiated services providers, which include service level agreements (SLAs) comprising end-to-end metrics of quality, such as an end-to-end delay threshold δ associated with a maximum violation probability π. If the fraction of packets experiencing an end-to-end delay greater than δ is smaller than π, the contract is fulfilled; otherwise, the supplier has not met the QoS objectives set in the SLAs. Our negotiation algorithm explicitly accounts for (δ, π) SLAs and, from this point of view, takes one significant step along the path marked by Zachariadis and Barria in [34], using effective bandwidth instead of raw-bandwidth to implement an advanced single-attribute negotiation.

We also account for traffic descriptors by referring to the two-parameter traffic conditioning description specified in the differentiated services standards [3] for media-oriented real-time applications, comprising the average traffic rate r and the burst size b. During the negotiation of end-to-end QoS (δ, π), the customer and the supplier also negotiate the traffic descriptors (r, b) which, in turn, will be enforced by traffic regulators when the negotiation is over and the traffic flows start the transmission phase. Thus, in this paper we present an application of our algorithm in a four-attribute negotiation scenario. However, our algorithm is general and it supports negotiation on an arbitrary number of attributes.

Our algorithm draws from the multi-agent negotiation literature the concepts of degree of concession and willingness to pay for a given set of QoS requirements, which represent the fundamental parameters of any negotiation strategy [6,16,29]. The exploration strategy of the space of possible negotiation outcomes is instead novel, as it is designed to accommodate non-linear QoS utility functions, as opposed to the traditional approach of selecting linear or quasi-linear utility functions. The degree of concession and willingness to pay are consistent with the policy-based network management paradigm [18]. According to this paradigm, the automation of transmission resource management is obtained by delegating to the network infrastructure the policies for bandwidth allocation and service provisioning specified by customers and suppliers.

In studying the runtime negotiation of QoS targets and traffic profiles in IP differentiated services [4,12,26], we proceed along the path followed in [1]. In this work, the authors study the pricing of differentiated services by considering both the TCP and non-controlled real-time traffic carried over the assured forwarding per hop behavior [13], where performance is mainly determined by the adaptive packet dropping executed by the active queue management modules in network buffers, as pointed out in [10,11]. In this paper, we complement this analysis by considering applications that are very sensitive to packet dropping, such as video broadcasting, real-time interactive and media streaming. For these applications, active queue management is not recommended [3], as the packet loss rate needs to be negligible (for media streaming, the assured forwarding per hop behavior with active queue management can also be used). A possible choice for the transport of media-oriented traffic is the set of class selector per hop behaviors, with the two-parameters (r, b) traffic conditioning and rate-based scheduling. In the differentiated services architecture, a rate-based scheduler assigns to each per hop behavior a fraction of the link capacity. Common rate-based schedulers guarantee a minimum share of capacity to each per hop behavior and if a service class does not consume its allocated capacity, other service classes can use it. With reference to [1], we consider an additional set of per hop behaviors suitable for the transport of media-oriented real-time applications.

The negotiation algorithm proposed in this paper is one of the elements of the framework that must be deployed to implement the end-to-end negotiation process. In particular, negotiation must be conveyed through a signaling protocol that, in turn, is implemented by architectural elements both on the user's and the provider's side. A possible framework for the implementation of an integrated negotiation, pricing, and QoS adaptation for multimedia applications has been proposed by Wang and Schultzrinne in [33]. This framework, referred to as resource negotiation and pricing (RNAP), specifies the architectural components and protocol mechanisms. The specific negotiation algorithm is not constrained by the RNAP framework, where our algorithm can be implemented by carrying our negotiation requests and responses in the RNAP's protocol messages.

Our algorithm is tested through simulation to analyze the agreement solutions that are reached by negotiating partners and verify their soundness within the IP differentiated services architecture. To assess the benefits of negotiation, the efficiency of the agreement solutions reached through negotiation and the utility of the negotiating parties are compared with those obtained with a traditional non-negotiated capacity allocation policy. Moreover, we compare the bandwidth efficiency of our multi-attribute one-to-one negotiation approach with that of the single-attribute negotiation of raw-bandwidth.

The presentation is organized as follows. In Section 2, we position our negotiation algorithm in the framework of the IP differentiated services architecture. In Section 3, we describe our negotiation algorithm. The non-linear quality and utility functions of customers and suppliers are described in Section 4. The parameters used in our simulation model are explained in Section 5. Section 6 presents simulation results and conclusions are drawn in Section 7.