Available online at www.sciencedirect.com
SCIENCE DIRECT®
Computer Networks 49 (2005) 593–619
Computer Networks
www.elsevier.com/locate/comnet

Inter-domain routing: Algorithms for QoS guarantees

Samphel Norden *

Lucent Bell Labs, Room 4F-529, 101 Crawfords Corner Road, Holmdel, NJ 07733, USA

Received 16 March 2004; received in revised form 16 November 2004; accepted 28 January 2005
Available online 23 March 2005

Responsible Editor: E. Ekici

Abstract

Quality-of-Service routing satisfies performance requirements of applications and maximizes utilization of network resources by selecting paths based on the resource needs of application sessions and link load. QoS routing can significantly increase the number of reserved bandwidth sessions that a network can carry, while meeting application QoS requirements. Most research on QoS routing to date, has focused on routing within a single domain. BGP, the de facto standard for inter-domain routing provides no support for QoS routing, and has well-documented performance related issues that lead to its inadequacy to support QoS. This paper proposes new approaches to inter-domain routing for sessions requiring guaranteed QoS. The performance-scalability tradeoff is explored via extensive experiments on the proposed algorithms. Our extensive experiments on realistic intra-domain ISP topologies as well as inter-domain settings, show that the proposed algorithms achieve at least an order of magnitude gain in performance (blocking probability) over current mechanisms, while remaining scalable and easy to deploy.

Keywords: Quality of service; Reservations; Inter-domain; Routing

1. Introduction

The need for timely delivery of real-time information over local and wide area networks is becoming more common due to the rapid expansion of the Internet user population in recent years, and the growing interest in using the Internet for telephony, video conferencing and other multimedia applications. Choosing a route that meets the resource needs of such applications is essential to the provision of the high quality services that users are coming to expect.

In this context, it is important to distinguish *datagram* and *flow* routing. In datagram routing, packets of a session may follow different paths to the destination. In *flow routing*, all packets belonging to an application session follow the same path,

* Fax: +1 732 332 6744.
E-mail address: norden@lucent.com

1389-1286/$ - see front matter
doi:10.1016/j.comnet.2005.01.017

allowing bandwidth to be reserved along that path, in order to ensure high quality of service. Because many thousands or even millions of packets are typically sent during a single application session, flow routing occurs far less often than datagram routing, making it practical to apply more complex decision procedures than can be used in datagram routing.

The current Internet follows the *datagram* routing model and relies on adaptive congestion control to cope with overloads. Internet traffic is forwarded on a best-effort basis with no guarantees of performance. This can result in wide variations in performance, resulting in poor service quality for applications such as voice and video. Furthermore, Internet routing is typically topology-driven instead of being load-driven. This approach does not allow traffic to be routed along alternative paths, when the primary route to a destination becomes overloaded. While the application of load-sensitive routing to datagram traffic can cause hard-to-control traffic fluctuations, it can be successfully applied to flow routing, since reserved bandwidths sessions typically have holding times of minutes, effectively damping any rapid fluctuations in routes. Note that the use of the term "flow" in the rest of this paper also applies to aggregates of smaller "micro-flows" which are often bundled, when routing across domains.

The most prominent inter-domain routing protocol in the current Internet is the Border Gateway Protocol (BGP) [1]. BGP is a path vector based protocol, where a path refers to a sequence of intermediate domains between source and destination routers. BGP suffers from a number of well-documented problems, including long convergence times [2,3] following link failures. BGP adopts a policy based routing mechanism whereby each domain applies local policies to select the best route and to decide whether or not to propagate this route to neighboring domains without divulging their policies and topology to others.

The immediate effect of the policy based approach is to potentially limit the possible paths between each pair of Internet hosts. BGP does not ensure that every pair of hosts can communicate even though there may exist a valid path between the hosts. Also, since every domain is allowed to use its own policy to determine routes, the final outcome may be a path that is locally optimal at some domains but globally sub-optimal due to the lack of a uniform policy or metric used to find an *end-to-end* route. This point is highlighted by [4,5], where a majority of paths that are picked by BGP do not represent the optimal end-to-end paths. The authors define "optimal paths" by *hop count* in [4]. Their results show that for 50% of the BGP paths, there exists an alternate path with at least 5 less hops. In [5], different measures of path quality such as loss rate, bandwidth and round-trip time, consistently indicate that 30-80% of paths actually have an alternate path with significantly superior quality than the default path chosen by BGP.

The sub-optimality of BGP is primarily due to most domains defaulting to *hot potato* routing, in which each domain in the end-to-end path, tries to shunt packets as quickly as possible to the next network in the path, rather than selecting routes that will produce the best end-to-end performance for users. This characteristic is clearly undesirable, even for datagram traffic, and is particularly problematic for sessions that require high quality of service. Thus, there is clearly a critical need for a QoS routing mechanism that allows guarantees across domains.

Most research in QoS routing has focused on routing within a single domain. While the intra-domain problem is important, it is arguably even more important to address the QoS routing problem at the inter-domain level. Also, it is not feasible to directly extend protocols for intra-domain routing to the inter-domain context. While scalability is an even larger concern due to the sheer number of nodes and domains, the issue of *peering relationships* can constrain the nature and periodicity of information exchanged between domains. Whereas, an intra-domain protocol operates in a smaller network where all routers cooperate such that information about the entire topology can be conveyed to each router, this is not possible when routing across domains. Providing QoS routes across domains is made harder by the fact that each domain has a constraint on the information that it exchanges with other domains. This can affect the routes that are advertized by BGP