Traffic Engineering Working Group
Internet Draft
Document: <draft-wlai-tewg-cap-eng-00.txt>
Expiration Date: September 2000

W. Lai
AT&T

March 2000

Capacity Engineering of IP-based Networks with MPLS

Status of this Memo

Abstract

Drawing from the experience and work done in ITU-T on traffic engineering, the feasibility of extending telecommunications network dimensioning methods to IP-based networks with MPLS is explored. This version of the document is preliminary, reporting work in progress.

1. Introduction

The routing protocols currently used in public IP networks are typically traffic-insensitive in that they do not include network resource utilization information in making routing decisions. As a result of this lack of traffic control, traffic tends to converge onto the same network segments, thereby causing unbalanced network loading with subsets of network resources being congested and other resources being underutilized. This shortcoming of network operation, coupled with the phenomenal growth of Internet usage, makes it very difficult to manage IP-based network performance. Hence, current engineering practice sometimes resorts to over-

provisioning of network capacity.

As IP-based network services evolve from a single best-effort class toward differentiation with multiple levels of quality of service

(QoS), a network infrastructure that offers consistent and predictable network performance is required. To achieve this, proper control and provisioning of network resources would be needed. Recently, the multiprotocol label switching (MPLS) technology [1, 2] has been developed to overcome the above limitations of existing destination-based routing protocols so that more effective traffic engineering of IP-based networks can be performed.

MPLS integrates the connection-oriented label-swapping operation with connectionless network-layer routing, thus gaining much of the control capabilities of a connection-oriented network such as ATM-based B-ISDN. In particular, MPLS facilitates the establishment of a set of paths satisfying certain constraints [3]. Such a kind of path, called a constraint-based routing label-switched path (CRLSP), is conceptually very similar to a virtual-path connection in an ATM network for traffic engineering purposes. For example, a CRLSP enables the implementation of connection control and resource allocation functions on aggregated groups of traffic flows, instead of on individual flows. This allows support for different classes of service [3, 4] and fits well with the Differentiated Services (DiffServ) architecture [5, 6].

Furthermore, the ability of MPLS to set up explicitly routed CRLSPs permits source-node control, i.e., source routing, to operate efficiently. This provides network administrators the capability to steer traffic on desired paths. Such flexibility to direct the traffic to follow engineered paths through a network offers a number of advantages, e.g., load sharing across multiple paths, and shifting traffic away from congested links. The ability to use fixed preferred paths for routing traffic, so-called route pinning, also gives the means to measure and capture the statistics of the traffic associated with each source-destination node pair. As a result, there is the potential to develop more refined measurement technology so as to obtain a more accurate estimation and characterization of the traffic demands for network planning purposes.

The rest of this document is organized into sections describing general aspects of traffic engineering, capacity engineering, traffic demand characterization, performance objectives and reference connections, and network dimensioning. Much of the

materials presented here are adapted and drawn from the experience and work done in ITU-T on traffic engineering of telecommunications networks.

2. Traffic Engineering of IP-based Networks

The issue of IP network traffic engineering has been addressed by two recent documents [7, 8]. As defined in [7], Internet traffic engineering is ôthat aspect of Internet network engineering that deals with the issue of performance evaluation and performance

optimization of operational IP networks.ö Since network congestion is a primary cause of performance degradation, a major objective of traffic engineering is to increase the efficiency of resource utilization while minimizing the possibility of congestion through capacity and traffic management. Reference [7] also includes a discussion of requirements for routing, traffic mapping, measurement, survivability, and other attributes, from the perspective of traffic engineering.

The interactions between traffic and capacity management in controlling a networkÆs response to traffic demands and network failures are further explored in [8]. Real-time traffic management functions such as routing control, path selection, and resource management, ensures that performance objectives are met under all conditions including load shifts and failures. Capacity management, through control of network design, ensures that network provisioning meets performance objectives for a given set of traffic demands at minimum cost.

A global problem of network engineering is the choice and implementation of a specific routing strategy for the selection of paths for carrying traffic. Reference [8] provides an overview and categorization of various routing methods, as well as routing table management methods for different network types and their interworking with each other. Also described are resource management methods to achieve QoS objectives, such as connection admission and bandwidth reservation.

Taking account of the provisioned network capacity, routing patterns (i.e., path sets and rules for path selection) are designed for different traffic streams between different source-destination node pairs. Periodically or possibly on a real-time basis, these routing patterns may be adjusted as necessary to correct service problems. An iterative network design process encompassing both routing-pattern design and capacity allocation is used to determine the

minimum network capacity required to meet the performance objectives for given traffic demands.

Reference [8] also presents a set of traffic engineering operational requirements. For example, network management controls such as call gapping can be used to assure acceptable network performance in case of overload or failures.

3. Capacity Engineering

According to ITU-T Recommendation E.600 [9], traffic engineering includes measurements, forecasting, planning, dimensioning, and performance monitoring. Traffic engineering has a goal of ensuring trafficability performance objectives for telecommunications services. Trafficability performance is defined in Recommendation E.800 [10] as follows. For each individual network element or functional subsystem, it is the ability of the element to meet a

traffic demand of a given size and other characteristics, under given internal conditions. These internal conditions may be, e.g., any combination of faulty and not faulty parts within the element.

Thus, trafficability performance has a direct impact on the accessibility, retainability, and integrity of any given service offered by a network; it is one of the major factors in QoS. Service accessibility refers to the ability of a service to be obtained, within specified tolerances and other given conditions, when requested by the user. Service retainability refers to the ability of a service, once obtained, to continue to be provided under given conditions for a requested duration. Service integrity refers to the degree to which a service is provided without excessive impairments, once obtained.

Generally, trafficability performance is an attribute of network performance and can be described in terms of measures such as losses and delays. However, E.800 has not explicitly provided any specific measures for trafficability performance.

In this document, the focus is primarily on the dimensioning aspect of traffic engineering and its related tasks. In mapping given user demands onto network resources, network dimensioning involves the sizing of the network elements, such as links and buffers, so that performance objectives can be met at minimum cost. Thus, two major inputs to the dimensioning process are characterization of user demands and specification of performance objectives. (While cost is an important factor, the development of cost models will not be

dealt with in this document.) The term capacity engineering is used here to cover this subset of traffic engineering tasks related to dimensioning. These tasks are covered in more detail in the next three sections.

4. Traffic Demand Characterization

Typically, dimensioning procedures are based on models that approximate the statistical behavior of network traffic in large populations of users. To allow straightforward characterization of the traffic demands, these models necessarily adopt some simplifying assumptions concerning the usually complicated traffic processes, such as the arrival patterns of flows and the distribution of flow sizes. For these assumptions to be relevant and applicable, they must give rise to statistical patterns that closely approximate the behavior of aggregate traffic flows in operational networks. Traffic data are collected to validate these assumptions, with modifications being made when needed. Additionally, traffic measurements are used to estimate offered load and to provide forecasting of future demands for capacity planning purposes. Forecasting and planning may result in capacity augmentation or may lead to the introduction of new technology and architecture.

Thus, a first step in the dimensioning process is to develop user demand models as input to characterize the offered load to the network. In the context of an ATM-based B-ISDN, ITU-T Recommendation E.716 [11] describes the characterization of user demand as manifested at the user-network interface. Since a CRLSP in MPLS is conceptually similar to a virtual-path connection in ATM for traffic engineering purposes, the methodology of E.716, when suitably modified and extended, may be applicable to MPLS-based IP networks. Further investigation of this feasibility is needed.

To allow the characterization of traffic offered to an IP-based network, user demands may be modeled as an arrival process of demands for IP-based services of different types. As each such service demand generates a set of flows, each service type may be defined by a set of flow attributes and by a flow pattern as follows. (1) Flow attributes are those attributes of the service demand that identify the resources needed by the service demand in the network. These may include, e.g., access channel rate, communication configuration such as point-to-point or multipoint, and traffic conditioning agreements as defined in DiffServ [6]. (2) A flow pattern describes the packet arrival process in a flow

through a set of traffic variables. E.716 presents four approaches for defining these traffic variables to describe the transient nature of rate variations. For example, they may be related to the burst structure of the packet flow, the number of packet arrivals in time intervals of specified length, packet interarrival time, or the number of packet arrivals exceeding a given rate. (Currently, models that describe the self-similar nature of IP traffic have been proposed in the literature. However, for capacity engineering purposes, models that capture the burstiness characteristics of source traffic will suffice.)

To recapitulate, a set of flow attributes and traffic variables may be used together to characterize a particular service type by setting appropriate values for the parameters. Depending on the different values chosen for these attributes and variables, the number of service types can potentially be very large. To minimize the effort of traffic engineering, especially in the initial stage of deployment, it may be desirable to limit the total number of service types. For example, service types with similar values may possibly be combined to form one representative type that captures the salient features essential for dimensioning.

5. Performance Objectives and Reference Connections

Performance objectives can be viewed from two perspectives: the user and the network service provider. Quality of service (QoS), which is performance perceivable by a user of a service, expresses the user's degree of satisfaction of the service [9, 10]. Thus, QoS parameters focus on performance effects that are observable at the service access points and network interfaces, rather than their causes within the network. Different service types usually have

different QoS requirements. This allows a network provider to provide different treatment to different service types, to gain higher resource utilization.

Grade of service (GoS) is a number of traffic engineering parameters to provide a measure of adequacy of a group of resources under specified conditions [9]. These GoS parameters may be probability of blocking, probability of delay, etc. They are essential for network internal design and operation, as well as specification of component performance. The latter is related to the trafficability performance described previously.

Based on a given set of QoS requirements, a set of GoS parameters are selected and defined on an end-to-end basis within the network

boundary, for each major service category provided by a network. The selected GoS parameters are specified in such a way that the GoS can be derived at well-defined reference points, i.e., traffic significant points. This is to allow the partitioning of end-to-end GoS objectives to obtain the GoS objectives for each network stage or component, on the basis of some well-defined reference connections.

As defined in E.600, for traffic engineering purposes, a connection is an association of resources providing means for communication between two or more devices in, or attached to, a telecommunication network. There can be different types of connections as the number and types of resources in a connection may vary. Therefore, the concept of a reference connection is used to identify representative cases of the different types of connections without involving the specifics of their actual realizations by different physical means.

Typically, different network segments are involved in the path of a connection. For example, a connection may be local, national, or international. The purposes of reference connections are for clarifying and specifying traffic performance issues at various interfaces between different network domains. Each domain may consist of one or more service provider networks. Recommendation E.651 [12] specifies reference connections for IP-access networks. Other reference connections are to be specified.

From the QoS objectives, a set of end-to-end GoS parameters and their objectives for different reference connections are derived. For example, end-to-end connection blocking probability and end-to-end packet transfer delay may be relevant GoS parameters. The GoS objectives should be specified with reference to traffic load conditions, such as under normal and high load conditions. The end-to-end GoS objectives are then apportioned to individual resource components of the reference connections for dimensioning purposes. In an operational network, to ensure that the GoS objectives have been met, performance measurements and performance monitoring are required.

6. Network Dimensioning

As discussed previously, network dimensioning involves the optimal (e.g., minimum cost) sizing of the network elements to accommodate given traffic demands, while meeting performance objectives. In using MPLS, this means the dimensioning of a set of pre-selected CRLSPs for carrying traffic, and mapping the logical network of

CRLSPs onto a physical network of links with capacity constraints. The dimensioning process also determines the link capacity parameters or thresholds associated with the use of some bandwidth reservation scheme for service protection. Service protection controls the GoS for certain service types by restricting access to bandwidth, or by giving priority access to one type of traffic over another. Such methods are essential, e.g., to guarantee a minimum amount of resources for connections with expected short duration, to improve the blocking probabilities for connections with high bandwidth requirements, or to maintain network stability by preventing GoS degradation in case of a local overload.

In performing the task of dimensioning, it is assumed that a network topology, both at the logical and physical levels, has been defined. This is because the layout of a network is usually influenced by other factors, such as the network providerÆs policy/administrative constraints, or considerations of an existing network. Also, it is assumed that the network is available, i.e., it does not consider network equipment in a failure state.

Routing deals with the selection of network paths for connection requests. To simplify the dimensioning process, a fixed routing-pattern with pre-determined paths for different traffic streams is usually assumed. As described previously, the process of routing-pattern design and dimensioning is iterated until an optimal design is reached.

The series of ITU-T Recommendations E.735-7 [13, 14, 15] presents a set of general principles and methods for dimensioning ATM-based B-ISDNs. The notion of Equivalent Cell Rate (ECR) has been used effectively therein for dimensioning purposes. The ECR captures the effects of expected traffic mix, cell-level control mechanisms, priority scheduling, bandwidth and buffer capacity limitations, thereby characterizing the estimated amount of resources that needs to be allocated to a connection to satisfy the specified cell-level GoS objectives. Borrowing from this technique, it may be useful to define some Equivalent Bandwidth parameter for dimensioning IP-based networks. Further studies are needed for its definition.

Assuming that connection blocking probability is the only GoS parameter at the connection level, and using ECR to account for cell-level performance, E.737 presents several iterative methods for dimensioning. To reduce computational complexity, approximation methods based on the independence assumption and the principle of decomposition may be developed. For example, if the capacity of a link can be considered as well delimited, independent of the traffic

carried by other links, then the global end-to-end decision on connection admission can be decomposed into local decisions. These techniques may be applicable to the dimensioning of CRLSPs in an MPLS-based network.

7. Security Considerations

Security considerations are not addressed in this version of the draft.

8. References

[1] E.C. Rosen, A. Viswanathan, and R. Callon, ôMultiprotocol Label Switching Architecture,ö Internet-Draft, Work in Progress, August 1999.

[2] R. Callon, P. Doolan, N. Feldman, A. Fredette, G. Swallow, and A. Viswanathan, ôA Framework for Multiprotocol Label Switching,ö Internet-Draft, Work in Progress, September 1999.

[3] D. Awduche, J. Malcolm, J. Agogbua, M. OÆDell, and J. McManus, ôRequirements for Internet Traffic Engineering Over MPLS,ö RFC 2702, September 1999.

[4] T. Li and Y. Rekhter, ôA Provider Architecture for Differentiated Services and Traffic Engineering (PASTE),ö RFC 2430, October 1998.

[5] K. Nichols, S. Blake, F. Baker, and D. Black, ôDefinition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers,ö RFC 2474, December 1998.

[6] S. Blake, D. Black, M. Carlson, E. Davies, Z. Wang, and W. Weiss, ôAn Architecture for Differentiated Services,ö RFC 2475, December 1998.

[7] D.O. Awduche, A. Chiu, A. Elwalid, I. Widjaja, and X. Xiao, ôA Framework for Internet Traffic Engineering,ö Internet-Draft, Work in Progress, January 2000.

[8] ITU-T Draft Recommendation E.TE, ôTraffic Engineering and QoS Methods for IP-, ATM-, and TDM-Based Multiservice Networks,ö March 2000 (Contact: G. Ash).

[9] ITU-T Recommendation E.600, ôTerms and Definitions of Traffic Engineering,ö March 1993.

[10] ITU-T Recommendation E.800, ôTerms and Definitions Related to Quality of Service and Network

Performance Including Dependability,ö August 1994.

[11] ITU-T Recommendation E.716, ôUser Demand Modelling in Broadband-ISDN,ö October 1996.

[12] ITU-T Draft Recommendation E.651, ôReference Connections for Traffic Engineering of IP Access Networks,ö to be approved.

[13] ITU-T Recommendation E.735, ôFramework for Traffic Control and Dimensioning in B-ISDN,ö May 1997.

[14] ITU-T Recommendation E.736, ôMethods for Cell Level Traffic Control in B-ISDN,ö May 1997.

[15] ITU-T Recommendation E.737, ôDimensioning Methods for B-ISDN,ö May 1997.

10. Acknowledgments

The review and comments of Gerald Ash is much appreciated.

11. Author's Addresses

Wai Sum Lai
AT&T Labs
Room D5-3D18
200 Laurel Avenue
Middletown, New Jersey 07748, USA
Phone: 732-420-3712
Email: wlai@att.com

Full Copyright Statement

 mailto:mikega@microsoft.co