Optical Switching and Networking 23 (2017) 52–66

Contents lists available at ScienceDirect

Optical Switching and Networking

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

On control plane algorithms for Carrier Ethernet networks: Unicast, multicast provisioning and control traffic reduction

Saurabh Mehta, Sandeep Joshi, Tamal Das, Ashwin Gumaste*

Department of Computer Science and Engineering, Indian Institute of Technology, Bombay, Mumbai 400076, India

A R T I C L E I N F O

Article history:
Received 13 January 2016
Received in revised form
14 June 2016
Accepted 6 September 2016
Available online 13 September 2016

Keywords:
Carrier ethernet
Control plane

A B S T R A C T

Carrier Ethernet systems are actively replacing SONET/SDH networks for transport of provider data. After much standardization in both the IEEE and the IETF, Carrier Ethernet has evolved as a rich carrier-class technology. The standards do not focus on specific implementation aspects of providing guaranteed service-oriented features that are intrinsic towards successful replacement of SONET/SDH by Carrier Ethernet solutions. In this paper, we consider the engineering aspects of the Carrier Ethernet software defined control plane for provisioning, managing and maintaining services. Specifically, we propose methods for provisioning guaranteed unicast services considering the underlying network state. The problem is difficult due to the complex nature of admission control coalesced with subjective user requirement that is not easy to rationalize across a network. We also investigate the problem of sub-50 ms protection of multicast connections and propose a unique, tractable algorithm within the control plane for enabling carrier-class multicast services. Finally, we focus on overall control traffic reduction schemes. An intuitive merging algorithm is proposed that minimizes control traffic in Carrier Ethernet domains. A simulations study extensively evaluates the proposed techniques.

1. Introduction

Packet networks present service-centric opportunities for residential and enterprise users that network operators are beginning to exploit especially in the transport space. The transport space has been traditionally dominated by SONET/SDH. A packet network-enabled backbone facilitates excellent interconnection for a 4 G mobile backhaul as well as supports dynamic cloud requirements. In the present scenario, several SONET/SDH ring networks are being replaced with packet-oriented Carrier Ethernet (CE) networks. CE technology standardized by both the IEEE and IETF as Provider Backbone Bridging-Traffic Engineering (PBB-TE) [1] and MPLS-Transport profile (MPLS-TP) [2] respectively assures carrier-class behavior in a packet network. The treatment of CE in the standards only define the protocol – leaving the implementation open to vendors. The standards do not consider the enormity of implementing a packet-core that has the onus of supporting circuit-like *operation, administration, maintenance and provisioning* (OAM&P) functionality. The probabilistic behavior of a packet core implies that delay and jitter are inherently present in the network. Despite the probabilistic behavior of a packet-core fabric, it is possible to keep the delay bounded and performance deterministic.

A key to achieving deterministic performance is to keep the admission control tractable especially in a managed transport network. *Given the domain of CE, an important problem that has never been addressed before, is to control the admission of carrier-class requests at both edge and core nodes – especially at core nodes, which become contention points for provisioned services*. One of the key goals of this paper is to propose implementation-oriented (software) control algorithms for load-balancing of traffic to control multiplexing of contending traffic streams. We propose a set of mathematical and fuzzy-logic based algorithms for load-balancing in the domain of CE that take into consideration multiple network parameters for ensuring service guarantees.

A second goal of this paper is to extend the ambit of carrier-class protection for multicast services. The problem of sub-50 ms restoration [3] of multicast services is challenging from an implementation as well as from a theoretical perspective. In this paper, we propose a multicast protection scheme and an algorithm for enabling point-to-multipoint (P2MP) and multipoint-to-multipoint (MP2MP) protection that is quintessential to the proliferation of CE as an enabling technology for video distribution. Our proposed algorithm facilitates the creation of two completely disjoint trees – a complex graph theoretic problem. The advantage of this scheme, which involves using two disjoint trees for protection, is the simplicity of detecting the failure conditions as well

* Corresponding author.
E-mail addresses: saurabhm@iitb.ac.in (S. Mehta), sandeepj@cse.iitb.ac.in (S. Joshi), ashwing@ieee.org (A. Gumaste).

http://dx.doi.org/10.1016/j.osn.2016.09.001
1573-4277/

as enhancing the overall availability of the network due to spatial redundancy. Once a network failure is detected, the affected services can be switched to the backup path with minimum delay – thereby resulting in preservation of Quality of Experience (QoE) for the user. Using the combination of the provisioning algorithms and the protection scheme, centrally managed CE networks like PBB-TE and MPLS-TP can enable fast protection switching with low implementation complexity and can be relevant in the present context of software defined networks (SDNs).

A key attribute of CE networks, especially in the transport domain, is that these have a distinct in-band control plane that is used for purposes of OAM&P. The role of the control plane in CE networks is in fact quite critical from various aspects: admission of services, maintenance and operability of network equipment, protection and restoration of services. With new services, the control traffic also grows. Also, the types of services imply that control traffic growth can be non-linear vis-à-vis data-plane growth. The growth of control traffic can in fact be catastrophic in the sense that it can take away a sizable amount of data-forwarding-plane bandwidth. A method to reduce control traffic in CE networks was proposed by us in [4]. It is based on the concept of grouping partially overlapped links, with aggregation and de-aggregation of the control traffic at the non-overlap boundary nodes. We further extend the concept in this paper by proposing an improved and tractable heuristic. We perform a rigorous simulation analysis validating the need for control traffic reduction and qualitatively as well as quantitatively studying the reduction process.

This paper is organized as follows: In Section 2, we introduce efficient algorithms for load-balancing and traffic engineering for unicast point-to-point services. Section 3 introduces a method for enabling carrier class protection for multicast services and discusses algorithms for computing redundant multicast trees. Section 4 explores mechanisms to enhance the scalability of in-band control plane required to enable carrier-class features, characteristic of CE networks. The performance of the algorithms mentioned are summarized in Section 5. Section 6 concludes the paper.

2. Engineering unicast services in Carrier Ethernet networks

E-Line services as defined by the Metro Ethernet Forum (MEF) provide a point-to-point connection between a pair of dedicated *user-network interfaces* (UNIs). E-Line services constitute the bulk of CE services [5] and are categorized by low frame delay, minimal frame delay variation and low frame loss. In CE networks, mechanisms such as circuit emulation and resource reservation [6–9] are deployed to enable provisioned services to meet stringent delivery constraints. These mechanisms are necessary to guarantee service quality, but are not sufficient by themselves for bounded latency and jitter constraints.

To elaborate the above point, consider the following scenario. In CE networks, for admission control, services are provisioned with two parameters – the *Committed Information Rate* (CIR) and *Committed Burst Size* (CBS) [10]. CIR specifies the average bandwidth for a virtual circuit, while CBS imposes a limit on the burst of data that can be transmitted by a service. Consider the provisioning scenario for an 8-port 1 Gbps per-port switch as shown in Fig. 1. Five services with CIR = 100Mbps and CBS = 10 ms are incident on to the same output port. The services are well within the capacity constraints of the output port as the combined bandwidth is 50% of the output link capacity. If it is assumed that the services multiplexed at a node will not send a burst of data at the same time, we can meet the CBS requirements of all services with limited memory buffers. However, it may happen that all the services transmit a burst of data simultaneously, i.e. the transmitted traffic is 1 Gbps on each input port. This will lead to an instantaneous

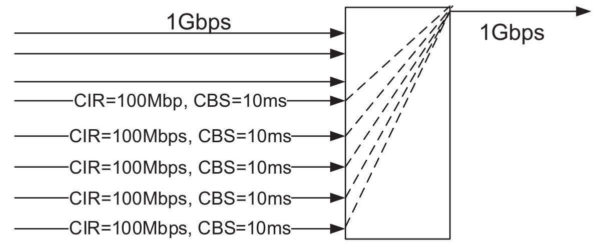

Fig. 1. Impact of high fan-in to latency and jitter in a network.

traffic burst of 5 Gbps, which is beyond the capacity of the output port. This will lead to buffering of data and the latency experienced by packets in this case can be as high as 40 ms.

This problem is more prominent when shortest paths algorithms are used to find the path for a given service. These algorithms calculate paths based on edge weights. The problem with a naive approach using static edge weights is that if sufficient resources are available, the same path is selected for a service between two nodes. This results in high resource utilization on shortest paths across the network, while other non-shortest paths are underutilized.

A simple solution to this problem is to allocate dedicated buffers to each service. However, this will allow us to provision fewer services or we will need large buffering capacity at a node. A better approach would be to provision services in a way that prefers under-utilized nodes/links while selecting paths. Like link utilization, there are many other factors to be considered while selecting the path for routing a service – Quality of transmission medium, node throughput, failure probabilities of links/nodes, statistical observations about traffic patterns, etc. We need traffic engineering techniques that are able to consider one or more of these factors together and enable us better control of the network.

The traffic engineering algorithms proposed in this section address this problem by providing the control required to balance the use of precious network resources. Additionally, they enable route diversity that minimizes the risk of a single link or device failure causing simultaneous disruption to both the primary and backup path in a network.

2.1. Related work

Algorithmic traffic engineering for unicast services is a well-investigated subject [11–15]. This area was investigated in the context of ATM circuit switched networks [16], load-balancing for IP networks [17], etc. With the advent of MPLS-TP and PBB-TE, new techniques have also been developed [18]. Widest shortest path [11] algorithm finds minimum hop paths and then chooses one with the maximum residual bandwidth. As it considers only shortest paths, it does not eliminate the above-mentioned problem. Another common practice is to dynamically assign edge weights and then use the shortest path algorithm to compute the route [12]. Some of the common techniques for assigning edge weights are:

1. *Minimum Interference Routing Algorithm (MIRA)*: MIRA [13,14] attempts to maximize the max flow between possible source-destination pairs by dynamically assigning edge weights using policy constraints to guide a shortest path algorithm – thus introducing path diversity in the network. However, MIRA considers only one parameter for traffic engineering at a given time.
2. *Threshold based edge weights* [15]: In this technique, a threshold is set for parameters being considered. All the edges that have values below a threshold are assigned an edge weight of 0,