Achieving Open Access in Ethernet PON (EPON)

Amitabha Banerjee and Biswanath Mukherjee
Department of Computer Science, University of California, Davis, California 95616
abanerjee@ucdavis.edu, mukherje@cs.ucdavis.edu

Glen Kramer
Teknovus Inc., Petaluma, California, 95954
glen.kramer@teknovus.com

Abstract: "Open access" is a regulatory requirement in many countries mandating that the residential access network infrastructure be competitively available to service providers. We propose dual Service-Level Agreements (SLAs) to enforce fairness in *open access* EPON.

OCIS codes: (060.4250) Networks

1. Introduction

Ethernet Passive Optical Networks (EPONs) are point-to-multipoint broadband access networks. An Optical Line Terminal (OLT) at the Central Office (CO) is connected to many Optical Network Units (ONUs) at remote ends using optical fiber and passive splitter. EPON has been standardized in the IEEE 802.3ah standards body [1]. Ethernet technology is inexpensive and ubiquitous, and hence EPON seems very promising today.

Government regulations in many countries have mandated a clear demarcation between the network operators who provide physical connectivity and transport data (e.g., telephone and cable companies), and the service providers (SPs) who deliver the content and the services (e.g., Internet SPs). The primary motivation is to have free-market competition between SPs, thereby making services cheap and available to end users; and also preventing monopoly by the network operator who owns the right-of-way to lay cable/fiber in a residential area. Such a framework is known as *open access* in which users are free to choose what services they need and the corresponding SPs, and the SPs are able to solicit subscribers through independent and competing marketing efforts [2].

Open access is a sustainable and profitable economic model for a broadband access solution such as EPON, because the cost of deployment may be recovered from multiple SPs [3]. Many residential communities and municipalities, seeing broadband data services as a key driver for economic growth, are deploying their own fiber networks. Viewing the broadband access network as a type of infrastructure similar to a town's plumbing/sewer system, the municipalities build and manage the infrastructure, facilitating access to broadband services, but not supplying the services themselves. *Open access* therefore is an extremely useful framework for future networks.

Open access networks should be designed to serve individual and non-cooperative users as well as independent and non-cooperative service providers. Different SPs may be delivering services to same or different sets of users. This leads to the requirement of fair allocation of network resources (especially bandwidth) amongst both SPs and users. This is a significant challenge, because SPs and users are located at opposite ends of the access channel. Fairness amongst users should guarantee some minimum degree of network performance (bandwidth) to all users, so that a user with killer bandwidth applications does not starve others out of bandwidth. Similarly fairness amongst SPs should ensure fair play in terms of bandwidth allocation amongst competing SPs. In this work, we investigate a dual Service-Level Agreement (SLA) method for achieving fairness in an *open access* framework, as applied to EPON. We note that EPON is considered here merely as an example due to its future promise; we strongly believe that our work may be applied to other access network technologies such as APON, GPON, and variants of DSL as well.

2. Fairness in an Open Access Framework

Currently deployed broadband access networks follow a best-effort model for sharing bandwidth amongst users. The main disadvantage of a best-effort model is that the service which the user receives is dependent on the current load in the network, which in turn depends on the bandwidth requirement of other users. Hence, a user can not be guaranteed some level of network performance. Although such a model has survived till now, primarily because of low bandwidth demand from users, it has discouraged growth of high-bandwidth application and services. The complex relationship between users, network operators, and SPs is depicted in Fig. 1. Because of such relationships, quality-of-service (QoS) aware services such as Video-on-Demand (VoD) are still not readily available in the current Internet. Access networks are the primary bottleneck, because metropolitan and long-haul networks are usually over-provisioned, so they are lightly loaded at most times.

A solution is to try to enforce Service-Level Agreements (SLAs) in the access network [4]. A typical SLA includes performance guarantees such as minimum bandwidth, maximum packet loss, maximum delay, etc. SLAs in the context of SPs, which define minimum bandwidth guarantees for a SP, shall enable a SP to provide QoS-aware services to its customers. For example, let us consider SP_a which provides a VoD service, which would require R_a bps average bandwidth per customer. Suppose SP_a negotiates a SLA of W_{MIN} bps guaranteed bandwidth. Then, SP_a would have a guarantee from the network operator to provide reasonable QoS to W_{MIN}/ R_a users at any time on average. If the number of users simultaneously requesting VoD from SP_a is greater than this amount, SP_a may block some requests to provide good service to the other users already in the system. Similarly, providing minimum guarantees to users will give the users reasonable services and would not starve any user out of bandwidth.

We define the following three categories of SLAs:

a. **Service Provider SLAs (SP SLAs):** These SLAs guarantee minimum amount of bandwidth to SPs. Guarantees may not be over-provisioned, i.e., they should not exceed channel capacity. This ensures that all SLAs may be met irrespective of the network load.
b. **User SLAs:** These SLAs guarantee minimum bandwidth to users. Guarantees may not be over-provisioned.
c. **Dual SLAs:** These SLAs incorporate both User and SP SLAs. Although User SLAs and SP SLAs are not individually over-provisioned, the sum of the User SLAs and SP SLAs may exceed the channel capacity. Hence, it may not be possible to meet both sets of SLAs for all traffic demands. Therefore, we need to distinguish between the two sets. The *primary SLA* is defined to be the one, whose specified minimum guarantees must be given the highest priority to be met. After the *primary SLA* has been met, the next priority is to meet the *secondary SLA*. If it is not possible to meet some *secondary SLAs*, the objective shall be to ensure fairness in the deficits in allocation of bandwidth from the secondary SLA.

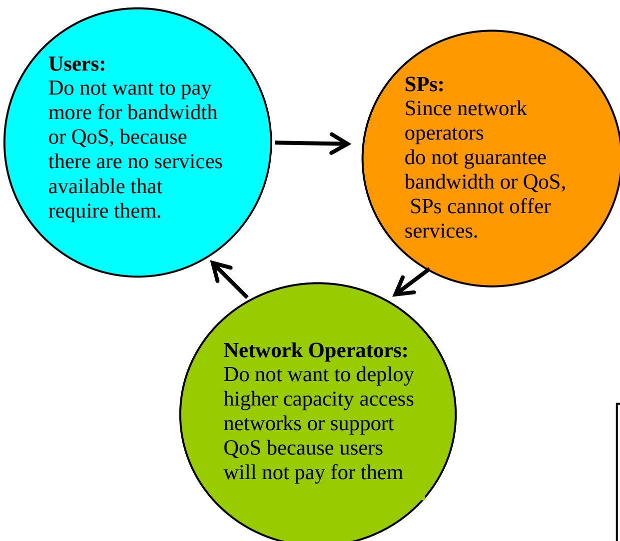

Fig. 1: Relationship between users, network operators, and SPs.

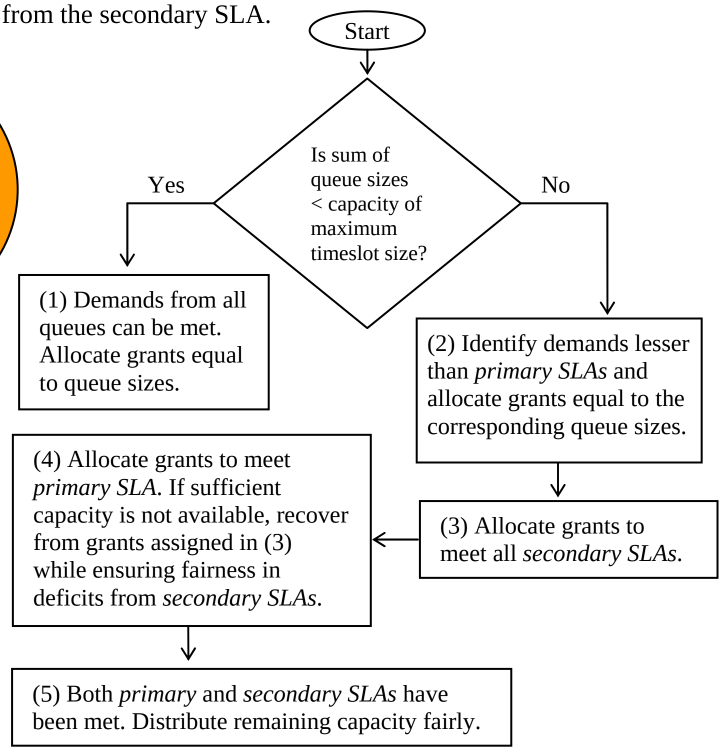

Fig. 2: Flowchart of proposed algorithm for meeting Dual SLAs.

3. Implementation and Illustrative Numerical Examples

In this paper, we consider the downstream traffic from SPs to users in an *open access* EPON. The OLT schedules the traffic for all the users. At the OLT, there exists a queue for traffic from each SP to each user. Considering a timeslot of a fixed size, the objective is to schedule traffic from each of the queues by granting some capacity in this timeslot, taking into consideration the SLA requirements, and distributing any surplus bandwidth after SLAs have been met fairly. This process is repeated for successive timeslots. The maximum size of the timeslot is fixed in order keep a bound on the delays. Implementing Dual SLAs is challenging because of the two different set of SLAs. In Fig. 2, we show a flowchart for our algorithm which meets Dual SLA demands. Due to constraints of space, we are unable to describe the algorithm in any greater detail here.

We simulate an EPON system with specifications in Table 1 for the three different categories of SLAs described above. The traffic matrix is shown in Table 2. SPs 2…6 have a total traffic demand of 150 Mbps. Therefore, a fair scheme should deliver equal bandwidth to the SPs. Similarly, we would expect Set III and Set IV of users to receive equal bandwidth. All traffic is generated to be self-similar with *Hurst parameter* of 0.8. Due to the bursty nature of Internet traffic, access networks are frequently subjected to very heavy load of traffic. Hence, in our simulations, the access network is heavily loaded at 1.39 (normalized to the line rate of 1 Gbps). We observe the performance of the EPON system for a period of 5 minutes, noting the bandwidth received by each user and the bandwidth allocated to each SP at intervals of 1 second. In Figs. 3(a) and 3(b), we show the average bandwidths assigned to SPs and user sets for the three different categories of SLAs.

Table 1: *Open access* EPON parameters

Parameter	**Value**
Rate of EPON	1 Gbps
Number of SPs	6
Number of users (ONUs)	32
SP SLA	150 Mbps per SP
User SLA	25 Mbps per user
Dual SLA	150 Mbps per SP (*primary*), 25 Mbps per user (*secondary*)

Table 2: Traffic matrix

SPs	**User Set**	**Traffic Rate**	**Example**
1	Set I (#1 – #20) Set II (#21 – #26) Set III(#27– #29) Set IV(#30– #32)	20 Mbps	2 DVD[1] quality video streams.
2, 3, and 4	Set II (#21 – #26)	30 Mbps	1 HDTV[2] + 1 DVD
5	Set III(#27– #29)	50 Mbps	2 HDTV + 1 DVD
6	Set IV(#30– #32)	50 Mbps	2 HDTV + 1 DVD

1 Average bandwidth for a MPEG-2 coded DVD video stream is 10 Mbps.
2 Average bandwidth for a compressed HDTV quality video stream is 20 Mbps.

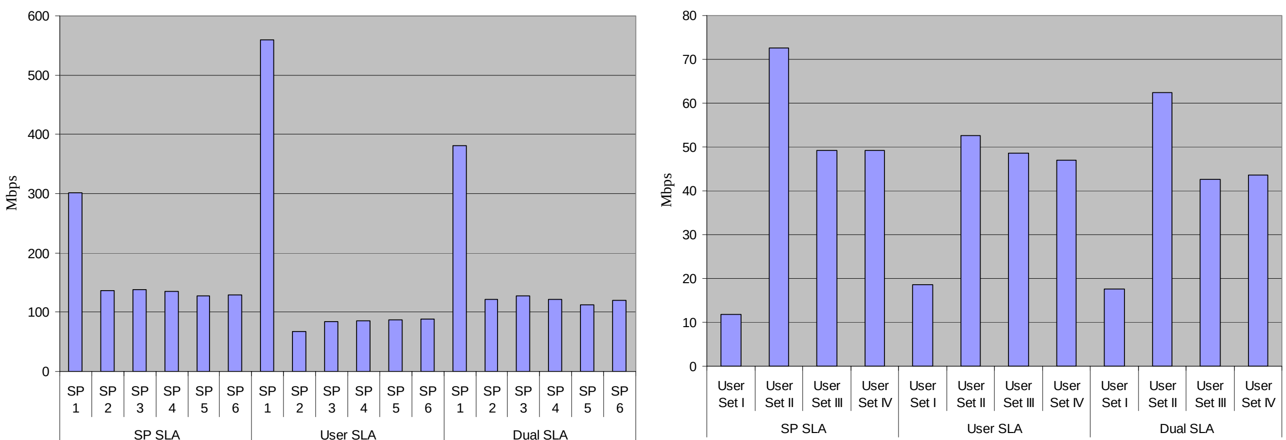

Fig. 3: Bandwidth assigned to (a) SPs and (b) users for different SLAs.

We observe that SP SLA grants almost equal amounts of bandwidth to SPs 2…6. However, User Set I receives an unfair share (close to 10 Mbps) of bandwidth in this scheme. The User SLA, while granting good distribution of bandwidth amongst the users, leads to poor distribution amongst SPs (SPs 2…6 receive less than 100 Mbps while SP 1 receives close to 600 Mbps). The Dual SLA achieves good bandwidth distribution amongst both SPs and users. We observe that Dual SLA performs close to SP SLA for the SPs and close to User SLAs for the users, while meeting both *primary* and *secondary* SLA requirements. More results are not included here due to space limitation.

4. Summary

In this paper, we showed that having SLAs in an access network will help in the deployment of QoS-aware services for residential customers. We proposed the concept of Dual SLA and demonstrated that it can ensure fairness to both SPs and users in an open access EPON (and the results can be applied to other access networks as well).

References:

[1] Clause 64, 65, IEEE 802.3ah standard, approved June 24, 2004.
[2] A. Banerjee and M. Sirbu, "Towards Technologically and Competitively Neutral Fiber to the Home (FTTH) Infrastructure," available at http://100x100network.org/papers/banerjee-tprc2003.pdf
[3] "Introducing the open access network," Telephony Online, June 9, 2004, available at http://telephonyonline.com
[4] NSF Workshop Report, "Residential Broadband Revisited: Research Challenges in Residential Networks, Broadband Access, and Applications," October 2003, available at http://cairo.cs.uiuc.edu/nsfbroadband/