The Importance of Dynamic Bandwidth Allocation in GPON Networks

White Paper

by Onn Haran, Fellow
and Amir Sheffer, Product Line Manager

Issue No. 1: January 2008

PMC-Sierra, Inc.

Legal Information

Copyright

PMC-2072146 (R1)

Disclaimer

Trademarks

Patents

Abstract

This article describes Bandwidth Allocation in the PON network and highlights the advantages of Dynamic Bandwidth Allocation (DBA). A well-defined DBA algorithm can significantly improve network performance, provide a means of flexibly tailoring network responsiveness and enable a service provider to generate more revenue from their FTTH networks without boosting raw bandwidth by increasing the percentage of acceptable oversubscription.

A GPON Bandwidth Allocation algorithm allows the Optical Line Terminal (OLT) to control upstream bandwidth allocation to the Optical Network Terminals (ONT). This article compares currently available Bandwidth Allocation methods and their impact on the service provided to the subscriber. By simulating a real-life TCP download scenario, it shows that Status Reporting (SR) DBA is a superior tool to provide a high level of quality of service.

When using Static Bandwidth Allocation or Non-Status Reporting (NSR) DBA, operator visibility into the upstream channel is limited. The PON access network can easily turn into a bottleneck, limiting overall network performance, as the inability to control upstream latency can result in latencies of 10mS and longer in the upstream channel alone. In addition, packet loss is inevitable using NSR DBA. As a result, the operator cannot oversubscribe the network, cannot provide symmetrical 100 Mbps services and cannot maximize its revenue potential.

SR-DBA provides excellent visibility into the upstream channel. When using SR DBA, tight control reduces upstream latency by 90% and limits it to the 1mS range, enabling symmetrical services at 100 Mbps and beyond. The SR-DBA supports oversubscription and can allocate bandwidth equivalent to 5 Gbps and beyond, leveraging the traffic burstiness characteristics.

About PMC

PMC-Sierra is a leading provider of high-speed broadband communications and storage semiconductors and MIPS-Powered processors for Enterprise, Access, Metro Optical Transport, Storage Area Networking and Wireless network equipment. The company offers worldwide technical and sales support, including a network of offices throughout North America, Europe and Asia. The company is publicly traded on the NASDAQ Stock Market under the PMCS symbol and is included in the S&P 500 Index.

Revision History

Issue No.	Issue Date	Details of Change
1	January 2008	Document created.

Table of Contents

List of Figures

1 Introduction

GPON is a Passive Optical Network (PON) operating at 2.488 Gbps downstream and 1.244 Gbps upstream rates, standardized by ITU-T G.984.

The GPON network includes an optical line terminal (OLT) residing in the Central Office and multiple Optical Network Terminals (ONT) residing at or near the customer premise — on the curb outside, in a building, or at the subscriber residence. As shown in Figure 1 , GPON is typically deployed as a tree or tree-and-branch topology, using a single, shared optical fiber and inexpensive optical splitters which divide the single fiber into separate strands feeding individual subscribers. The OLT is connected to the IP network and backbone of the network operator. Through this interface, multiple services are provided to the access network.

One of the key advantages of PON networks over point-to-point fiber networks is that a single fiber can be shared between many subscribers, greatly reducing the cost of deploying and maintaining the broadband access network. This makes PON technologies a very attractive option for service providers replacing or augmenting costly copper networks with fiber-based networks. A key feature of the point-to-multipoint network is that it is practical for up to 64 ONTs to share a single fiber (the standard supports up to 128). However, sharing of a single fiber requires careful design of the bandwidth allocation in the network. The GPON downstream arm is a natural point-to-multipoint topology and supports services like any other Ethernet network. The upstream channel is TDM based; the OLT controls the transmission of each ONT using a Bandwidth Allocation mechanism or algorithm. This mechanism determines the length and location of each ONT upstream burst. An explicit Bandwidth Allocation specification was intentionally left out of the ITU-T G.984 standard, which was limited to describing the mechanisms used for allocating upstream bandwidth.

DBA algorithms are well explained in many well-known technical papers. They have been used in high-volume EPON networks with very strict requirements. HFC networks use DBA to maximize the upstream of the COAX shared access. The performance and stability of DBA algorithms are field proven.

Figure 1 GPON Network Simplified Diagram

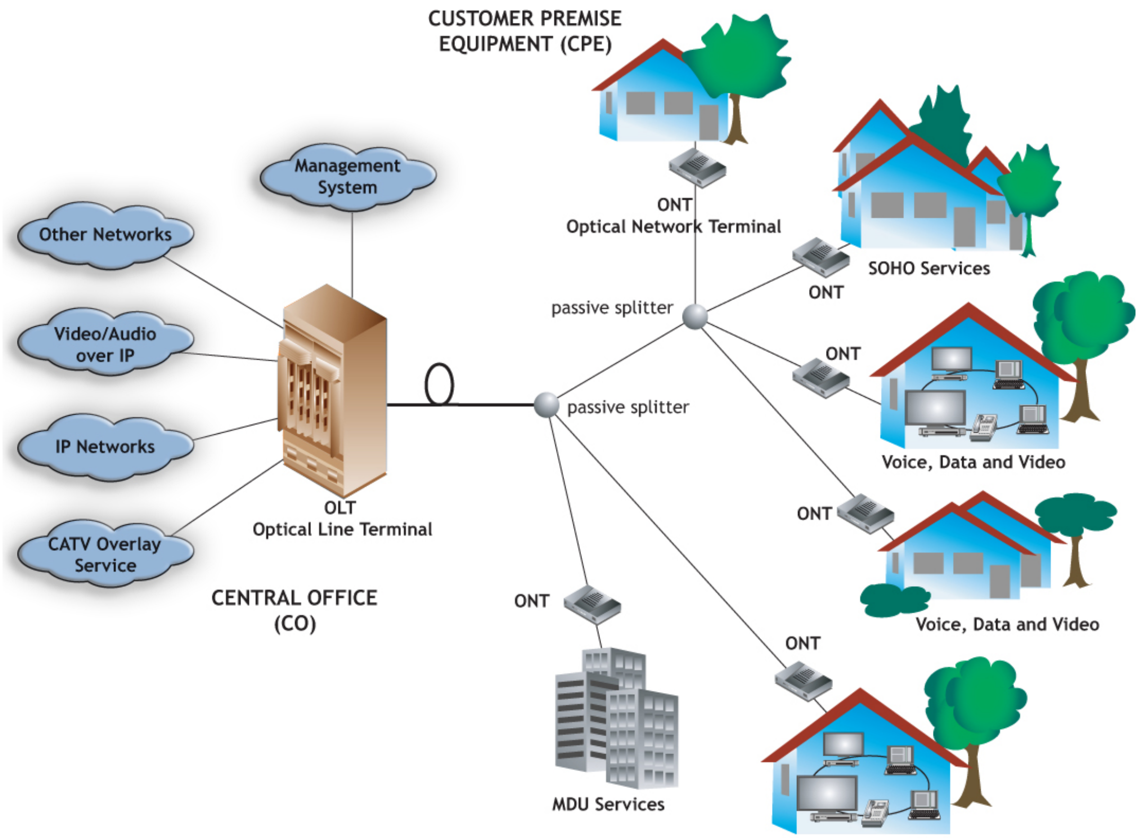

As the nature of the network traffic and services continues to change, an efficient upstream bandwidth management mechanism becomes ever more important to the service provider, who cannot afford to manually alter the broadband access network characteristics as requirements change. The popularity of peer-to-peer (P2P) applications is growing, and P2P traffic is forecasted to quadruple between 2006 and 2011[1]. Also, video and picture upload to sites such as YouTube or Flickr is growing rapidly. While in the past, network usage was asymmetric – downstream required much more bandwidth than upstream – this is changing. As seen in Figure 2 , upstream bandwidth consumption is increasing over time. In addition, bandwidth usage is changing as exchanged files are getting bigger and bigger. The service provider needs an efficient mechanism such as dynamic bandwidth allocation to tailor network performance based on ever changing user requirements.

[1] Source: Cisco Systems, "The Exabyte Era", 2007

Figure 2 Global Consumer Internet Traffic Forecast

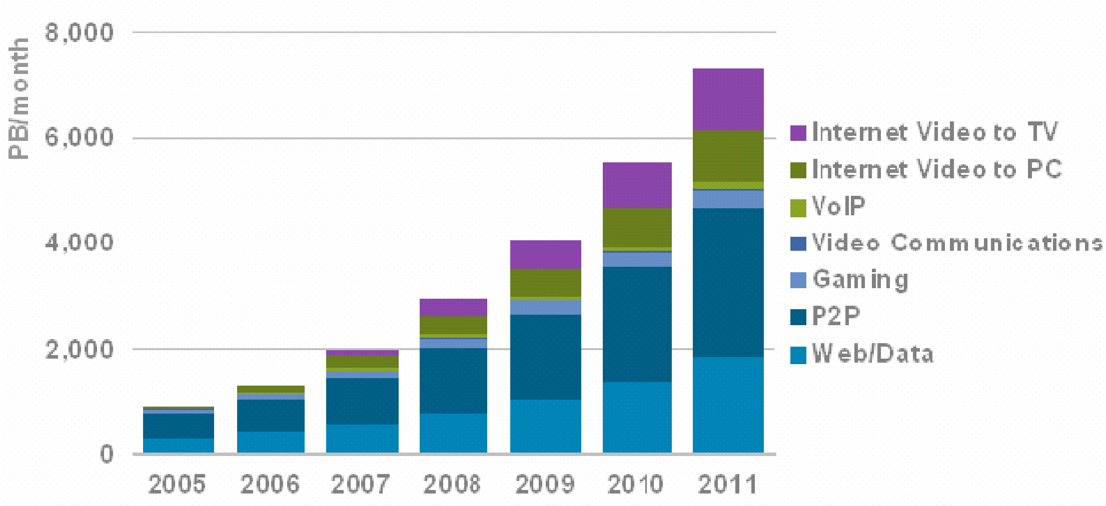

Source: Cisco Systems, "The Exabyte Era", 2007

2 Bandwidth Allocation Overview

The GPON network includes a mix of services. Some services, such as VoIP or native TDM, require constant upstream bandwidth, and the OLT may statically allocate the bandwidth for these services.

Other IP-based services, such as Internet browsing, streaming video, file sharing and file download, are bursty by nature. For highest upstream bandwidth utilization, the OLT should allocate the upstream bandwidth for these services dynamically, using a dynamic bandwidth allocation algorithm (DBA). With a good DBA algorithm, the GPON network upstream channel can be oversubscribed, thus increasing the number of ONTs that can connect to the network. A simple example is a network with 32 subscribers, where each may receive up to 100Mbit/s. The required capacity for this network is 3.2Gbit/s, nearly 3 times more than GPON upstream network capacity, but with a good DBA, these data rates can be supported and the service provider can charge for the full bandwidth service.

With the transition to a full IP traffic model, applications are expected to become even more bursty. As a result, the efficiency of the DBA algorithm becomes even more important. Efficiency directly impacts latency. For example, a burst of requests arriving from all ONTs accumulating to 3.75Gbit/s over 10mS would be cleared after 30mS with 100% efficiency. In contrast, if the efficiency is only 50%, it would take 60mS to clear. Not having an efficient DBA means that all of the active ONTs notice an additional 30mS delay in this case.

Latency is an important parameter, where the maximum limit plays a more significant role than the average value. Average latency of 1.5mS and a maximum value of 5mS provide a different quality of experience than an average latency of 1.5mS and a maximal latency of 50mS. Latency is mostly negatively impacted, i.e., increased, by the ability of the DBA to quickly adjust to the varying traffic.

In the GPON network, the OLT informs the ONTs of the upstream bandwidth allocation by transmitting Bandwidth Mapping messages (BWMAP), which are built of multiple bandwidth allocations for the individual ONTs or the ONT Transmission Containers (T-CONT). Each bandwidth allocation is an indication to an ONT to transmit in a defined time slot. The essence of DBA is dynamically calculating the BWMAP to allocate the right bandwidth for each ONT.

3 Static Bandwidth Allocation

Earlier generations of PON networks allocate their upstream bandwidth using static, TDM-like allocation. Each ONT gets its predefined bandwidth allocation, whether it uses it or not. This is ideal if all services in the network require constant allocation (VoIP or TDM), or when the provisioned upstream bandwidth is low enough, but it has low efficiency in the high-speed, high-utilization GPON network. As long as the ONT traffic keeps arriving at a fixed rate, upstream utilization is good. Once the ONT goes idle, as shown for ONT B and ONT C in Figure 3 , its statically allocated bandwidth is unavailable to other ONTs in the network, and the overall upstream utilization degrades. The latency of ONT B is higher than it would have been had the data been transmitted in the available slots.

This inability to use bandwidth prevents the carrier from earning revenue from the unused bandwidth. As long as the network is not congested, and the total upstream bandwidth required to support all ONTs at any given time is less than 1.244Gbps, the upstream channel available bandwidth is sufficient to service all ONTs with virtually no queuing.

If the unused bandwidth could have been allocated to other ONTs, it could have served the bursty services they run, improved the overall user experience and service level, and lowered the risk of queue congestion at these ONTs (refer to ***Bandwidth Allocation and ONT Upstream Queuing*** section).

In order to utilize the unused bandwidth to offer higher speed connections and better upstream QoS to residential and business users, the bandwidth allocation mechanism must be dynamic.

Figure 3 Static Bandwidth Allocation Example

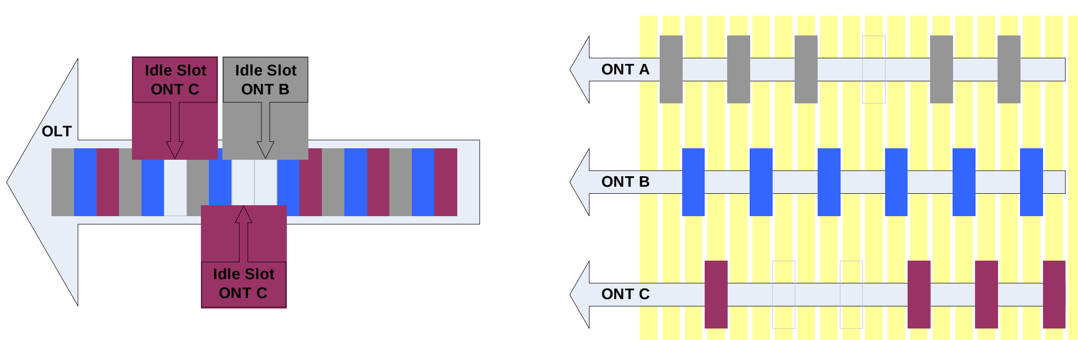

4 Dynamic Bandwidth Allocation

The GPON standard provides the tools to implement Dynamic Bandwidth Allocation (DBA) and leaves the actual bandwidth allocation scheme open to different implementations. Using these tools, the OLT can allocate bandwidth either per-ONT or per-ONT-per-service (T-CONT) and can base the bandwidth allocation on ONT requests, on measuring upstream traffic, or on any combination of the two, taking into consideration the Service Level Agreement (SLA) of the subscriber.

As an example, Figure 4 describes a scenario where bandwidth not used by ONTs A and C is allocated to other ONTs that request it.

Figure 4 Dynamic Bandwidth Allocation Example

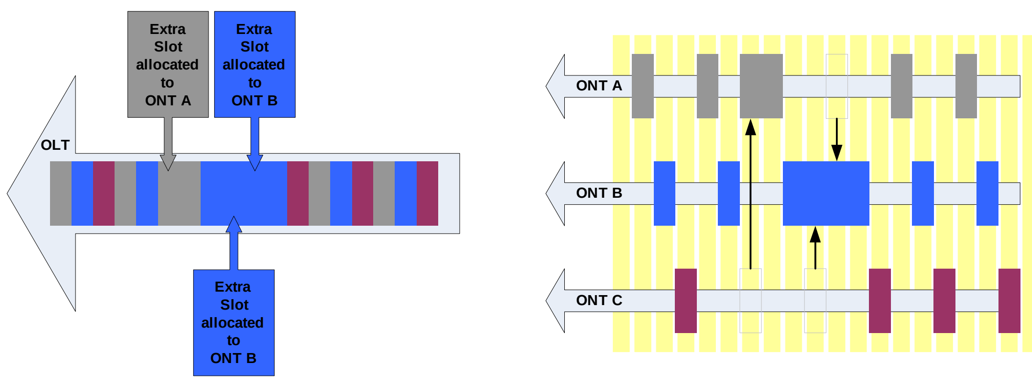

A good DBA algorithm quickly adjusts the upstream bandwidth allocation to the ever-changing traffic scenarios. At any given time slot, sufficient bandwidth would be allocated to certain ONTs (or certain services). In other time slots, the bandwidth would be allocated to other ONTs or other services, maximizing utilization of the PON bandwidth.

5 Oversubscription and Service Level Agreement (SLA)

Oversubscription, where the served bandwidth is theoretically higher than the physical capacity, is a key factor in the profitability of access networks. Typically, in other broadband access networks (cable/DSL), an operator oversubscribes between 4:1 and 20:1, and the oversubscription ratio depends on the subscribers' usage profiles. To enable oversubscription, the network must provide Quality of Service (QoS) to its users.

The Service Level Agreement (SLA) is a service contract between the provider and the subscriber. In this contract, the subscriber is charged for the services and bandwidth provided, regardless of actual usage. Usually, the SLAs consist of a *committed information rate* (*CIR*), which is committed to the subscriber, and an *excess information rate (EIR)*, which is additional bandwidth a subscriber may use, if available.

For example, a user may subscribe to a *CIR* of 40 Mbps with an *EIR* of 150 Mbps. If 32 ONTs need to share the same GPON channel, all with the same SLA, the theoretical required upstream bandwidth is 1.28 Gbps (1280 / 32 = 40 Mbps), which is slightly more than the available 1.244 Gbps upstream bandwidth.

If the bandwidth allocation scheme in the network is static, each ONT would get exactly its *CIR* at any given time, whether utilized or not. It would be impossible to serve an ONT with higher bandwidth even for short periods of time; therefore the E*IR* becomes irrelevant. The Static bandwidth allocation, which is QoS-unaware, does not allow for oversubscription, which translates to higher costs per subscriber.

When a DBA algorithm is used, the bandwidth allocation to an ONT can peak at higher rates, up to the *EIR*. Bandwidth that is not consumed by some ONTs is allocated to others, resulting in faster network behavior and a better customer experience. This leads to additional revenue to the operator, as statistically there is always excess bandwidth that the operator can charge for.

The DBA algorithm should implement *statistical multiplexing*, which leverages the bursty nature of the data services, considering that typically all applications of all users do not peak at the same time. Statistical multiplexing allows oversubscription, where the theoretical upstream bandwidth required is higher than the channel capacity. In this case, the sum of the maximum rates of all subscribers is higher than the physical PON capacity, and the sum of the committed rates is below upstream capacity. The DBA can provide the required QoS control, enabling oversubscription.

It should be noted that the PON downstream channel also relies on statistical multiplexing to provide services, and the upstream channel should not be different.

6 SR vs. NSR DBA

DBA algorithms can be divided into two categories:

- **Status Reporting (SR)**
  All ONTs report their upstream data queue occupancy, to be used by the OLT calculation process. Each ONT may have several Transmission containers (T-CONTs), each with its own traffic class. By combining the queue occupancy information and the provisioned Service Level Agreement (SLA) of each T-CONT, the OLT can optimize the upstream bandwidth allocation.
- **Non Status Reporting (NSR)**
  ONTs do not provide explicit queue occupancy information. Instead, the OLT estimates the ONT queue status, typically based on the actual transmission in the previous cycle. For example, if an ONT has no traffic to send, it transmits idle frames during its allocation period. The OLT would then observe the idle frames and decrease the bandwidth allocation to that ONT in the following cycle. In the opposite case, the OLT constantly increases the allocation size until idles are detected, slowly adjusting to growing traffic. There is no optimal point for the trade-off between slow response to traffic growth and utilization. Improving one aspect always comes at the expense of the other. Regardless of whether allocating more or less bandwidth, the latency is increased and degraded.

SR-based algorithms are superior to NSR-based algorithms in all respects. SR-based algorithm utilization is higher because the OLT does not overestimate or underestimate the queue occupancy, as NSR-based algorithms always do. SR-based algorithm latency is lower because traffic can be granted for transmission once reported by the ONT, while NSR-based algorithms typically increase the grant length gradually, requiring more grants until all pending data is transmitted. As a result, SR-DBAs grant the ONT report faster.

For all of these reasons, an efficient DBA algorithm should be SR-based, while the GPON network allows implementation of NSR-based algorithms. As typical oversubscription in other broadband access networks is between 4 and 20 to 1, only the SR-DBA enables oversubscription and high upstream efficiency.

7 DBA Cycle-based Allocation

Most DBA algorithms use the *cycle* concept. A DBA cycle is the duration covering transmissions from all ONTs. In some cases, several ONTs would not transmit during certain cycles, or cycle duration may vary, but the general concept of a time-wheel used for DBA processing appears in all algorithms. As schematically shown in Figure 5 the OLT collects DBRu messages from the ONTs in every given cycle, processes them in the following cycle, and uses the processing result to map the following cycle. The cycle length has a major impact on IP service performance as it determines the ability to change the bandwidth allocation maps, adjusting the allocation to the changing traffic.

Figure 5 Schematic Upstream Mapping Process

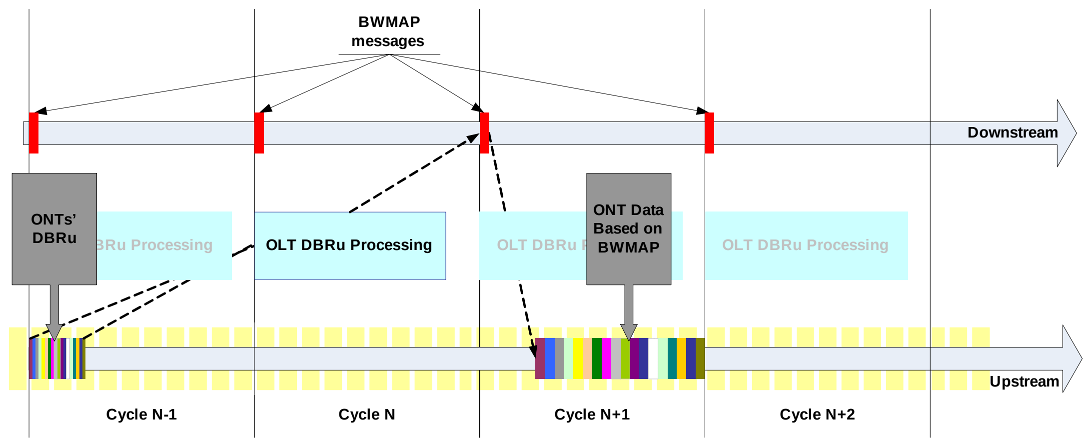

8 DBA Evaluation Parameters

There are three major parameters for evaluating DBA performance:

- Latency, or the time a packet waits at an ONT upstream queue before transmission
- Fairness, or the ability to equally satisfy / dissatisfy Service Level Agreement (SLA) for all ONTs
- Utilization, or the percentage of usable bandwidth in the upstream channel

The three criteria are interrelated. As described below, reduced utilization results in increased latency because the available bandwidth to empty an ONT queue is decreased; hence more time is required to empty the queue. Similar impact is noticeable for fairness. Low fairness performance indicates that some ONTs will be served slower than others. Consequently, the latency of the slowly-served ONTs increases.

In a sense, latency is the most important parameter, and the best indicator of the quality of the DBA algorithm. Utilization and fairness are optimized in order to improve latency. Low utilization and a poor fairness DBA do not mean that low latency cannot be reached, especially in synthetic test cases. However, it strongly indicates that under real-life scenarios, where traffic tends to be bursty, the latency is likely to increase for some or all of the ONTs.

8.1 Latency

Fundamentally, low latency is essential to fast downstream operation of the TCP protocol. Delaying an upstream TCP ACK message would limit the effective download rate. From the operator's point of view, increasing the download rate directly increases customer satisfaction.

Figure 6 Impact of RTT on TCP throughput

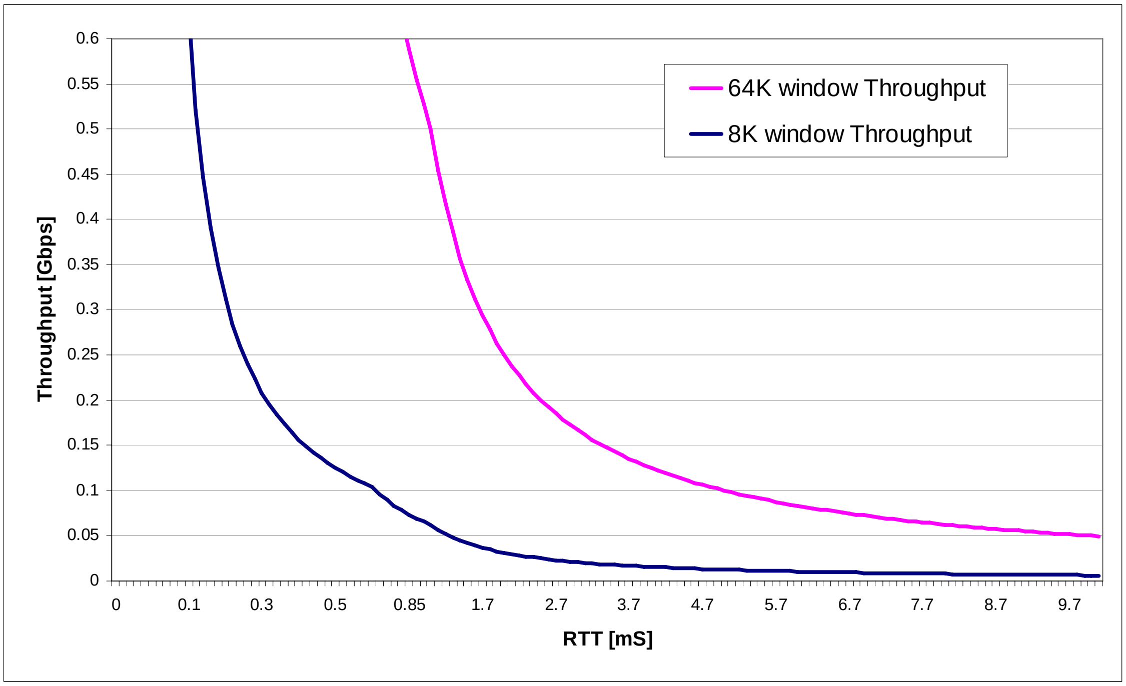

The graph in Figure 6 illustrates the impact of Round Trip Time (RTT) on TCP throughput. Two windows have been used for illustration: 64Kbytes and 8Kbytes. RTT includes the total time from the PC to the server and back, including the processing time at each end, with the PON network being just one element of the delay. In a network optimized for low latency, the PON portion of the delay must not be the major contributor. For example, in order to reach 100Mbit/s throughput, the total RTT should be 5mS. Assuming the delay external to the PON is 3mS, 2mS are allocated to the PON delay for reaching the target. A PON delay of 7mS, for example, will limit the throughput to 50Mbps, despite the network capability to provide twice that throughput. When the access network becomes the bottleneck, the operator will be forced to spend and enormous amount of effort on limiting the overall network delay to 3mS, effort and expens that could have been saved. In such cases, this huge effort can not be leveraged, as the PON delay will overwhelm the total network delay.

The latency impact on performance is discussed in the following section

8.2 Bandwidth Allocation Impact on Upstream Latency

With the shift to more symmetrical networks, it is critical to correctly design the ONT upstream queue depth. It is important to find the balance between sufficient queuing that prevents packet loss, but does not add latency that impacts TCP performance (latency is discussed above in the ***Latency Consideration*** section). The Bandwidth Allocation algorithm has a major impact on minimizing latency and on the required buffer size.

In the following simulation of ONT upstream traffic, we assume an upstream SLA of 50Mbps. In this example, the subscriber generates P2P, video upload and picture upload applications. These applications consume the entire bandwidth with some burstiness, but they average 50Mbps. For the purpose of the simulation, the upstream queue of the ONT is infinite, thus no packet loss occurs and the upstream latency accumulates. Three upstream bandwidth allocation methods by the OLT are simulated: static, dynamic NSR and dynamic SR. The network is assumed to be congested, so no 'spare' bandwidth is available.

Figure 7 describes the correlation between the bursty ONT upstream traffic and the allocation. The incoming upstream traffic through the ONT UNI port averages 50 Mbps. When static bandwidth allocation is used, the OLT constantly allocates 50Mbps, regardless of the incoming traffic, and does not track ONT needs. When NSR-based allocation is used, the OLT is able to slowly respond, with a time offset, to the incoming ONT traffic, but having no visibility to the ONT queue occupancy, the OLT cannot allocate bandwidth in a way that will keep the ONT upstream queue relatively empty and the latency low. As shown in Figure 8 , upstream packets that build up in the ONT upstream queue are transmitted with increasing latency. In all practical scenarios, the ONT upstream queue would overflow. Packets would be dropped and re-transmitted with no clear difference between the NSR DBA and static allocation.

Figure 7 Bandwidth Allocation with Different Algorithms

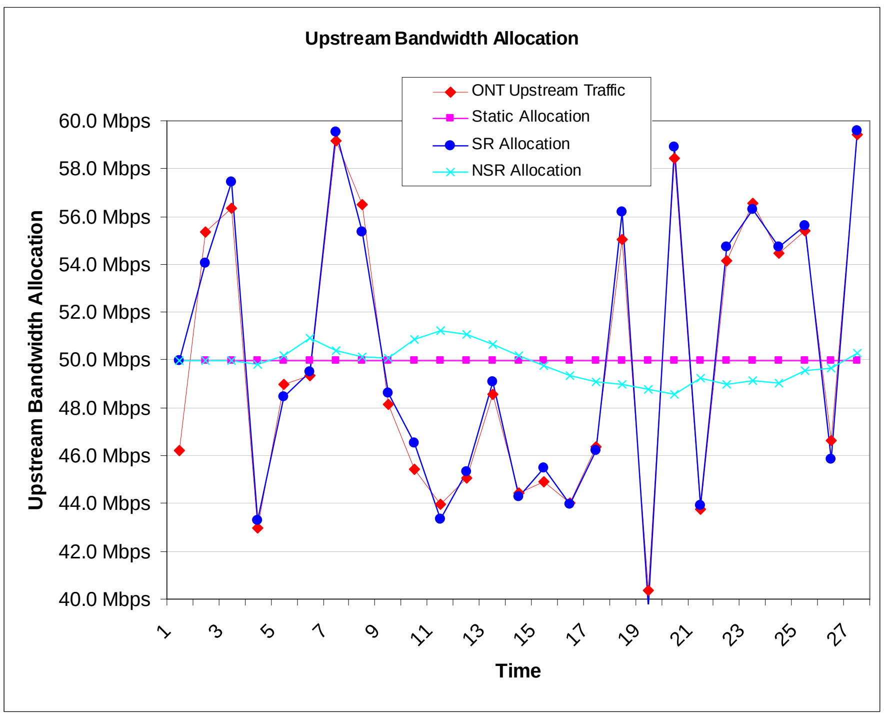

When implementing SR-based allocation, the OLT can measure the incoming traffic from the ONT and has full visibility into the upstream queue occupancy of the ONT. The SR-based allocation tracks ONT needs well and empties the ONT upstream queue in a timely manner. As seen in Figure 8 and Figure 9 , the upstream latency is kept low and relatively constant, and there is no risk of queue overflow and packet loss.

Figure 8 Average Upstream Latency with Different Bandwidth Allocation Algorithms

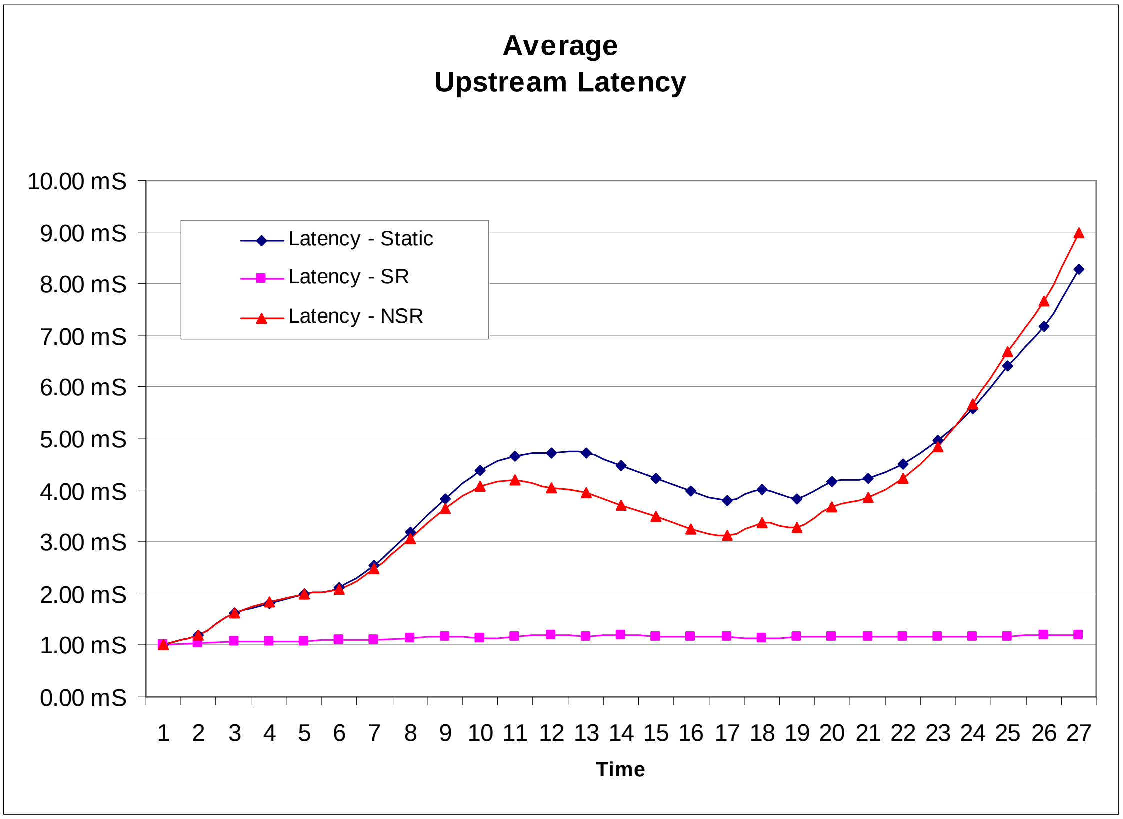

Figure 9 Average Upstream Latency with SR DBA

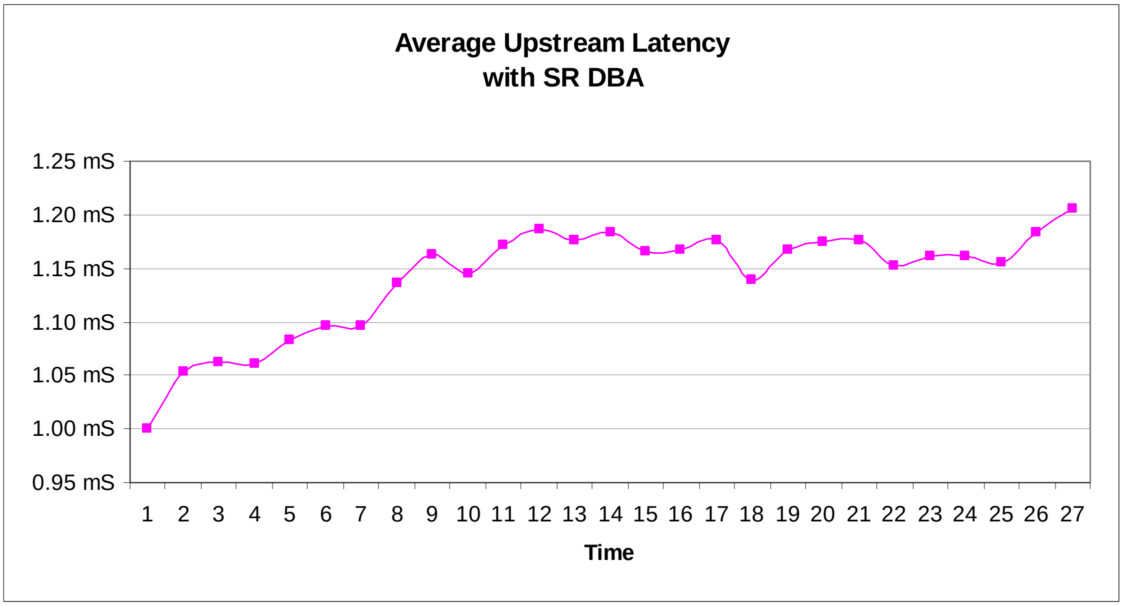

8.3 Fairness

Fairness is an elusive parameter. A fair algorithm would guarantee that SLA deprivations are equally shared among all ONTs. As an example, two ONTs with identical SLA receiving 60% of their provisioned SLA is considered to be fair, compared to one ONT receiving 100% of the SLA and the other ONT only 20%.

Fairness can only be guaranteed by considering queue status information from all ONTs. The less ONTs considered, the harder it is to provide fairness. Since queue information is unknown in NSR-mode or with static allocation, fairness is unreachable.

From the operator point of view, higher fairness indicates tighter fulfillment of the SLA, increasing customer satisfaction.

9 PMC-Sierra's GPON OLT

The PAS5211, PMC-Sierra's four-port GPON OLT SoC, provides the core functionality of an ITU-T G.984 GPON-compliant OLT solution. The PAS5211 integrates four GPON channels, consisting of a GPON Media Access Control (MAC) and protocol management function, a GPON compliant serializer / deserializer (SERDES) and burst-mode clock-data-recovery (CDR) function, an advanced packet classification engine, on-chip packet buffers and a selection between four serial 2.5 Gbit/s interfaces, 1 per each channel, and a single 10 Gbit/s XAUI interface to the core network.

The PAS5211 works in combination with PMC-Sierra's family of Optical Network Terminal (ONT) solutions, and also accommodates other GPON-compliant ONT solutions.

The DBA implementation is a combined hardware-software solution. It is based on various DBA software algorithms that use hardware accelerators to provide programmable, QoS-capable DBA solutions. These solutions support any combination of Status Reporting and Non Status Reporting ONTs.

As the PAS5211 offers the complete functionality required in a GPON OLT device, it includes an embedded processor sub-system for real-time PON management and programmable Dynamic Bandwidth Allocation (DBA) algorithms. Alternatively, DBA development platform can be provided to users, for the ability to adapt the programmable DBA engine to specific QoS requirements.

10 Summary

Carriers can increase revenues by providing oversubscription, and can improve customer satisfaction by providing low latency PON. The quality of the DBA used is proportionately translated to increased carrier revenues.

The only way to achieve both oversubscription and low latency PON is by implementing a DBA algorithm. A DBA algorithm minimizes latency, improves utilization and should respond quickly to the changing traffic patterns. The DBA algorithm should be SR-based, to make the best gains possible.

When implementing SR-DBA in the GPON upstream channel, the operator can oversubscribe and provide 5 Gbps and beyond of allocated upstream bandwidth by leveraging the traffic burstiness. Through statistic multiplexing and ONT reports, latency is reduced by as much as 90%.

Contacting PMC-Sierra

PMC-Sierra
100-2700 Production Way
Burnaby, BC
Canada V5A 4X1

Tel: +1 (604) 415-6000
Fax: +1 (604) 415-6200

Document Information: document@pmc-sierra.com
Corporate Information: info@pmc-sierra.com
Technical Support: apps@pmc-sierra.com
Web Site: http://www.pmc-sierra.com