

QoS Control Schemes For Optical Ethernet Over SONET Transport

Aarti Ahuja

U.I.E.T, Panjab University,
Email: ahuja.aarti@gmail.com

Ravi Tandon

U.I.E.T, Panjab University
Email: robinTandon_007@yahoo.co.in

Abstract - Ethernet passive optical networks (EPONs) have emerged as the one of the most promising candidates for next-generation access networks. This paper introduces a new optical transport. These new architectures couple low-cost optics with advanced edge electronics to offer vastly improved scalability over competing digital subscriber line and cable modem offerings. This paper proposes several novel architectural enhancements for EPON, which will help increase the viability of optical access over a broader range of subscriber access scenarios. Specifically, this paper proposes two-stage EPON architecture that allows more end-users to share an optical line terminal link, and enables longer access reach/distances (beyond the usual 25 km distance). More importantly, it is competitive with SONET protection without its 100% bandwidth overhead.

Keywords : Ethernet, metropolitan area network (MAN), network, optical services, access network, dynamic bandwidth allocation algorithm (DBA), Ethernet-based passive optical network (EPON), quality-of-service (QoS), PESO (Protection for Ethernet over SONET transport) protocol in EPON.

I. INTRODUCTION

Recently, there has been a dramatic increase in data traffic, driven primarily by the explosive growth of the Internet, as well as the proliferation of corporate virtual private networks (VPNs) [1]. As traffic demands have grown, many carriers have been prompted to add capacity quickly and in the most cost-effective way possible. As result, new core optical networks have been extensively deployed, and in particular, the use of dense wavelength-division-multiplexing (DWDM) technology has dramatically increased the capacity of these networks [2], [4]. At the same time, enterprise local-area networks (LANs) technologies have steadily scaled tributary speeds progressively from 10 and 100 Mb/s upwards toward multi gigabit speeds, e.g., 1.0, 10 Gb/s Ethernet. Overall, the above developments have led to a growing “access bottleneck,” where metro/regional and backbone capacities are vastly out-scaling last-mile bandwidths. For example, many renditions largely limit end-users to speeds under 10 Mb/s and distances under 5 km. Hence, these technologies lack broader universality for business (or small business) settings. Clearly, new and improved access solution technologies are required.

These offerings must be inexpensive, yet still be capable of scaling to delivering bundled data, voice and video over the same high-speed connection. Additionally, other prime concerns are quality-of-service (QoS) guarantee provisions and the ability to purchase bandwidth on an as needed basis [6]. It is here that Ethernet passive optical networks (EPONs) [6] have emerged as the best candidate for next-generation access networks. Propelled by rapid price declines in fiber optics and Ethernet components [3], [5], these new EPON architectures combine the latest in optical and electronic advances and are poised to become the dominant means of delivering bundled services over a single platform [3], [5].

EPON is basically a point-to-multipoint (1:N) optical access network with no active elements in the signal path. The network provides two-way operation (Fig. 1), in which traffic from an optical line terminal (OLT) is sent to/from multiple optical network units (ONUs). Namely, OLT-ONU traffic is called “downstream” (point-to-multipoint) and meanwhile, reverses ONU-OLT direction traffic is called “upstream” (multipoint-to-point) [3]. The latter requires contention resolution (arbitration) mechanisms to avoid upstream transmission collision between ONU senders.

The OLT typically resides in a central office (CO) location and connects the optical access network to the metro (backbone) network. Meanwhile, the ONU is usually located at or near end-user locations and must support a wide array of services—broadband video, voice, data, etc. In particular, various ONU deployment possibilities exist, as per different architectures such as fiber-to-the-curb (FTTC), fiber-to-the-building (FTTB), and fiber-to-the-home (FTTH) [6]. Overall, the operational costs of these setups are minimal since no active elements are placed in the outside fiber plant, e.g., no maintenance is needed in the field. Moreover, by sharing the network equipment among the maximum number of customers, operators can amortize the cost of installation and operation in a much more economical manner.

Along with EPON we also know that Ethernet over SONET (EoS) is increasingly being deployed as the foundation for next-generation data services in service provider networks. Both Ethernet and SONET are the dominant transport technologies — Ethernet for data transport in local area networks (LANs) and SONET for reliable voice transport in metropolitan and wide-area networks (MANs and WANs). EoS and EPON also make good business sense – providers

leverage their SONET infrastructure to deliver new services and thereby, create new revenue streams from their legacy hardware. EoS and EPON are being also driven today by the availability of multi-service switches that can support both (Gigabit) Ethernet i.e. EPON and SONET. In addition to traditional SONET rings, these switches also support more efficient mesh topologies. As providers deploy next-generation SONET, it is expected that mesh architectures will increasingly become commonplace. Riding the SONET infrastructure over MANs and WANs, it is thus possible to deliver Ethernet data services seamlessly over regional and national geographic areas. This includes Ethernet private-line services providing dedicated bandwidth and virtual private-line services that use statistical multiplexing to share bandwidth among various streams. Applications for these services include voice and other enterprise applications such as storage networks and Transparent LANs

Ethernet and SONET use very different mechanisms for data protection due to their packet (or, frame) and circuit switched nature. Ethernet uses (Rapid) Spanning Tree Protocol [7],[8] as a protection mechanism. Since Ethernet switches (or bridges) use a spanning tree to forward frames, these protocols dictate how to reconfigure the tree as quickly as possible after a failure. In practice, depending on the size of the network, the reconfiguration can take 10-60 secs during which time there may be traffic disruption. SONET, on the other hand, uses some variant of the 1+1 Automatic Protection Switch (APS) such as UPSR and BLSR, where primary and backup paths are pre provisioned. On failure, the system switches from one path to the other [9]. SONET APS typically takes about 50 ms and is considered the gold standard of reliability. However, the pre-provisioning of two paths imposes at least a 100% protection bandwidth overhead.

II. PROTOCOLS AND ARCHITECTURE DESIGN

In general, an EPON network cannot be treated as a basic shared medium network, i.e., one using carrier sense multiple access with collision detection (CSMA/CD) type protocols. At the same time, neither can an EPON network be treated as a point-to-point network. Instead, it is a combination of both types. For example, consider upstream ONU-OLT communications.

Here, due to large propagation delays across EPON infrastructures (which can easily exceed 20 km), the effectiveness of regular CSMA/CD protocols is greatly reduced. Instead, these protocols are more suited for SONET domains, where links are short and traffic predominantly comprises “best-effort” data. Since EPON access [10] must support much more stringent service requirements (QoS), related architectures must provide strict guarantees on such as packet delay and jitter performance. In order to accommodate diverse traffic types and to achieve bandwidth sharing and flow isolation, clearly efficient bandwidth allocation algorithms must be developed. This approach allows the ONUs to share a single upstream wavelength in which the OLT allocates timeslots to each ONU to transmit its

backlogged traffic. Overall, this yields a very cost-effective solution, and Fig. 1 illustrates time-shared data flow in an EPON. Time-sharing techniques can either be static or dynamic. In the former, each ONU is allocated a fixed timeslot to transmit data. Although this is a rather simple approach, its implementation is not contingent with EPON’s requirements, e.g., QoS support, OLT link efficiency in SONET. Therefore, more effective and dynamic schemes will be needed in order to successfully implement service guarantees in the next generation access networks.

Along these lines, [11] presents a simple algorithm for dynamic bandwidth allocation based upon a time interleaving method, i.e., the interleaved polling scheme with an adaptive cycle time (IPACT) scheme. Here, an in-band signaling approach is used to allow the use of a single wavelength for both downstream data and grant transmission. Here, notable examples include the limited, gated, linear credit, and elastic allocation schemes.

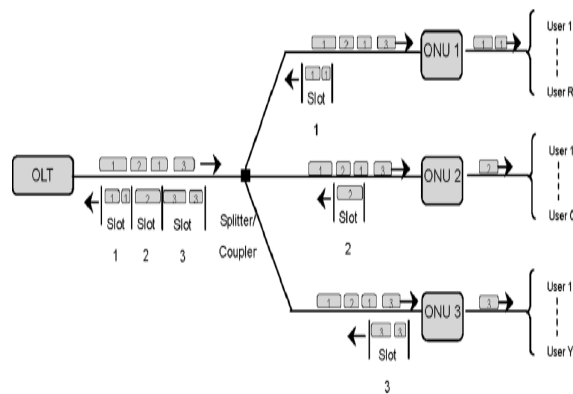


Fig 1. EPON network architecture

A. Dynamic Bandwidth Allocation with QoS Support for Ip-Epons:

In the upstream direction an EPON network acts like a shared medium setting in which all ONU devices can potentially contend while transmitting data. Ideally, at any given time only one ONU should be allowed to occupy the medium. It is also important to state that there is no direct communication between

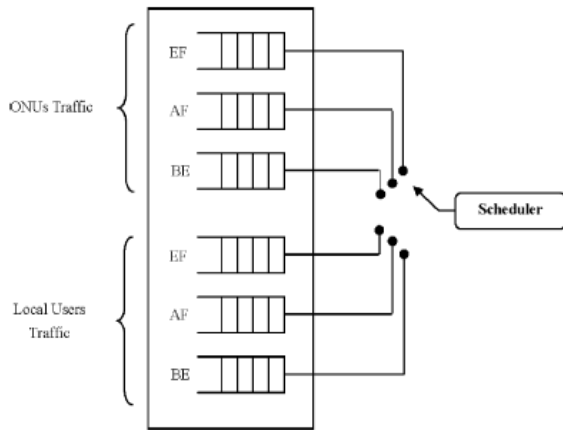


Fig 2.. Sub-OLT queue structure.

B. ONUs in an EPON-based network.:

As such, the OLT is the only network element that is able to communicate with the other units in an EPON setup. Note that in the basic EPON standard, the new multipoint control protocol (MPCP) is used to implement bandwidth arbitration by the OLT [12]. At the beginning of each transmission cycle, the OLT sends a GATE message to each ONU through the downstream connection. This GATE message contains the following information: time when the ONU should start transmission and the length of its transmission window. Upon receiving its GATE message, the ONU performs synchronization and updates its local parameters. When its local clock matches the transmission start time sent by the OLT, the ONU starts sending information packets to the OLT. At the end of its transmission window, the ONU sends a REPORT message to the OLT to report its current buffer occupancy and request bandwidth for the next transmission cycle.

Upon receiving REPORT message from the ONU, the OLT updates its report table and passes the message to the DBA module responsible for bandwidth allocation decision. At the end of the transmission cycle and upon receiving REPORT messages from all ONUs in the network, the DBA module in the OLT calculates the new window size for every ONU in the network.

Thereafter, a series of GATE messages are generated and broadcast continuously to the ONUs by the OLT to initiate the next transmission cycle. The MPCP only defines the mechanism for the OLT to arbitrate the transmissions of its attached

ONUs and does not specify any details about the bandwidth allocation amongst ONUs. Thus, the development of efficient EPON scheduling algorithms (i.e., dynamic bandwidth allocation schemes) that can accommodate multiple, diverse traffic types, e.g., voice, video, and data are critical for the deployment of such networks.

C. LCAS Protocol:

In this section, we explain a key shortcoming of the VC protocol and describe how LCAS address it. We also briefly describe some operational details of LCAS relevant to this work. When one or more members of a VCG are adversely affected due to a network or link failure, the data can be corrupted even if some of the VCG members are still active. Consider an example where a packet stream of "123456781234..." was byte interleaved on to four members A, B, C and D. Now, consider a network element failure resulting in the failure of member D. Since, source node is not aware of the failure of member D; it'll keep mapping the data to all four members. However, failure of member D will force the sink node to perform packet assembly only with the three active members. Hence, the absence of member D will make the reconstructed data look like "123567123..." result in a malformed packet. Therefore, even though, VC theoretically provides resiliency by routing the members diversely, no practical benefits can be achieved by it. LCAS protocol as described by the [10] remedies this problem. It provides a mechanism for the sink to notify the failure of a VCG member to the source using the still active members. After receiving such notification, the source node temporarily removes the failed member from the VCG group and starts sending data only on remaining active members. LCAS also enables scheduled addition and deletion of traffic from a VCG in a hitless manner. It can also be used to detect the restoration of a failed member and add it back to the VCG without requiring any operator intervention.

D. PESO (Protection for Ethernet over SONET transport) protocol in EPON:

PESO consists of three key components a) routing) failure notification and c) protection switching. Depending on the reliability requirement being a ROP or, a NOP, each of these components function slightly differently. The PESO protection scheme can be designed to handle any specific failure model, however, assume a single link or node failure.

Consider a NOP scenario (e.g., Scenario A and B). In these cases, PESO routing determines the number of VCG members necessary to provide the appropriate reliability and suggest routes for them. On failure, LCAS resizes the bandwidth in LRT time. Even though no additional bandwidth is over provisioned, a weaker form of reliability is achieved by enabling data flow even after network failure. This is in contrast to a traditional all-or-nothing protection where the circuit would go down without any additional bandwidth over-provisioned. The ROP scenario (e.g., Scenario C) is similar to the more traditional reliability requirements — the operator is willing to overprovision in order to be continue at full throttle even after a failure. The PESO approach in this case is to pre provision additional bandwidth as "backup" members in the VCG in addition to the "primary" members that would normally carry traffic. The PESO routing component provides the necessary routes for all the members. On failure, LCAS is used to switch traffic from the primary to the back up members. Assuming the backup bandwidth suffices, after a

disruption of the LRT time, the circuit is restored. The effectiveness of PESO, in terms of protection bandwidth overhead, is dependent upon the availability of diverse routes in the network. In the traditional SONET network where most of the deployment is in UPSR/BLSR rings, the network is limited to two diverse routes. As a result, PESO also will require a 100% protection bandwidth. However, PESO scheme will be extremely effective in mesh architectures, where availability of diverse paths is high. As service providers move to next-generation SONET networks, they are migrating towards mesh due to the efficiencies it provides over rings. we propose a novel routing scheme which minimizes the protection bandwidth overhead requirement. PESO's protection speed is dependent upon the failure notification mechanism used by sink to notify source node about the failure of a primary member. In the current LCAS standard [10], time taken by the sink to detect and report failure is of the order of 64ms/128ms for Higher Order and Lower Order concatenation respectively. PESO proposes a faster version of protocol, FLCAS, presented, which substantially brings down failure notification time, to less than 50ms, for most of the cases.

E. Protection Switching

In this section, we describe how PESO recovers from a network or link failure. As described above, in the NOP case, PESO simply uses the LCAS protocol for the member failure detection and their removal. The novelty in the switching component of PESO is for ROP scenarios such as the Scenario C i. For such cases, VCG members are partitioned into primary and backup members. Once the primary and ackup members have been identified and routed, the source node starts sending traffic




Fig3.PESO before Failure

on the primary members. The backup members do not carry any traffic during normal operation. The primary members carry NORM in their CTRL field while backup members carry DNU to ensure that the sink does not pickup any data from them. When a link or network element failure results in failure of a primary member, the LCAS protocol at sink detects and reports the failed member's status FAIL back to the source. The PESO protection switching kicks in after the source node receives the notification of a member failure. It ensures that the failed primary members are temporarily removed from the

VCG and instructs already provisioned backup members to take over. Upon notification of a member failure, PESO redirects the source to start sending normal (NORM) in the CTRL field of backup member and DNU on the failed member. This swapping of CTRL fields is achieved in same multi frame header. Once this multi frame header information is completely transmitted, the source switches the data previously transmitted on failed primary members to backup members. Since, primary members can share routes; a single failure (network or link) can affect several primary members. Again, it is the responsibility of the routing algorithm to ensure that sufficient number of backup members is setup to support any failure. Figure 3 and 4 shows an example of the PESO protection scheme. To transport a 600M Ethernet service, a 4 member VCG is setup. Each VCG member is a STS-3c circuit routed diversely from one another. The backup member (SQ 5) is also diversely routed from primary members (SQ1-4). The protection bandwidth required to protect this VCG is 25%, an extremely low bandwidth overhead. As shown in Figure 3, backup members do not carry traffic during normal operation and instead have DNU in their CTRL field. Figure 4 shows the swapping of the CTRL field of primary and backup members after failure. Clearly, PESO requires some support from the network element to enable this protection switch mechanism. The element has to provide some means to mark members as primary and backup and have the necessary logic to do the switch on failure. However, this is a relatively minor requirement compared to the complexity of supporting VC and LCAS. One may also conceive of a variation of PESO scheme for the ROP case. Instead of partitioning into primary and backup members, one can spread the traffic equally among all the members. On failure of a member (active or backup), it is simply removed from the VCG using LCAS. The advantage of this scheme is that there is no special processing required once the source is notified — it simply follows the LCAS protocol on failure. However, it has a fundamental drawback.




Fig.4 PESO after failure

In PESO, traffic is not impacted if a backup member goes down since it does not carry any valid data (has a DNU flag).However, in this case, any member going down has a LRT time hit in traffic. Thus, PESO provides an extra level of reliability which plays a critical role since the diversity of VCG members increase the probability of introducing failures. Finally, we briefly focus on the protection switch time .For both the NOP and the ROP case, the failure detection

and notification is via LCAS requiring the LRT of 64ms. Swapping the backup and primary members takes another 2ms for the ROP case and the same for the removal of failed members in NOP case when done by LCAS. So, total worst case switch time can be 66ms which is acceptable for most data applications. Note that we are ignoring signal propagation delay ($\approx 5\mu\text{s}/\text{km}$) as they are negligible compare to switching times and are fixed for all the protection schemes. For the services which have even more stringent requirements, It is also proposed that an enhancement to lower the switch time.

F. Peso Routing:

In this section, we propose a novel routing scheme to enable the routing for virtually concatenated circuits. As mentioned in previous section, VC provides a unique opportunity of splitting the traffic flow on multiple paths (VCG members) carrying smaller rate traffic. PESO routing algorithms are intended to exploit this flexibility. The PESO routing algorithm accounts for both the NOP and ROP scenarios. The routing to address the scenarios is considered below

G. Routing for Scenario A (Algorithm α)

This is very likely be a very common case as service providers may not be willing to put additional bandwidth to protect data services. However, they may be interested in limiting the extent of the damage on failures. Moreover, critical services tend to be provisioned at their peak rates and thus, a temporary failure may not necessarily impact the end user performance. Algorithm α shown on next page addresses this scenario. Consider the network in Figure 5, where the requirement is to transport a 120Mbps Ethernet Service from source S to sink D such that single failure does not impact more than $2/3^{\text{rd}}$, or, 67% of the traffic. Transporting a 120Mbps Ethernet service requires a STS-3c ($\approx 156\text{Mbps}$) equivalent frame rate on SONET side. As per the VC standard, it can be achieved by either one STS-3c circuit or three STS-1 circuits. For the moment, we assume this service is transported on a three member STS-1 VCG. We will highlight later the trade-offs involved in choosing between STS-3c or STS-1 as members. Since, the requirement is that at least 40Mbps traffic (33% of 120Mbps) is protected against one failure, it is necessary for at least one STS-1 member to survive any failure. Now consider the network of Figure 6, which represents the network of Figure 5 with link capacities altered. These new link capacities reflect the largest SONET rate (STS-Nc) they can carry. For instance, link S-A has available bandwidth of 100Mbps, which makes it large enough to carry only a STS-1 ($\approx 52\text{Mbps}$), hence a capacity of 1 unit. Thus, routing of a 120Mbps service in Figure 5 is equivalent to routing (or pushing) 3 units of flow in the network of Figure 6. However, to ensure that no link failure results in failure of more than two members (or two units of flow), its necessary that no link is allowed to carry more than two units of flow. To capture this constraint, we restrict the link capacities to a maximum two units. For example, though the link S-C has three units of capacity (as it can support a STS-3c), it has been assigned two

units. For routing F (or, 3 units in this case) units of flow, any of the standard flow routing algorithms can be used. For example, path augmentation based maximum flow algorithms from Ford




Fig 5: Original Network




Fig 6: Transformed Network

& Fulkerson [14] or, Edmonds & Karp [15] can be used to route the flow. As our requirement is to route only F units of flow, these algorithms can be stopped after the sufficient flow is routed. In any given network, there may be various distinct solutions for routing F units of flow; therefore it may desirable to obtain the smallest cost solution. Such smallest cost feasible flow solutions can be easily computed using Minimum Cost Flow algorithms [16]. In Section VI-D, we analyze the overall complexity of algorithm α based on one such min-cost flow algorithm. Once F units of flow is routed, F paths of unit flow are extracted and each path is used to route a VCG member made up of a STS-1 circuit. Note that the Algorithm α only handles link failures and not node failures. For example, in Figure 6, failure of the node E will result in complete failure of the entire VCG. However, we can address this by doing a standard graph transformation [13] of splitting each node into an ingress and egress node and inserting a link of requisite capacity between them. Therefore, for rest of the paper we only address link failures and assume that node failures can be accounted for using standard transformations.

H. PESO Routing Algorithm α

Input: Network $G(V, E)$, new demand D for bandwidth B and the maximum bandwidth X allowed to be impacted on failure.

Problem: Route D in G such that a single link failure does not affect more than X amount of the traffic.

Output: A set of routes for members of the VCG carrying D.

Algorithm 1

Let STS-Fc and STS-Yc be the smallest SONET frame rate that can carry B and X respectively. For all edges in E: Set

their capacity to highest SONET rate (N units for STS- N_c) they can carry or, to Y units whichever is smaller. Find minimum cost flow of F units in G .

I.. Routing for Scenario B (Algorithm β)

This problem is similar to Scenario A except that the requirement is to minimize the extent of damage on failure.

In a network of high route diversity, all the flows can be routed on disjoint paths where any failure will affect only unit flow. On the other extreme, in a network with no diversity where the all flow is carried on one route, a failure will bring the entire traffic down. Therefore, the problem of minimizing Smallest Y for which F units of flow can be routed in G , is the desired solution. The damage on failure requires finding a solution in between these two extremes. Algorithm β above achieves that. Once the required value of flow (F) is determined from the bandwidth (B), Algorithm β chooses a value of Y (damage on failure) by doing a binary search between 1 and F . For each value of

Y , it first alters the link capacities as in algorithm α and then attempts to route the flow of F units. For each value of Y , algorithm β finds a solution (if there exists one) where VCG circuit of bandwidth B can be routed such that no link failure will affect more than STS- Y_c (assuming STS-1 members) amount of bandwidth. The smallest value of Y for which F units of flow can be routed in G , is the best solution.

J. PESO Routing Algorithm β

Input: Network $G(V, E)$, a traffic demand D of bandwidth B .

Problem: Route the demand D in G such that a single link failure affects the minimum amount of traffic.

Output: A set of routes for members of the VCG carrying D .

Algorithm 2

Let STS- F_c be the smallest SONET frame rate that can carry B . Choose Y between 1 and F by binary search.

For all edges in E : Set their capacity to highest SONET rate (N units for STS- N_c) they can carry or, to Y units whichever is smaller. Find minimum cost flow of F units in G .

Smallest Y , for which F units of flow can be routed in G , is the desired solution.

CONCLUSION

It is concluded that EPON designs help in increasing the universality and making better the coverage within the access domains. DBA algorithms help in validating the architecture while EOS helps in reliable data transfer in PESO. It is been reviewed that the 50 ms protection is being possible. The study of protocols helps in splitting of the traffic and it is also possible to operate at lower bandwidth capacity for shorter period of time. PESO provides failure notification which is the best feature.

REFERENCES

[1] K. G.Koffman and A. M. Odlyzko, "Internet growth: Is there a "Moore's law" for data traffic?," in Handbook of Massive Data Sets. Norwell, MA: Kluwer, 2001.

[2] B. Mukherjee, "WDM optical communication networks: Progress and challenges," IEEE J. Sel. Areas Commun., vol. 18, no.10, pp. 1810–182, Oct. 2000.

[3] B. Lung, "PON architecture 'Future proofs' FTTH," Lightwave, vol. 16, pp. 104–7, Sep. 1999.

[4] N. Ghani, S. Dixit, and T.-S. Wang, "On IP-over-WDM integration," IEEE Commun. Mag., vol. 38, no. 3, pp. 72–84, Mar. 2000.

[5] Alloptic. "Ethernet Passive Optical Networks," Whitepaper. International Engineering Consortium (IEC). [Online]. Available: <http://www.iec.org>

[6] G. Kramer and G. Pesavento, "Ethernet passive optical network(EPON): Building a next-generation optical access network," IEEE Commun. Mag., pp. 66–73, Feb. 2002.

[7] I. Standard, "Spanning Tree Protocol," ANSI/IEEE Std 802.1D, 1998Edition.

[8] IEEE Standard, "Rapid Spanning Tree Algorithm and Protocol," ANSI/IEEE Std 802.1W, 2001.

[9] W. J. Goraliski, SONET/SDH. McGraw-Hill, 2002.

[10] C. Assi, Y. Ye, S. Dixit, and M. A. Ali, "Dynamic bandwidth allocation for quality-of-service over Ethernet PONs," IEEE J. Sel. Areas Commun., vol. 21, no. 9, pp. 1467–1477, Nov. 2003.

[11] G. Kramer and B. Mukherjee, "Ethernet PON (EPON): Design and analysis of an optical access network," Photonic Netw. Commun. J., vol. 3, pp. 307–319, Jul. 2001.

[12] G. Kramer, B. Mukherjee, and G. Pesavento, "Interleaved polling with adaptive cycle time (IPACT): A dynamic bandwidth distribution scheme in an optical access network," IEEE Commun. Mag., vol. 40, pp. 74–80, Feb. 2002.

[13] ITU-T Standard, "Link Capacity Adjustment Scheme for Virtual Concatenated Signals," ITU-T standard, 2001.

[14] J. L. R. Ford, "Flows in network," Princeton University Press, 1962.

[15] J. Edmonds and R. M. Karp, "Theoretical improvements in algorithmic efficiency for network flow problems," Journal of ACM, vol. 19, No. 2, 1990.

[16] T. L. M. R. K. Ahuja and J. B. Orlin, Network Flows: Theory, Algorithms, and Applications. Prentice Hall, 1993.

[17] J. B. Orlin, "A Faster Strongly Polynomial Minimum Cost Flow Algorithm," Proc. of the 20th ACM Symposium on the Theory of Computing, pp. 377–387, 1988.

[18] W. S. Jewell, "Optimal Flow through Networks," Interim Technical Report 8, Operations Research Center, MIT Cambridge, MA.

[19] S. Acharya, B. Gupta, P. Risbood and A. Srivastava, "PESO: Low Overhead Protection for Ethernet over SONET Transport," Bell Labs. Technical Report, 2003.

[20] J. Conover, "Networking for the next generation," Network Computing Magazine, <http://www.networkcomputing.com>, 2001.