Bandwidth management for WDM EPONs

Michael P. McGarry and Martin Reisslein

Department of Electrical Engineering, Arizona State University, Tempe, Arizona 85287
michael.mcgarry@asu.edu, reisslein@asu.edu

Martin Maier

EMT, Institut National de la Recherche Scientifique, Montréal, Canada
maier@ieee.org

Ahmet Keha

Department of Industrial Engineering, Arizona State University, Tempe, Arizona 85287
ahmet.keha@asu.edu

Received March 15, 2006; revised July 11, 2006; accepted July 15, 2006;
published August 18, 2006 (Doc. ID 68980)

We explore the problem of bandwidth management for the evolutionary upgrade of WDM EPONs. We divide the bandwidth management problem into two subproblems: (1) grant sizing and (2) grant scheduling. We then apply a scheduling theoretical approach to find a best scheduler for WDM EPONs. We show by means of extensive simulations that a multidimensional scheduling approach using results from scheduling theory can provide much better bandwidth management by means of better wavelength utilization than a static wavelength assignment. We also show that an online scheduling approach can provide lower queueing delays than a cyclical offline scheduling approach. We conclude with some specific guidance on future research on bandwidth management for WDM EPONs. © 2006 Optical Society of America

OCIS codes: 060.4250, 060.4510.

1. Introduction

As network traffic increases, service providers will scramble to provide more bandwidth to their customers. Currently, service providers are bringing fiber closer and closer to the end user. An evolution is taking place where fiber penetration in the access network gets deeper as costs come down and subscriber demand goes up. This evolving penetration will eventually result in fiber-to-the-subscriber, i.e., business or home. At this point, there will exist a potential for very large transmission capacity. Current technology is realizing 128 10 Gbps wavelength channels [1] on a single fiber strand. These channels exist only in a region of the optical spectrum for which we currently have effective transceivers. This number of channels is sure to increase as we can effectively tap into the other regions of the optical spectrum. We need to also be aware that there are several strands of fiber co-existing in a fiber optic cable. These fibers serve little purpose for fault tolerance because they are in close proximity to other fiber strands in the same cable. To offer effective protection capabilities the fiber strands need to be spatially diverse. For this reason, these fiber strands can simply be viewed as media for more transmission capacity. Naturally, as the access network is currently evolving in the sense of gaining deeper fiber penetration, it will continue to evolve by activating more transmission channels. These additional channels will be realized through wavelength-division multiplexing (WDM) within a fiber strand and space-division multiplexing across fiber strands. This evolution is sure to be driven by demand and cost.

In this article, we explore the topic of bandwidth management for WDM EPONs that takes advantage of an evolving number of available optical transmission channels. The article is organized as follows. Section 2 reviews related work for bandwidth management for WDM EPONs. In Section 3 we introduce the topic of bandwidth management for WDM EPONs and outline the problem of grant sizing. Section 4 explores the grant scheduling approach of treating time and wavelength assignment as separate problems. Section 5 examines a multidimensional scheduling approach to grant scheduling, whereby we make use of results from scheduling theory. Section 6 presents a discussion of our simulation results. Finally, Section 7 concludes our work with a final summary of our results and a brief outline of future research on bandwidth management for WDM EPONs.

2. Related Work

The topic of medium access control (MAC) protocols, which includes bandwidth management, for WDM EPONs is still in its infancy [2]. Recently, attention has been given to MAC protocols for WDM EPONs. A WDM PON architecture, called SUCCESS, that provides migration paths from time-division multiplexed (TDM)-PON to WDM-PON through the use of tunable transmitters in an optical line terminal (OLT) was proposed in Ref. [3]. The SUCCESS architecture provides an economical migration path by sharing the tunable transceivers at the OLT among multiple separate PONs. SUCCESS was enhanced to add protection capabilities through the use of a ring topology in SUCCESS-HPON [4]. In Ref. [5] the authors developed batch and sequential scheduling algorithms to handle the sharing of the tunable transceivers at the OLT that are used for both downstream and upstream transmission.

The WDM PON architecture presented in Ref. [6] integrates both APON and EPON through the use of the proposed byte-size clock medium access protocol. This protocol requires all nodes to be synchronized and its frame format is not compliant with EPON.

A bandwidth management algorithm for WDM EPONs was proposed in Ref. [7]. In the WDM EPON considered, all Optical Network Unit(s) (ONU)s are assumed to have a transceiver for each upstream/downstream wavelength channel. The proposed WDM IPACT-ST algorithm in Ref. [7] is a multichannel extension of Interleaved Polling with Adaptive Cycle Time (IPACT) [8] that schedules transmission of grants on the first available upstream wavelength channel. In Ref. [9] refinements of the limited grant sizing technique in the context of WDM EPONs were developed and evaluated.

In the existing studies [7,9], grants were scheduled in a next-available-first-fit manner. Our study is complementary to the existing studies in that we explore grant scheduling in detail and relate the grant scheduler in a WDM EPON to general scheduling theory. Throughout, we consider a heterogeneous WDM EPON in which the ONUs differ in their WDM capabilities, as will typically be the case in EPONs that are in the process of being upgraded.

3. Bandwidth Management for WDM EPONs

Bandwidth management for a WDM EPON can be broken into two subproblems: grant sizing and grant scheduling. Grant sizing determines the size of a grant to an ONU, and grant scheduling determines when and on which wavelength channel to schedule the grant. [Many works in this area refer to dynamic bandwidth allocation (DBA) as consisting of both the grant sizing and scheduling.] In single-channel EPONs, scheduling is greatly simplified because bandwidth requests can only arrive one at a time. This is because MutiPoint Control Protocol (MPCP) REPORT messages cannot arrive concurrently, and therefore immediately scheduling the request at the next available upstream time is the natural policy. This results in a round robin scheduler. On a multichannel EPON, there are multiple transmission channels, which makes it possible for transmissions to be scheduled on several wavelengths. Efficiently scheduling grants across multiple wavelengths is thereby necessary to fully utilize the network resources.

We will first briefly investigate grant sizing methods and then delve into the topic of scheduling these grants across multiple wavelengths with constraints imposed by each ONU's WDM architecture. Existing methods of grant sizing [8] can be employed in the multichannel case. Let G_i denote the grant size to ONU i, R_i the requested grant size from ONU i, and $G_{\max}$ the maximum grant size:

- Fixed-$G_i = G_{\max}$,
- Gated-$G_i = R_i$,
- Limited-$G_i = \min(R_i, G_{\max})$.

All these grant sizing schemes simply account for information from a single MPCP

REPORT message which is appropriate for single channel EPONs. However, multiple MPCP REPORT messages can be received concurrently on a multichannel EPON. Hence, it is prudent to have grant sizing mechanisms account for knowledge of multiple ONU's requests in an attempt to provide better bandwidth allocation decisions. This essentially implies offline grant sizing schemes.

Two fundamental approaches to grant scheduling are investigated in this paper:

- Separated time and wavelength assignment.
- Joint time and wavelength assignment (i.e., multidimensional scheduling).

In this paper we consider only time and wavelength for multidimensional scheduling. However scaling the scheduling techniques to more dimensions, such as space to support multiple fiber strands in a single fiber optic cable, can be achieved with straightforward extensions to the mechanisms presented here.

4. Separated Time and Wavelength Assignment

In this approach, scheduling is broken into two separate problems:

- Wavelength assignment.
- Bandwidth management in time on each wavelength.

We can view this approach as partitioning a physical EPON into multiple virtual EPONs separated by wavelength. Suppose we have a three-wavelength EPON with six ONUs. Table 1 shows the ONU wavelength support and load (high, medium, low). Figure 1 shows how these ONUs could be partitioned into multiple virtual EPONs through wavelength assignment. ONUs 1 and 2 support only λ_1, so they are forced to be assigned to this wavelength. Since ONU 6 has a high traffic load, it is assigned exclusively to λ_3. This leaves λ_2 for ONUs 3, 4, and 5.

Any of the DBA algorithms presented in the literature, e.g., Refs. [8,10–12], can be selected to manage the scheduling in time. The selected algorithm will then be instantiated for each wavelength (virtual EPON). For example, in the case of bandwidth-guaranteed polling proposed by Ma *et al.* [10], bandwidth-guaranteed polling tables are maintained for each wavelength (virtual EPON) on the physical PON. This separates the time and wavelength assignments, thereby breaking the scheduling into two subproblems, time scheduling and wavelength assignment. This results in a total of three subproblems for the bandwidth management: (1) grant sizing, (2) ONU wave-

Table 1. ONU Wavelength Assignment Criteria

Node	Supported Channels	Traffic Load
1	λ_1	low
2	λ_1	medium
3	λ_1, λ_2	low
4	λ_1, λ_2	low
5	λ_1, λ_2, λ_3	low
6	λ_1, λ_2, λ_3	high

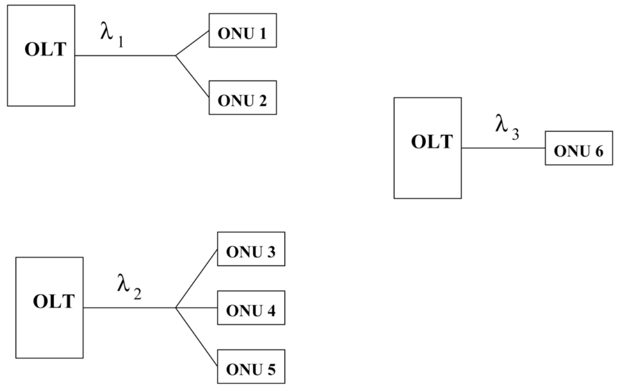

Fig. 1. Illustration of multiple virtual EPONs separated by wavelength.

length assignment, and (3) time scheduling. Subproblems 1 and 3 are solved by single-channel DBA algorithms proposed in the research literature, e.g., Refs. [8,10–12].

We now explore the ONU wavelength assignment problem by outlining a few heuristic approaches to the wavelength assignment problem. We have adapted these heuristics from the literature on wavelength-routed optical WDM networks [13]. Unlike the light-path networks discussed in Ref. [13], EPONs require no routing because the links are single hop (ONU to OLT and OLT to ONU). For this reason, wavelength assignment is all the more critical in WDM EPONs. The two heuristic paradigms are:

- Static wavelength assignment—once wavelengths have been assigned for upstream and downstream for each ONU, the assignment remains static.
- Dynamic wavelength assignment—periodic adjustments are made to ONU wavelength assignments to adapt to changes in PON traffic loading.

4.A. Static Wavelength Assignment

Once a wavelength has been assigned for transmission to, and another wavelength has been assigned for reception from an ONU, this assignment is never changed. This makes the wavelength assignment very simple to implement but does not allow for adaptation to changing bandwidth requirements. Some heuristics for selecting a wavelength for static assignment when an ONU registers are: [13],

- Random—randomly select a wavelength supported by the ONU.
- Least assigned—select the wavelength that is supported by the ONU and has the fewest ONUs already assigned to it.
- Least loaded—select the wavelength that is supported by the ONU and has the least load assigned to it (assuming that this information is available when a new ONU is registered).

4.B. Dynamic Wavelength Assignment

This approach requires more logic but is still very simple to implement. Dynamic wavelength assignment allows for ONUs to be reassigned to different wavelengths. The ideal goal is to adapt the wavelength assignment to the bandwidth requirements. The OLT keeps track of the utilization of each wavelength and uses this information with a heuristic for deciding on ONU wavelength assignment changes. Two possible dynamic wavelength assignment approaches are:

- Load sifting—when the utilization of a wavelength reaches a certain threshold, *loadedThresh*, the OLT attempts to move an ONU assigned to this wavelength to another wavelength. The OLT selects the first ONU assigned to this wavelength and assigns it to the wavelength supported by the ONU that has the lowest current utilization. If this utilization is higher than another threshold, *addThresh* (the threshold for adding new ONUs), then this ONU will not be reassigned and the algorithm moves to the next ONU. Once a single ONU has been reassigned to a different wavelength, the wavelength has been successfully load shifted and the algorithm moves to the next heavily loaded wavelength.
- Quality of service (QoS) load shifting—this method adjusts the loading similarly to load shifting but instead uses a QoS measure instead of actual load. Therefore shifting takes place to maintain certain QoS levels.

Dynamic wavelength assignment has the limitation that, when large time scales are used for wavelength assignment changes, the assignment changes will not allow the bandwidth management to adapt quickly enough to traffic load changes (traffic in the access network is known to be bursty). This means that for large time scales of assignment changes, the performance of dynamic wavelength assignment will not provide much of an improvement over static wavelength assignment. Further, as the time scale of assignment changes approaches that of an assignment change for each grant, dynamic wavelength assignment becomes a multidimensional or joint time and wavelength scheduling problem that is solved more efficiently by the approaches discussed in the next section.

5. Joint Time and Wavelength Assignment (Multidimensional Scheduling)

For conventional time-division multiplexed single-channel EPONs, the bandwidth management problem is limited to scheduling the upstream transmissions on the single wavelength channel. For WDM EPONs, the bandwidth management problem is expanded to scheduling the upstream transmissions on the different upstream wavelengths supported by the ONUs. In other words, in WDM EPONs not only decisions on when and for how long to grant an ONU upstream transmission but also on which wavelength channel to grant the upstream transmission are required to make efficient use of the transmission resources.

The multichannel EPON scheduling problem can be formulated using the scheduling theory notation defined in Ref. [14]. Theoretical analysis of all the scheduling models discussed in this article can also be found in Ref. [14]. Scheduling theory is concerned with scheduling a set of jobs with specific processing times to be executed on a set of machines as efficiently as possible with respect to an optimization criterion. We can view each ONU as representing a job, its grant size as defining its processing time, and the channels used for transmission on the EPON as representing the machines. In scheduling notation, a scheduling problem is defined by a triple $\alpha|\beta|\gamma$, where α describes the machine environment (e.g., single machine, parallel machines), β describes the processing characteristics and constraints, and γ describes the objective to be minimized.

For the formulation of the multichannel EPON grant-scheduling problem in the scheduling theory notation, we let P denote the number of identical parallel machines (channels) that defines our machine environment. Our only processing characteristic or constraint is M_j, which refers to machine (channel) eligibility constraints. Specifically, M_j is the set of machines (channels) on which job (ONU) j can be executed (transmitted). Let C_j denote the time at which the transmission for ONU j is complete, w_j the priority weight of ONU j, and $C_{\max}$ the maximum completion time, or the make-span, of the schedule produced. With this notation we define the scheduling problem with the objective to minimize the unweighted sum of the completion time as $P|M_j|\Sigma_j C_j$. Similarly, we define the scheduling problem with the objective to minimize the weighted completion time as $P|M_j|\Sigma_j w_j C_j$ and the problem with the objective to minimize the make-span of the schedule as $P|M_j|C_{\max}$.

In the above models, the M_j processing constraint is required because each ONU has its own subset of supported channels. If all ONUs supported transmission on all wavelengths we could remove the machine eligibility constraint to obtain models $P\|\Sigma_j C_j$, $P\|\Sigma_j w_j C_j$, and $P\|C_{\max}$.

Our performance objective in designing a scheduler for a WDM EPON is to increase resource (i.e., channel) utilization and lower queueing delays experienced by frames in transit across the EPON. To see how these performance objectives relate to the objectives from scheduling theory we first explore in detail all the component delays in a scheduling cycle. We start by defining cycle length, which we also refer to as the GTG, or GATE-to-GATE delay, as the time between back-to-back grants to an ONU. All the component delays of a scheduling cycle are visualized in Fig. 2. GTR is the GATE-to-REPORT delay (since we append the REPORT at the end of the transmission window, GTR is equal to the transmission time of the grant), and RTG is the REPORT-to-GATE delay, which includes the propagation delay from ONU to OLT. Using GTR and RTG, we express the cycle length as GTG=GTR+RTG.

The GTR delay is completely dependent on the grant size, and the RTG delay is related to the quality of the scheduling. The STG is the schedule-to-GATE delay, which is the time between the OLT scheduling an ONU's next grant and the time the grant

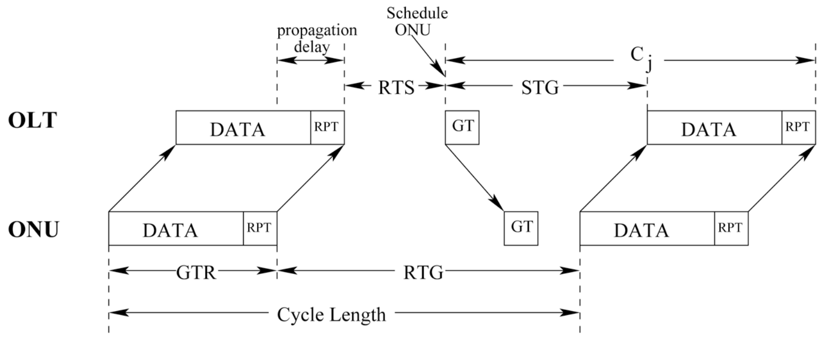

Fig. 2. Illustration of scheduling cycle.

starts. The STG includes a propagation delay from OLT to ONU, a GATE message transmission time and propagation delay from ONU to OLT. We note that the OLT to ONU propagation of the GATE message can be masked through interleaving with an online scheduler, to be discussed later in Subsection 5.B. The STG along with the grant time represents the completion time of an ONU's transmission from the point in time it is scheduled, i.e., C_j=STG+GTR. Since the grant time (or size) is not determined by the scheduler, the scheduler can only work to minimize the variable portion of completion time, i.e., the STG. Minimizing C_j will minimize STG. The RTS is the REPORT-to-schedule delay and is the delay from the OLT receiving a REPORT from an ONU to when the ONUs' REPORT is considered for scheduling by the OLT. Thus, REPORT-to-GATE delay is composed of the RTS and STG delays, i.e., RTG=RTS+STG.

We will illustrate in our simulation results that the RTS delay depends on the type of scheduler used as well as the length of the overall schedule. The length of the schedule is $C_{\max}$ (note: $C_{\max}=\max_j C_j$) for a specific scheduler type (offline as will be discussed later in Subsection 5.A). We notice that, when we examine offline schedulers, minimizing $C_{\max}$ will help minimize RTS.

With a GATED grant sizing mechanism (i.e., a grant sizing scheme that grants every ONU's requested transmission time), the queueing delay can be anywhere between one and two times the cycle length. An Ethernet frame that arrives at the ONU just before the REPORT message is prepared for transmission is counted in that REPORT message and is transmitted at the end of the next grant (i.e., one cycle length of queueing delay). However, an Ethernet frame arriving just after the REPORT has been generated will not be accounted for until the next REPORT message (a cycle length later); it will then be transmitted a cycle length after that REPORT is sent to the OLT, for a total queueing delay of two cycle lengths. Lowering the cycle length, therefore, will inherently lower the queueing delay.

We now outline and compare two broad paradigms for scheduling grants for upstream transmissions on the different upstream wavelengths in a WDM EPON, namely, *cyclical offline scheduling* and *online scheduling*.

5.A. Cyclical Offline Scheduling

In a cyclical offline scheduler, scheduling decisions are made with full knowledge of all the jobs to be scheduled, including their processing times for a particular scheduling cycle. This is in contrast to an online scheduler, where scheduling decisions are made without complete knowledge. In an offline scheduler for a multichannel EPON, the ONUs are scheduled for transmission once the OLT has received current MPCP REPORT messages from all ONUs, allowing the OLT to take into consideration, in the scheduling (and grant sizing), the current bandwidth requirements of *all* ONUs. Since an offline scheduler makes scheduling decisions for all ONUs at once, all of the REPORTs, which are usually appended to the end of the data stream of a gated transmission window from the previous cycle, must be received, as illustrated in Fig. 3.

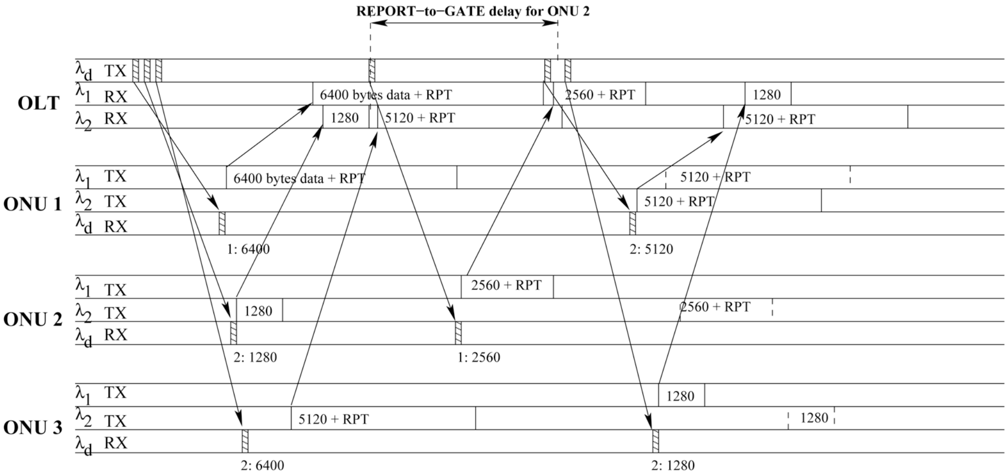

Fig. 3. Illustration of the offline scheduler that introduces the REPORT-to-schedule (RTS) delay.

This requires that the scheduling algorithm be executed after the OLT receives the end of the last ONU's gated transmission window. Because of this, all ONU's grants are scheduled after $C_{\max}$, or the make-span of the preceding schedule. The RTS delay for the last ONU will be negligible, however, the RTS may not be negligible for the other ONUs. This RTS will introduce further queueing delays in the ONUs because it introduces additional delay in the cycle length (GTG) for an ONU.

In the example illustration of Fig. 3, the worst-case RTS is for ONU 2. ONU 2 is the first ONU to complete its granted transmission and send its REPORT to the OLT. Wavelength λ_1 was available for ONU 2 right after ONU 1's transmission and could have accommodated ONU 2's next request a guard time after that. Since the OLT waits until all REPORTs have been received before performing all the grant scheduling for the next cycle, ONU 2's grant scheduling is delayed past the end of ONU 1's transmission. Specifically, ONU 2 has to wait until the REPORT is received from ONU 3, which in this case is the last ONU to send its REPORT to the OLT.

5.A.1. Parallel Machine Models and Solutions

We now utilize results from scheduling theory to find the best offline scheduler for the multichannel EPON. As mentioned above, if all ONUs support transmission on all wavelengths, we can remove the machine eligibility constraints and obtain the following models:

- $P \| C_{\max}$ is NP-hard [14]; however, the longest processing time (LPT) first rule provides a good upper bound on performance. This rule is $4/3 - 1/3m$, competitive with the optimal. For an algorithm to be ρ competitive means that in the worst case this algorithm is ρ times worse than optimal. With a small number of machines this approaches 1-competitive (i.e., as good as optimal); however, for a large number of machines, this approaches 4/3-competitive.
- $P \| \Sigma_j C_j$ is solved to optimality by the shortest processing time (SPT) first rule. However, when we add the weights, the problem can be solved only by the heuristic weighted shortest processing time first rule [14]. This heuristic is $1/2(1+\sqrt{2})$-competitive with the optimal.

If we include the machine eligibility constraints, least flexible job (LFJ) first scheduling is proved optimal for $P|M_j, p_j=1|\Sigma_j C_j$ and $P|M_j, p_j=1|C_{\max}$ (p_j is the processing time of job j) if M_j have a special nesting structure. The special nesting structure between the machine eligibility constraints for two ONUs holds if one and only one of the following relationships holds for ONUs j and k:

- $M_j = M_k$.
- $M_j \subseteq M_k$.
- $M_k \subseteq M_j$.
- M_j and M_k do not overlap.

This nesting structure of the supported wavelengths is not guaranteed for all WDM upgrade scenarios of EPONs.

The $p_j = 1$ component means that the bandwidth requirements of all the ONUs are equal, or that each bandwidth unit is considered as a separate job. This could produce fragmentation if the individual bandwidth units are not scheduled consecutively (which would increase the number of gaps in the schedule due to guard times required between transmission grants). If we remove the $p_j = 1$ requirement and/or the nesting structure of M_j, then LFJ is a heuristic for the problem.

5.A.2. Unrelated Machine Models and Solutions

Another possible approach is to loosen our model and see that $P|M_j|\Sigma_j w_j C_j$ can be viewed as a special case of $R \| \Sigma_j w_j C_j$, where R refers to an unrelated machine environment in which each machine executes a job at a different speed. As well, $P|M_j|\Sigma_j C_j$ can be viewed as a special case of $R \| \Sigma_j C_j$, and $P|M_j|C_{\max}$ can be viewed as a special case of $R \| C_{\max}$. For machines that are in M_j, we set the execution time on these machines to one per processing unit, and for machines not in M_j, we set the execution time on these machines to ∞.

We will now pursue solution methods for these alternative unrelated machine environment models for the WDM EPON scheduling problem.

- $R \| \Sigma_j w_j C_j$ is strongly NP-hard [14] and can be formulated as an integer program solvable by branch and price methods (a form of branch and bound) [15,16].

- $R\|\Sigma_j C_j$ can be formulated as an integer program with a special structure that yields an integer solution under linear program relaxation. A common method used to solve this problem is *weighted bipartite matching*. A weighted bipartite matching problem in which the number of jobs and number of machines are equal is referred to as an assignment problem. The time complexity of weighted bipartite matching is $O[n(m+n \log n)]$, where in our case m is the number of wavelengths and n is the number of ONUs. The integer program for weighted bipartite matching is formulated as follows:

$$\text{minimize} \sum_{i=1}^{m} \sum_{j=1}^{n} \sum_{k=1}^{n} k p_{ij} x_{ikj},$$

subject to

$$\sum_{i=1}^{m} \sum_{k=1}^{n} x_{ikj} = 1, \forall_j,$$

$$\sum_{j=1}^{n} x_{ikj} \leqslant 1, \forall_i, \forall_k,$$

where k is the scheduling position, p_{ij} is the grant processing time for ONU j on channel i (either ONU grant time for supported channel i or ∞ for nonsupported channel i), x_{ikj} are binary variables representing whether position k on machine (channel) i is selected for job (ONU) j, m is the number of machines (channels), and n is the number of jobs (ONUs).

- $R\|C_{\max}$ is NP-complete [14] and can be solved by heuristics proposed by Davis and Jaffe [17].

5.A.3. Application of Scheduling Algorithms to Grant Scheduling

The simplest (in terms of computational complexity) offline scheduler to implement for an EPON is the LFJ heuristic. This basically requires that for each scheduling cycle we schedule the ONUs that support the fewest channels first. We can improve on this scheduler by breaking ties using LPT first (this is proved to provide at worst 4/3-competitive solutions for minimizing the make-span, $C_{\max}$). We refer to this scheduler as LFJ w/LPT. Alternatively, we can break ties using SPT first, which provides an optimal solution for $P\|\Sigma_j C_j$. We refer to this scheduler as LFJ w/SPT.

Note the following relationship between RTS and $C_{\max}$: $\text{RTS}_{\max} = C_{\max} - C_{\min}$. Since the worst-case RTS is dependent on $C_{\max}$, schedulers that minimize $C_{\max}$ will minimize the worst-case *RTS*. For this reason, using LFJ w/LPT is preferred for an offline scheduler.

To compute a schedule, the OLT creates an ordered list of ONUs that depends on the scheduling algorithm used. The OLT, traversing the list in order, schedules the next ONU's transmission on the first avaliable channel supported by that ONU. For example, if the LFJ-LPT scheduling algorithm is to be used, the OLT orders ONUs in increasing order of the number of channels they support, ONUs supporting the fewest channels would appear first. The OLT follows up by breaking ties with the grant size. For LPT, larger grant sizes would appear first. Alternatively, for SPT the smaller grant sizes would appear first. Once a sorted list of ONUs is produced this list is used to generate the schedule as discussed above. Essentially, application of these scheduling algorithms is implemented by a sorting operation.

5.B. Online Scheduling

In an online scheduler, scheduling decisions are made without complete knowledge of all the jobs and/or their processing times. Online scheduling problems can be classified by their incomplete specifications as follows [18,19]:

- Scheduling jobs one by one.
- Jobs arriving over time.
- Unknown processing times.

The scheduling jobs one by one classification refers to scheduling problems where each job must be scheduled without the knowledge of any other jobs. The jobs arriving over time classification (also known as *nearly online scheduling*) refers to scheduling problems where jobs arrive over time, i.e., all jobs do not become available for scheduling at the same time. This is similar to the one-by-one classification, except that if more than one job arrives at the same time they can be scheduled together using an offline technique. The unknown processing time classification refers to situations where the processing time of a job is unknown even after it arrives. In other words, a job must be executed to completion before its processing time is known.

In an EPON, jobs (i.e., MPCP REPORTs) arrive over time and the processing time of the job (i.e., grant size) is known in full once the job arrives (actually, once the bandwidth management or DBA algorithm in the OLT computes the grant size from the MPCP REPORT message). Although MPCP REPORTS arrive over time, the OLT of a single-channel EPON can receive only a single REPORT message at a time because of the single wavelength channel used for reception. However, an OLT of a multichannel EPON can receive multiple REPORT messages within a short time interval. Each REPORT message is received on a separate wavelength channel.

5.B.1. Immediate Online Scheduling

We first examine an online scheduler that schedules ONUs one at a time without considering the bandwidth requirements of other ONUs. For this type of online scheduler (i.e., scheduling jobs one by one), a given ONU is scheduled for upstream transmission as soon as the OLT receives the REPORT message from the ONU. A basic online scheduling policy for the WDM EPON is to schedule the upstream transmission for an ONU on the wavelength channel that is available the earliest among the channels supported by the ONU, which we refer to as *next available supported channel* (NASC) policy. This is our variation on an algorithm proposed by Graham [20] nearly 40 years ago called the List algorithm for identical parallel machines. This algorithm schedules jobs one by one and assigns them to the next available machine. Since each of our nodes has a different set of supported channels, we had to develop a variation on this algorithm where we consider only supported machines (i.e., channels). The list algorithm has been proved to be $2-1/m$-competitive for the make-span optimality criterion, where m is the number of machines. With a single machine (channel) this is 1-competitive (i.e., the same as optimal), and with a large number of machines (channels) this approaches 2-competitive (i.e., 2 times worse than optimal).

Figure 4 illustrates online scheduling using NASC for an EPON with three ONUs. Notice that the 2560 byte upstream transmission from ONU 2 (its second grant) is scheduled on the earliest available supported wavelength, namely, wavelength 1 and is timed by the OLT such that it is separated from the preceding transmission on wavelength 1 by ONU 1 by the guard interval. In an offline scheduler, ONU 2 would have to wait until the REPORT message is received from ONU 3 before it would be

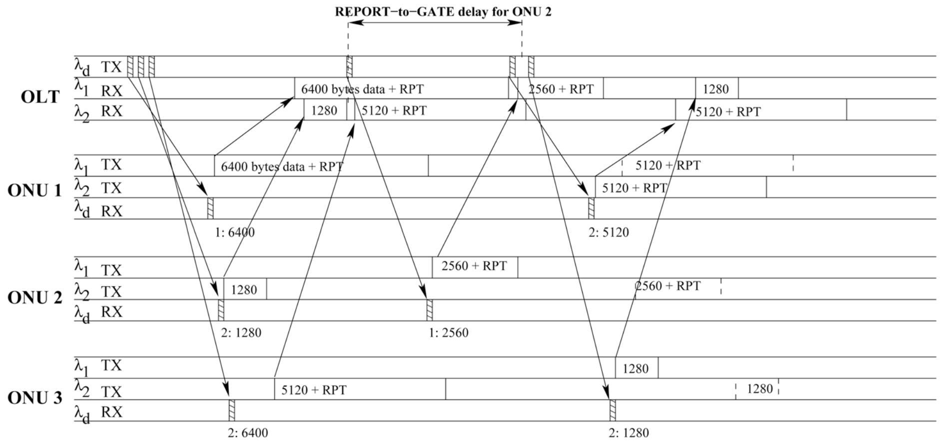

Fig. 4. Illustration of online scheduling with the next available supported channel (NASC) policy. The illustration includes one downstream wavelength, λ_d, and two upstream wavelengths, λ_1 and λ_2, which are supported by all three ONUs. Each ONU reports its queue occupancy in the REPORT (RPT) message, which is appended to the current upstream transmission. Upon receipt of a RPT message the OLT immediately schedules the next upstream transmission for the corresponding ONU and sends a GATE message (illustrated by the dashed message) indicating the wavelength and length (in bytes in the illustration) of the granted transmission to the ONU.

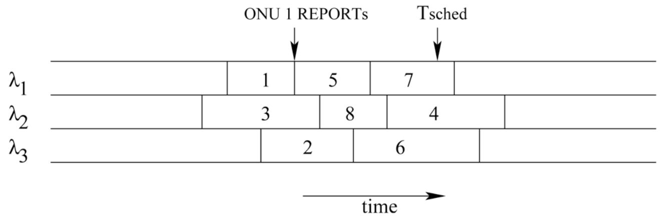

Fig. 5. Illustration of the online interval scheduler. ONU 1 is scheduled a short time period before its NASC is available; during this time ONUs 3, 2, 5, and 8 report. The online interval scheduler will schedule all these ONUs together using an offline technique at time T_{sched} noted in this figure.

scheduled. The dashed line boxes show the grants for cycle 2 for the offline scheduler. We can see all the cycle 2 grants are pushed further out in time compared with the online scheduling. ONU 2, the first ONU to complete cycle 1, experiences the largest difference in delay between online and offline scheduling.

We can also note that the RTS delay noted for offline scheduling does not exist in an online scheduler since an ONU is scheduled as soon as its REPORT is received. This means that RTG delay simply reduces to the STG delay or C_j (i.e., ONU transmission completion time). See Fig. 2 for an illustration of these delays.

5.B.2. Just-in-Time Online Scheduling

We note that it is no detriment (in terms of queueing delay) to the ONU to schedule its next grant in just-in-time fashion (i.e., just before the NASC becomes free). This is the latest time that the GATE message can be transmitted without incurring any additional unnecessary queueing delay. During this time period additional ONUs can potentially communicate their requests through REPORT messages, and we will have a larger pool of ONUs to schedule with an offline technique, such as weighted bipartite matching or LFJ w/SPT to minimize $\Sigma_j C_j$. Specifically, we want to allow enough time to send a GATE message to the ONU (i.e., propagation delay+GATE message transmission time) and start receiving data from the ONU (i.e., propagation delay) on the NASC. Let T_{sched} be the latest time that we can schedule an ONU, T_{free} be the free time of the NASC, RTT be a round-trip time from the OLT to the ONU, and T_{GATE} be the transmission time of a GATE message. For just-in-time scheduling we need to schedule an ONU at a time expressed as $T_{\text{sched}}=T_{\text{free}}-RTT+T_{\text{GATE}}$.

We have implemented and provide simulation results for this type of scheduler using LFJ w/SPT. We refer to this scheduler as the online interval scheduler, signifying that it attempts to schedule in an offline fashion ONUs whose REPORTs arrive within a short time interval of each other. We illustrate this approach in Fig. 5. We notice that ONU 1 finishes its transmission on wavelength 1 some time before its NASC becomes available, in this case wavelength 1. An online scheduler would schedule this ONU as soon as its REPORT is received. However, there is no additional delay incurred if the scheduling is delayed to T_{sched} as expressed above. During the time that ONU 1 finishes its transmission and a short period before wavelength 1 becomes free, we can see in Fig. 5 that ONUs 3, 2, 5, and 8 finish their transmissions and send their REPORT messages. The online interval scheduler now schedules these ONUs together using an offline technique. This scheduler has the potential to outperform NASC because for instances where several ONUs are ready for scheduling within a short time interval, this scheduler can maker better scheduling decisions, since it uses an offline scheduling technique (in this case LFJ w/SPT).

6. Simulation Results

The CSIM [21] simulation library was used to develop a behavioral model of the MAC layer aspects of an EPON in the C programming language. Our EPON consisted of 10 ONUs (unless otherwise noted). Each wavelength on our EPON supports 1 Gbps link speed. All loads presented in our results are measured with respect to this link speed. A load of 1.3 represents a traffic load of 1.3 Gbps. Note that preamble and interpacket gap between packet transmissions are not factored into the load calculation, and neither are control packets. The simulations were conducted using self-similar traffic sources [22–24] with a Hurst parameter of 0.75. The ONU to OLT round-trip time was randomly generated according to a uniform distribution $U[100\ \mu s, 200\ \mu s]$, which corresponds to ONU distances of 15 – 30 km from the OLT.

In our simulations we use the GATED technique for the sizing of grants. The GATED grant sizing technique grants each ONU its full bandwidth request. By fixing the grant sizing technique we are comparing the scheduling aspects of the WDM EPON.

In this comparison, our EPON consisted of ten ONUs (five lightly loaded ONUs, five heavily loaded ONUs). The heavily loaded ONUs support transmission on all wavelengths supported on the EPON. The heavily loaded ONUs have double the traffic of lightly loaded ONUs. In other words, the five lightly loaded single-channel ONUs generate 1/3 of the traffic load, and the five heavily loaded multichannel ONUs generate the remaining 2/3 of the total traffic load on the network.

The static wavelength assignment scheme that we simulated used an online scheduling paradigm. The wavelength assignment was fixed using the least-assigned wavelength assignment heuristic when ONUs were registered on the EPON.

6.A. Static Wavelength Assignment versus Multidimensional Scheduling

We have simulated a five-wavelength EPON with the ONU WDM configuration discussed previously. For each of the simulations we vary the grant scheduling technique. Figure 6 plots the queueing delay experienced by Ethernet frames at the ONU for different grant scheduling techniques. Figure 7 plots the cycle length or GATE-to-GATE delay. Figure 8 plots the wavelength utilization under each grant scheduling technique for each of the five wavelengths on the EPON. Comparing static wavelength assignment with the multidimensional scheduling theoretical approaches, we notice that the static wavelength assignment scheme, as the loads increase, overutilizes wavelength 1 [see Fig. 8(a)]. This overutilization is caused by the fact that the static

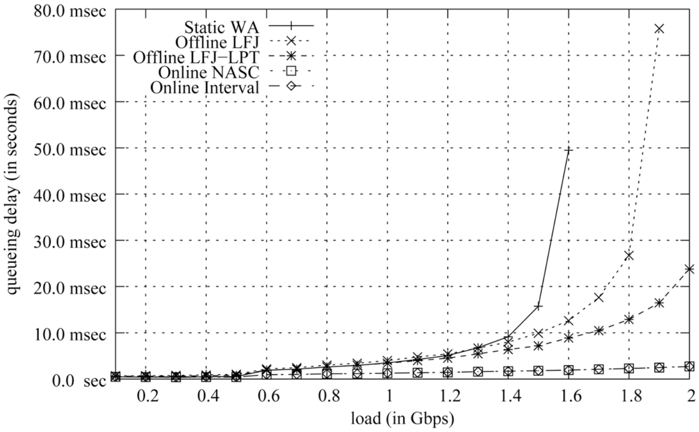

Fig. 6. Comparison of average queueing delay with online and offline scheduling on a five-wavelength EPON as a function of load.

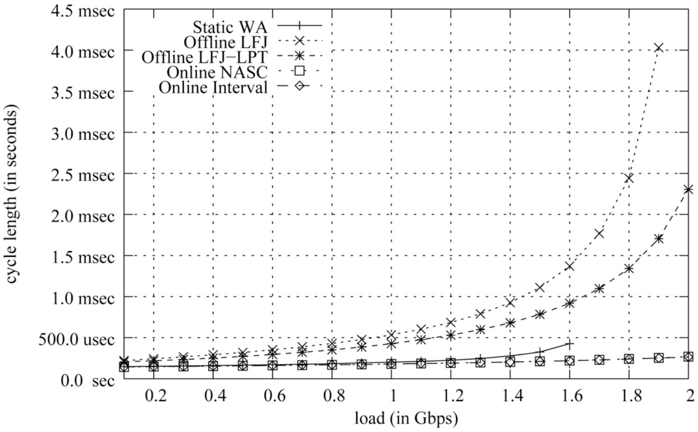

Fig. 7. Comparison of average cycle length with online and offline scheduling on a five-wavelength EPON as a function of load.

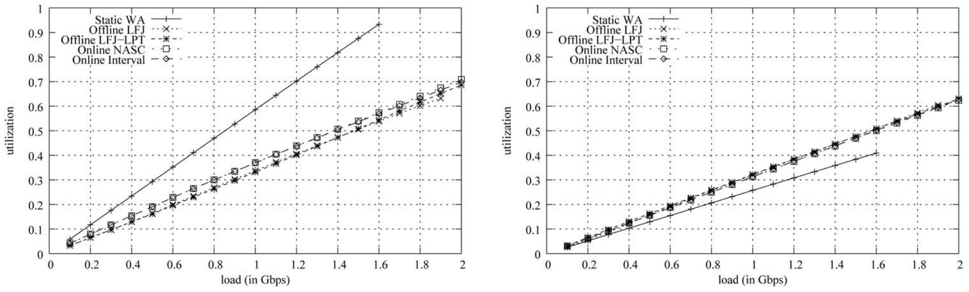

a) Wavelength #1

b) Wavelength #2

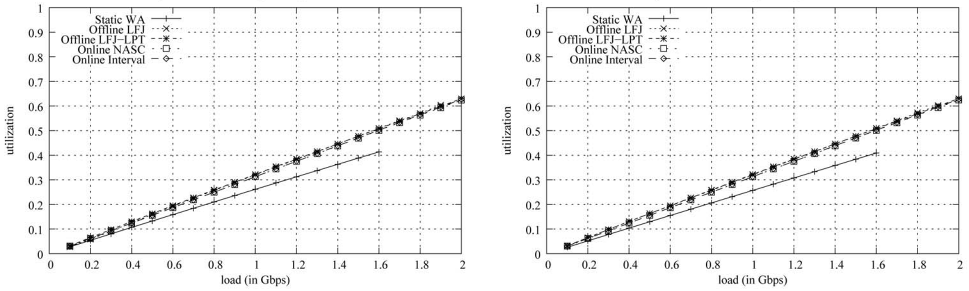

c) Wavelength #3

d) Wavelength #4

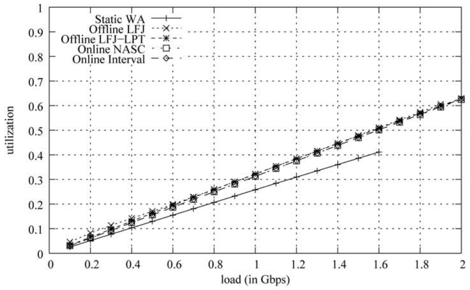

e) Wavelength #5

Fig. 8. Comparison of wavelength utilization under each grant scheduling technique for each of the five wavelengths on the EPON.

wavelength assignment scheme cannot reassign ONUs to other wavelengths as their bandwidth requirements increase to balance the load with the other wavelengths. This overutilization causes ONUs assigned to wavelength 1 to experience very high queueing delays, and at a load of 1.6 Gbps wavelength 1 utilization reaches saturation, causing the queueing delays to approach infinity (see Fig. 6 for average queueing delay). Simulation results for load values above 1.6 Gbps are not shown for the static wavelength assignment scheme because the EPON cannot accomodate traffic above 1.6 Gbps for this type of grant scheduling. Throughout, we have chosen to simulate only load values that can be accomodated by the EPON using the specified grant scheduling technique.

Since the static wavelength assignment scheme uses an online scheduling paradigm, it is possible for certain ONUs to be polled more frequently than other ONUs. Specifically, ONUs assigned to less-utilized wavelengths can be polled more frequently than ONUs that are assigned to more-utilized wavelengths. This allows ONUs that are polled more frequently to record their cycle lengths more frequently, which drives the average cycle length to lower values. In contrast, when the queueing delay is recorded, each Ethernet frame's queueing delay is recorded. By comparing the values of queueing delay and cycle length for static wavelength assignment in Figs. 6 and 7 , respectively, we can see that, although the average cycle length is lower than the offline multidimensional scheduling techniques, the queueing delays are much higher.

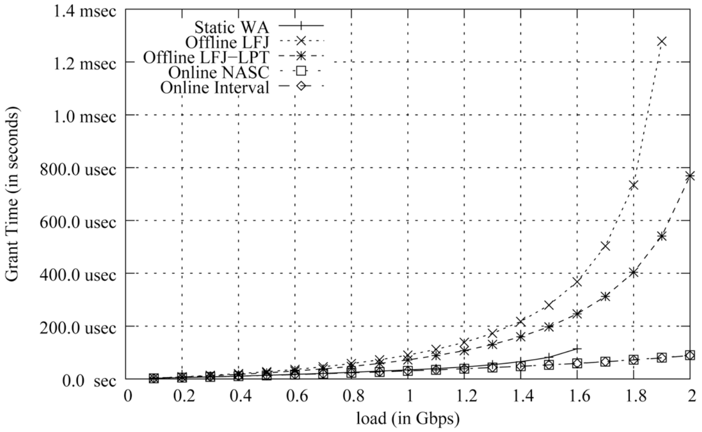

Fig. 9. Comparison of average grant time (i.e., GATE-to-REPORT delay) between the different grant scheduling approaches.

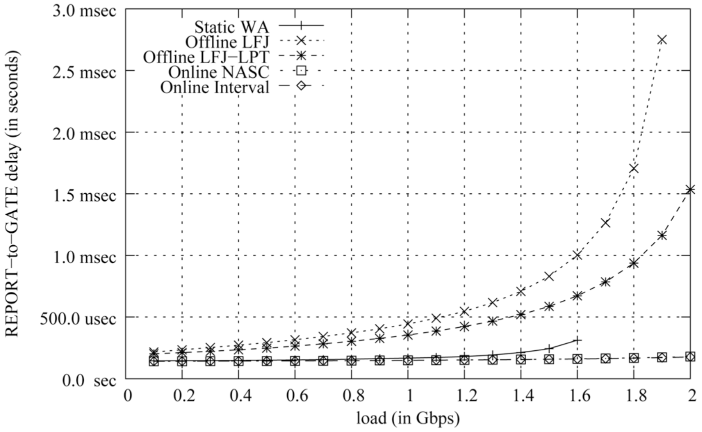

Fig. 10. Comparison of average REPORT-to-GATE delay (includes a propagation delay, average propagation delay is 75 μs) between the different grant scheduling approaches.

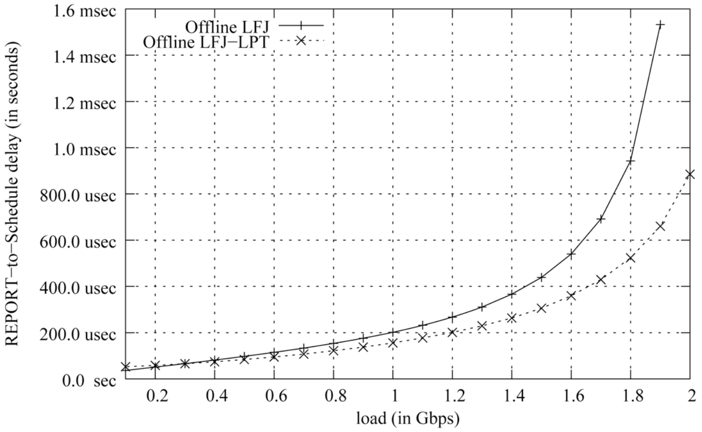

Fig. 11. Comparison of average REPORT-to-schedule delay between the different grant scheduling approaches

6.B. Offline Scheduling versus Online Scheduling

Figures 6 and 7 show average queueing delay and average cycle length or GTG delay. To examine the constituent delays of cycle length, we also plot average grant time or GTR delay, average RTG delay, average RTS, and average STG delay in Figs. 9–12, respectively.

From Fig. 6 we observe that even at low loads the online scheduling techniques (NASC and interval) provide lower queueing delays than the offline techniques (LFJ

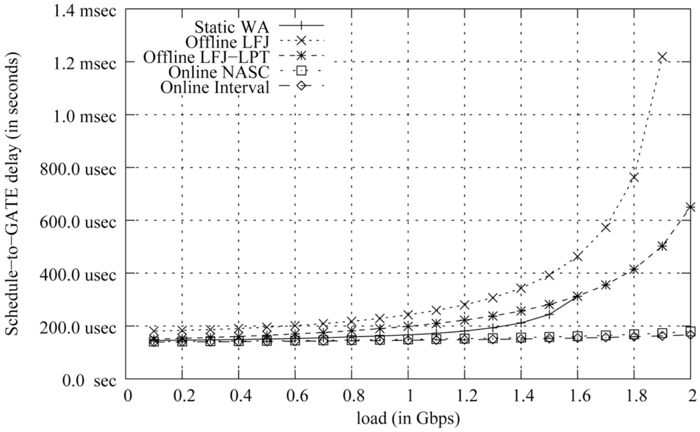

Fig. 12. Comparison of average schedule-to-GATE delay (includes a propagation delay, average propagation delay is 75 μs) between the different grant scheduling approaches.

and LFJ-LPT). We notice that augmenting the LFJ offline technique with LPT provides lower queueing delays than simply LFJ. This is because LPT is a good heuristic for minimizing the make-span or length of schedule, which minimizes the RTS delay. Figure 11 shows that the RTS delay is smaller for the LFJ-LPT scheduler for loads above 0.3 Gbps.

To understand why queueing delays are shorter for online schedulers, we refer back to Fig. 4 and see how the online scheduler, by scheduling ONUs immediately and avoiding the RTS delay, can more efficiently use the available wavelengths, resulting in faster ONU transmission completion. The schedule produced by the offline scheduler for cycle 2 is shown in Fig. 4 by the dashed boxes for comparison. We can see further evidence of this by looking at the average cycle length data shown in Fig. 7. We see that the online schedulers provide lower average cycle lengths, which is more pronounced at higher loads. For further insight into what causes the difference in cycle length between online and offline schedulers, we first recall the delays that contribute to the cycle length (or GTG), namely, the GTR, or size of the grant, and the RTG delay, which consists of the RTS delay and the STG delay. The STG delay, along with the grant time (GTR delay), constitutes the completion time of an ONU's transmission (i.e., C_j), measured from the time it is scheduled. See Fig. 2 for an illustration of these delays.

From Fig. 10 we see again that there is a growing difference in the RTG delay between the offline and online schedulers as the load increases. This is due to the fact that the RTS exists for offline schedulers but not for online schedulers. This RTS causes the RTG to be longer, which in turn makes the cycle length longer. A longer cycle length causes more frames to queue up at an ONU between grants, which causes grant sizes to increase. From Fig. 9 we see that an EPON load of 1.8 Gbps causes an average grant time of 724 μs for offline LFJ, 401 μs for offline LFJ-LPT and 72.6 μs for online NASC. This relationship is a result of the difference in RTS (see Fig. 11). The increasing grant times further cause the other component of RTG, namely, STG, to increase, because schedules become longer as a result of larger grant sizes (see Fig. 12). Further, an offline scheduler cannot mask the propagation delay from OLT to ONU of the GATE message, since it is unable to send GATE messages while other ONUs are still transmitting their grants. So, we can attribute the increasing disparity in RTG between offline and online scheduling to the existence of RTS and the inability of the offline scheduler to mask the propagation delay from OLT to ONU of the GATE message. Figure 14 shows that increasing the number of ONUs on the EPON exacerbates the RTS delay (Fig. 14 is shown with a fixed traffic load.) This in turn increases the disparity in cycle length and queueing delay, which can be seen in Fig. 13. Increasing load, as well as an increasing number of ONUs on the EPON, increases the RTS delay for offline schedulers.

One other observation we make is that, although we expect the queueing delay to be between one and two times the cycle length for a GATED grant sizing mechanism, we see a multiple that approaches 6. Figure 6 shows the queueing delay, and Fig. 7 shows the cycle length. This discrepancy is due to the differences in the manner that

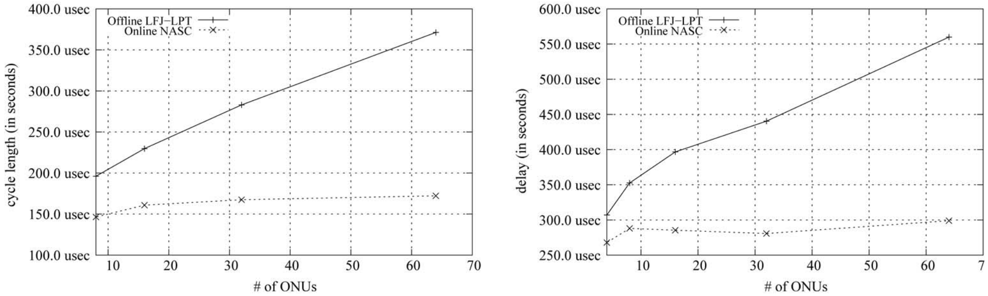

a) Average Cycle Length b) Average Queueing Delay

Fig. 13. Average cycle length and queueing delay versus number of ONU.

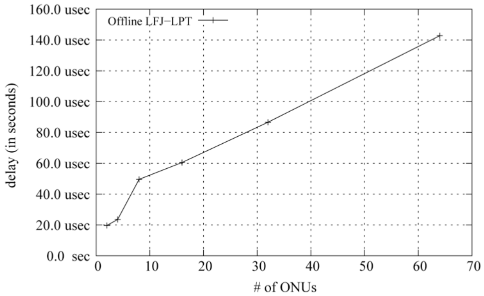

Fig. 14. Average REPORT-to-schedule delay versus number of ONUs.

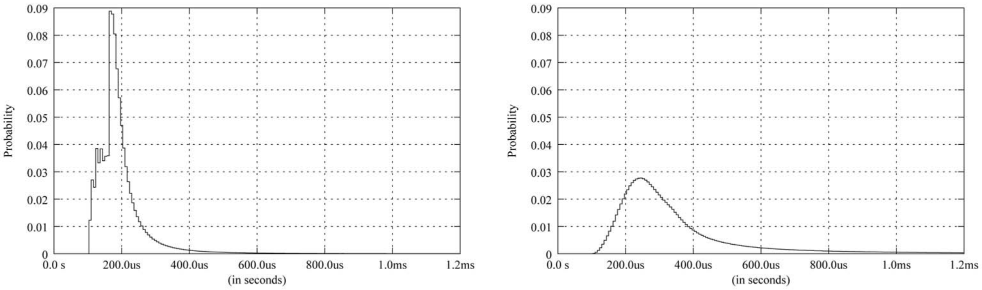

a) Cycle Length b) Queueing Delay

Fig. 15. Comparison of cycle length and queueing delay PDFs for online load =1.4 Gbps.

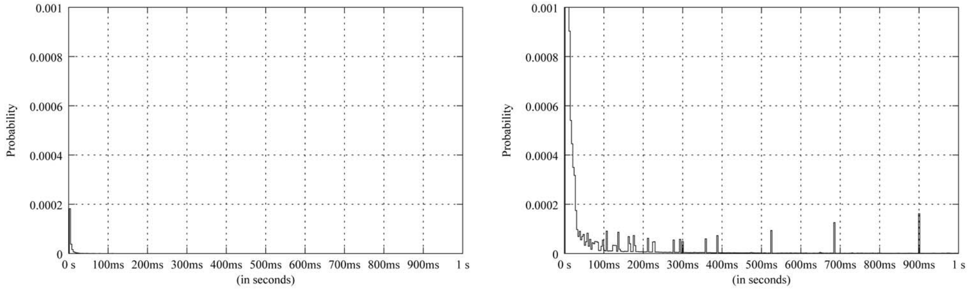

a) Cycle Length b) Queueing Delay

Fig. 16. Comparison of cycle length and queueing delay PDFs (tail) for online load =1.4 Gbps.

cycle length and queueing delay values are recorded. Cycle length values are recorded as the time between back-to-back grants to an ONU; these values are therefore recorded once per ONU per scheduling cycle. Queueing delay values are recorded for every Ethernet frame in transit, and the number of Ethernet frames per grant depends on the size of the grant.

Table 2. Cycle Length and Queueing Delay Values for Certain Percentiles of Their PDFs (Online NASC, load = 1.4 Gbps)

Percentile	Cycle Length (s)[a]	Percentile	Queueing Delay (s)[b]
0.25	0.000167	0.25	0.000240
0.50	0.000186	0.50	0.000300
0.75	0.000220	0.75	0.000420
0.90	0.000273	0.90	0.000766
0.95	0.000333	0.95	0.001365
0.975	0.000420	0.975	0.002598
1.00	0.932954	1.00	1.60489

[a] Expected cycle length value 204.3 μs.

[b] Expected queueing delay value 1.68 ms.

Therefore, we recognize that for longer cycles the queueing delay associated with that cycle is experienced by more Ethernet frames than for shorter cycles. This is a result of the fact that longer cycles create a larger backlog of Ethernet frames. We can see evidence of this in the probability distribution functions (PDFs) of queueing delay and cycle length in Figs. 15 and 16. We see that the PDF of the queueing delay has more weight in the tails. Table 2 shows the cycle length and queueing delay values for certain percentiles of their PDFs. We notice that up to the 75th percentile, the values match the expected one to two times relationship. The relationship does not hold at the 90th and 95th percentiles because of the heavier tail in the PDF of the queueing delay, but it does hold for the 100th percentile, or the maximum value.

6.C. Online NASC versus Online Interval

For this comparison, we have simulated a 64 ONU EPON with four wavelengths, where each ONU supports all four wavelengths. Table 3 shows the difference in queueing delay between online NASC and online interval scheduler, and Table 4 shows their difference in average cycle length. In Table 5 we show the number of ONUs scheduled together in our original 10 ONU, 5 wavelength example. In Table 6 we show the number of ONUs scheduled together for NASC and online interval for the 64 ONU EPON just described.

In our results for the 10 ONU, 5 wavelength EPON, there was no distinguishable difference between the online NASC and online interval schedulers. This was due to the fact that the average number of ONUs scheduled was slightly above 1 (see Table 5), meaning it was very rare that there was an opportunity to schedule ONUs together. To examine the case where the probability of scheduling ONUs together without incurring RTS is high, we need to look at EPONs with a large number of ONUs and where the loads are close to capacity. For this reason we have simulated a WDM-EPON with 64 ONUs that all support the four wavelengths available on the network and looked at loads approaching full capacity. What we noticed, by looking at Table 6, is that the number of ONUs able to be scheduled simultaneously leveled out

Table 3. Comparison of Queueing Delay between NASC and Online Interval Scheduler (64 ONUs, 4 Wavelengths)

Load	Online NASC	Online Interval LFJ-SPT
3.95	12.04	11.94
3.96	15.79	15.18
3.97	19.92	19.41

Table 4. Comparison of Cycle Length between NASC and Online Interval Scheduler (64 ONUs, 4 Wavelengths)

Load	Online NASC	Online Interval LFJ-SPT
3.95	3.55	3.03
3.96	4.01	3.58
3.97	5.05	4.31

Table 5. Number of ONUs Scheduled for Online Interval Scheduler (10 ONUs, 5 Wavelengths)

Load	Number of ONUs Scheduled
0.1	1.13
0.2	1.13
0.3	1.13
0.4	1.13
0.5	1.13
0.6	1.13
0.7	1.14
0.8	1.15
0.9	1.16
1.0	1.16
1.1	1.17
1.2	1.18
1.3	1.19
1.4	1.20
1.5	1.21
1.6	1.23
1.7	1.24
1.8	1.26
1.9	1.28
2.0	1.31

Table 6. Number of ONUs Scheduled for Online Interval Scheduler (64 ONUs, 4 Wavelengths)

Load	Number of ONUs Scheduled
3.95	60
3.96	60
3.97	60

at 60 ONUs, because even at high loads there will be an ONU per wavelength, in this case four that are unavailable for scheduling. We can see by looking at Table 4 that the there is an improvement in cycle length as the load approaches full network capacity. This translates into a half-second improvement in the average queueing delay as the network reaches full capacity. This can be seen in Table 3. The NASC algorithm, since it is derived from the list algorithm, is $(4/3-1/3m)$-competitive with optimal, where m is the number of wavelengths. This means that the higher the number of wavelengths, the less competitive with the optimal the NASC algorithm will be. This will give the online interval scheduler an advantage that grows with the number of wavelengths used on the WDM EPON.

7. Conclusion

In conclusion, we have formulated the problem of bandwidth management for WDM EPONs and utilized scheduling theory (a branch of optimization theory) to find the best solution for the scheduling part of the problem. Our simulation results indicate that the scheduling used on a WDM EPON can severely impact the network's performance with respect to queueing delays experienced at the ONUs. We have kept our solutions general enough to support evolutionary WDM EPON upgrades [25]. Simulation results demonstrate that for the GATED grant sizing scheme, a scheduling theoretical approach to multidimensional scheduling can allow for better bandwidth management by adapting wavelength assignment to instantaneous bandwidth requirements. Further, an offline scheduling approach induces a delay, called the REPORT-to-schedule (RTS) delay, that increases with increasing network load and increasing number of ONUs. This RTS delay causes longer delays between back-to-back grants, which results in larger queueing delays experienced by Ethernet frames waiting for transmission at an ONU. For this reason an online scheduling approach appears to be generally preferable. However, some grant sizing techniques with incomplete knowledge of all the ONUs requests, which is the case for an online scheduling approach, will result in worse grant sizing decisions than an offline approach.

We have also shown that the online scheduling problem for the WDM EPON is properly classified as a jobs arriving over time classification, where some ONUs could be scheduled together with an offline technique making our pursuit of an appropriate result for offline scheduling from scheduling theory worthwhile. We conclude that the best scheduler for a WDM EPON is a hybrid online–offline scheduler. We have pointed out that an ONU can be scheduled in just-in-time fashion (for example, the NASC available time minus a RTT to the OLT) without imposing any unnecessary dead time on a wavelength. We use this insight in our design of a hybrid online–offline scheduler that we have shown to provide lower queueing delays as the network approaches full capacity. Using a hybrid scheduler also benefits the grant sizing technique by allowing for more-comprehensive knowledge of the ONUs' requests when one is making grant sizing decisions.

References and Links

1. Lucent's LambdaXtreme transport, http://www.lucent.com.
2. A. Banerjee, Y. Park, F. Clarke, H. Song, S. Yang, G. Kramer, K. Kim, and B. Mukherjee, "Wavelength-division-multiplexed passive optical network (WDM-PON) technologies for broadband access: a review," JOM **4**, 737–758 (2005).
3. Y.-L. Hsueh, M. S. Rogge, S. Yamamoto, and L. G. Kazovsky, "A highly flexible and efficient passive optical network employing dynamic wavelength allocation," J. Lightwave Technol. **23**, 277–286 (2005)
4. F. T. An, D. Gutierrez, K. S. Kim, J. W. Lee, and L. G. Kazovsky "SUCESS-HPON: a next-generation optical access architecture for smooth migration from TDM-PON to WDM-PON," IEEE Commun. Mag. **43**(11), S40–S47 (2005).
5. K. S. Kim, D. Gutierrez, F. T. An, and L. G. Kazovsky, "Design and performance analysis of scheduling algorithms for WDM-PON under SUCCESS-HPON architecture," J. Lightwave Technol. **23**, 3716–3731 (2005).
6. C. Xiao, B. Bing, and G. K. Chang, "An efficient reservation MAC protocol with preallocation for high-speed WDM passive optical networks," in *Proceedings of IEEE INFOCOM* (IEEE, 2005), pp. 444–454.
7. K. H. Kwong, D. Harle, and I. Andonovic, "Dynamic bandwidth allocation algorithm for differentiated services over WDM EPONs," in *Proceedings of the IEEE International Conference on Communications Systems (ICCS)* (IEEE, 2004), pp. 116–120.
8. G. Kramer, B. Mukherjee, and G. Pesavento, "IPACT: A dynamic protocol for an Ethernet PON (EPON)," IEEE Commun. Mag. **40**(2), 74–80 (2002).
9. A. R. Dhaini, C. M. Assi, and A. Shami, "Dynamic bandwidth allocation schemes in hybrid TDM/WDM passive optical networks," in *IEEE Consumer Communications and Networking Conference* (IEEE, 2006), Vol. 1, pp. 30–34.
10. M. Ma, Y. Zhu, and T. Cheng, "A bandwidth guaranteed polling MAC protocol for Ethernet passive optical networks," in *Proceedings of IEEE INFOCOM* (IEEE, 2003), Vol. 1, pp. 22–31.
11. C. M. Assi, Y. Ye, S. Dixit, and M. A. Ali, "Dynamic bandwidth allocation for quality-of-service over Ethernet PONs," IEEE J. Sel. Areas Commun. **21**, 1467–1477 (2003).
12. M. P. McGarry, M. Maier, and M. Reisslein, "Ethernet PONs: a survey of dynamic bandwidth allocation (DBA) algorithms," IEEE Commun. Mag. **42**(8), S8–S15 (2004).
13. H. Zang, J. Jue, and B. Mukherjee, "A review of routing and wavelength assignment approaches for wavelength-routed optical WDM networks," Opt. Networks Mag. **1**(1), 47–60 (2000).
14. M. Pinedo, *Scheduling: Theory, Algorithms, and Systems*, 2nd ed. (Prentice-Hall, 2002).
15. Z. L. Chen and W. B. Powell, "Solving parallel machine scheduling problems by column generation," INFORMS J. Comput. **11**, 78–94 (1999).
16. J. M. Van Den Akker, J. A. Hoogeveen, and S. L. Van De Velde, "Parallel machine scheduling by column generation," Oper. Res. **47**, 862–872 (1999).
17. E. Davis and J. M. Jaffe, "Algorithms for scheduling tasks on unrelated processors," Acta Math. Acad. Sci. Hung. **28**, 721–736 (1981).
18. A. Fiat and G. J. Woeginger, *Online Algorithms: The State of the Art* (Springer, 1998).
19. A. P. A. Vestjens, "Online machine scheduling," Ph.D. thesis (Eindhoven University of Technology, 1997).
20. R. L. Graham, "Bounds for certain multiprocessing anomalies," Bell Syst. Tech. J. **45**, 1563–1581 (1966).
21. CSIM (Mesquite Software), http://www.mesquite.com.
22. G. Kramer, "On generating self-similar traffic using pseudo-Pareto distribution," UC Davis Technical Brief (University of California, Davis), http://wwwcsif.cs.ucdavis.edu/~kramer/papers/self_sim.pdf.
23. M. S. Taqqu, W. Willinger, and R. Sherman, "Proof of a fundamental result in self-similar traffic modeling," ACM SIGCOMM Comput. Commun. Rev. **27**, 5–23 (1997).
24. K. Park and W. Willinger, *Self-Similar Network Traffic and Performance Evaluation* (Wiley-Interscience, 2000).
25. M. P. McGarry, M. Maier, and M. Reisslein, "WDM Ethernet Passive Optical Networks," IEEE Commun. Mag. **44**(2), S18–S25 (2006).