

Document downloaded from:

<http://hdl.handle.net/10251/92335>

This paper must be cited as:

García Roger, D.; Artundo Martínez, I.; Ortega Tamarit, B. (2011). On the Conditions that Justify Dynamic Reconfigurability in WDM-TDMA Optical Access Networks. *Journal of Optical Communications and Networking*. 3(4):259-271. doi:10.1364/JOCN.3.000259



The final publication is available at

<http://doi.org/10.1364/JOCN.3.000259>

Copyright The Optical Society

Additional Information

On the Conditions That Justify Requiring Dynamic Reconfigurability in WDM-TDMA Optical Access Networks

David Garcia-Roger, Iñigo Artundo, and Beatriz Ortega

Abstract—In a passive optical network with a hybrid wavelength division multiplexing time division multiple-access scheme, implementing reconfigurable wavelength assignment is complex; hence the need to determine the conditions for which the capacity improvements justify requiring reconfigurability over adopting a more inexpensive fixed wavelength assignment. Fixed and reconfigurable approaches to wavelength assignment are modeled and evaluated under nonstationary traffic conditions. The performance improvement is obtained in terms of bit rate gain relative to the nominal bandwidth and depends on the number of wavelength channels as well as the magnitude of the load offered by the optical network units. In addition, frame delay and frame loss in relation to the bit rate performance are obtained for Pareto and exponentially distributed traffic. Simulations show that when introducing reconfigurability, typical peak bit rate gains with respect to the fixed case are 17%, and maxima of 175% are potentially possible when traffic demands are particularly uneven.

Index Terms—Access networks; Fiber optics; Optical communications; Wavelength assignment; Wavelength routing.

I. INTRODUCTION

The growth in the worldwide demand for high-capacity communication services compels network operators to offer their subscribers faster bit rates while preserving optimum guarantees of quality of service (QoS). Because of its vast intrinsic bandwidth, fiber-optic communication has become a mature technology to rely on when deploying the future-proof access networks that will satisfy the needs of existing and emerging intensive data services (e.g., peer-to-peer file transfers of emerging disk storage formats) and multimedia services (e.g., high-definition or even 3D video streaming).

The infrastructure considered that connects subscribers to the service provider hence is an optical fiber access network, which comprises an optical line terminal (OLT) and several optical network units (ONUs). ONUs are situated close to the subscriber's premises. The OLT is located at the operator's central office (CO) and controls a specific pair of wavelength channels: downlink (OLT to ONU) and uplink

(ONU to OLT); OLT and ONUs send and receive modulated pulses of light carrying information simultaneously. The passive optical network (PON) is a widespread architecture, a point-to-multipoint topology built from passive optical filters and couplers, which needs only a single (but shared) feeder link to connect multiple ONUs to an OLT. A PON avoids the cost of deploying multiple individual fibers and its low-cost passive components reduce the energy consumption and maintenance expenses characteristic of active optical devices.

With the purpose of handling the anticipated rise in network traffic and to provide further access capacity, wavelength division multiplexing (WDM) upgrade of PONs—from a single wavelength channel to multiple channels—is being investigated. A WDM-PON makes possible the distribution of wavelength channel pairs to serve specific ONUs. In fact, there are several OLTs and each ONU is registered to a particular OLT. The wavelength channel pair between the OLT and its associated ONUs is assigned by a multiplexer/demultiplexer located at an intermediate remote node (RN). Since several ONUs share the capacity of each pair, a degree of statistical multiplexing is expected.

Certain WDM-PON architectures have been conceived as a cost-effective, scalable, and smooth capacity upgrade from an already installed time division multiple-access (TDMA-) PON: hybrid solutions known as WDM-TDMA-PONs [1]. While downlink traffic is broadcast by the OLT to its ONUs, uplink traffic here requires a collision-free multiplexing method like TDMA to allow proper sharing of the medium among ONUs, according to a dynamic bandwidth allocation (DBA) scheme. TDMA-PONs have been deployed extensively: widespread standards are the IEEE 802.3ah, which specifies a 1 Gb/s Ethernet PON (1G-EPON)—1.250 Gb/s symmetrical—and the ITU G.984 family of recommendations, which define gigabit PON (GPON)—up to 2.488 Gb/s; more recently IEEE 802.3av, which defines 10G-EPON, and ITU G.987 recommendations, which define 10G-PON, have been proposed as next generation PON 10 gigabit standards.

The wavelength resources of a WDM-TDMA-PON are managed according to an assignment scheme. Fixed wavelength assignment (FWA) schemes determine the wavelength channel pairs beforehand and then allocate the specific pairs statically to groups of ONUs. An improved use of resources may be achieved by dynamic and reconfigurable wavelength assignment (DRWA) schemes. DRWA schemes allow flexible wavelength routing from the OLT to the appropriate ONUs in response to traffic variability.

The authors are with the Optical and Quantum Communications Group (iTEAM), Universidad Politécnica de Valencia, 46022 Valencia, Spain (e-mail: dagarro@iteam.upv.es).

The study of DRWA schemes is part of the current research in WDM-PONs [2–5]. Approaches depend on the specific architecture, but they may be classified according to objective as either extra capacity or load balancing schemes. Extra capacity schemes provide supplementary wavelength channels when necessary to fix situations of traffic congestion (e.g., sudden needs for bandwidth to serve new traffic sources at unpredictable locations, to cope with link faults, etc.). Load balancing schemes manage the set of wavelength channels to get the most out of the system capacity (e.g., to obtain a uniform amount of traffic load per wavelength channel, and, in general, to achieve significant statistical multiplexing gains). Remarkable examples of architectures that use extra capacity [2,3,6], and load balancing schemes [1,4,5] have been proposed to date.

To the best of our knowledge, only a few previous contributions deal with the performance of DRWA in comparison to FWA [3–5], but in such a manner that it is not straightforward to quantify their conclusions or extrapolate them to completely different circumstances (e.g., the scenarios in [3,4] are static and no evidence is provided about their performance under evolving traffic demands). Since the introduction of reconfigurability is a complex and costly issue, a key subject still pending proper study is the assessment of the degree of convenience of DRWA within a range of traffic conditions. This paper addresses the subject from the load balancing viewpoint, focusing on providing general insights for evaluating whether the benefits of DRWA in terms of bit rate gain are worthwhile and the subsequent impact on frame loss and frame delay.

The rest of this paper is organized as follows. Section II describes the generic optical access network scenario modeled. Section III presents the two approaches to wavelength assignment considered in this paper. Section IV introduces the results of the simulation work, Section V briefly discusses additional issues of wavelength reconfigurability, and finally Section VI concludes the paper, and highlights the original contributions of the work.

II. THE WDM-TDMA OPTICAL ACCESS NETWORK

The proposed hybrid WDM-TDMA access network scenario is shown in Fig. 1. There are N ONUs (one per household, assuming that each household is a multi-dwelling unit), that both send and receive users' traffic. Specifically, ONU- i is located at or near the premises of household i (with $i = 1, \dots, N$) which is served by the ONU. The WDM resources of the network comprise M pairs of wavelength channels $\{\lambda_j^d, \lambda_j^u\}$ (with $j = 1, \dots, M, M < N$) with wavelength pair $\{\lambda_j^d, \lambda_j^u\}$ managed by OLT- j and multiplexed and demultiplexed at the CO, which assigns a certain wavelength channel pair $\{\lambda_j^d, \lambda_j^u\}$ to one or several ONUs (e.g., ONU- i) for the downlink and uplink communications. The RN connects the OLTs to the ONUs. Theoretically, there is no limit to the number of ONUs assigned to a channel pair, although achievable data rates per customer and power losses in the splitting process at the RN are aspects that will limit this ratio. It is important to emphasize that for generalization reasons, the model does not specify the explicit physical details of the underlying TDMA-PON infrastructure. The physical architecture and implementation of the RN

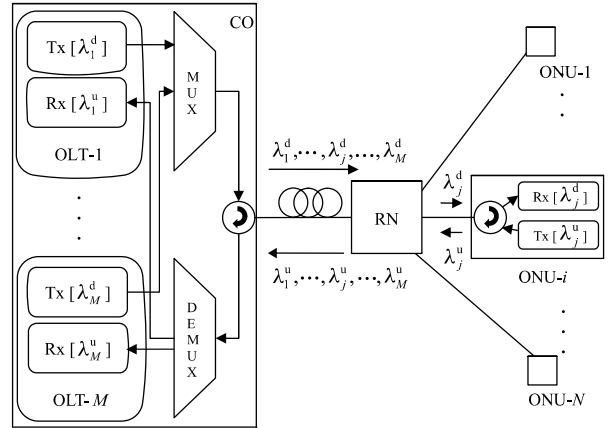


Fig. 1. The WDM-TDMA optical access network scenario.

will depend on the access network architecture and topology. For a fully passive RN, only filters and splitters/combiners will be needed in a broadcast-and-select topology, or arrayed waveguide gratings (AWGs) and AWG routers to multiplex or demultiplex the different wavelengths from/to the ONUs in a WDM-PON topology. For large splitting configurations, cascaded devices could be used with repeaters and amplifiers in an active RN, to compensate for the losses at every stage and extend reach if needed. For this, Raman amplification or erbium doped fibers pumped from the CO could be used. To provide reconfiguration capabilities at the RN (allowing for protection and balancing), tunable power splitters or optical switches (latched, thermal, etc.) could be used for tree topologies, or reconfigurable add-drop multiplexers (ROADMs) for ring topologies concatenating several RNs. Recent works have proposed the placement of ROADMs closer to the network edges, to provide higher flexibility as their cost is being progressively reduced [7,8].

Transmissions in the downlink direction from OLT- j to its ONUs use broadcasting; in this case the receiver at ONU- i tuned to λ_j^d carefully selects the data aimed at it. Communications in the uplink direction follow a TDMA approach that assigns non-overlapping transmission time slots to ONUs in order to control the access of the ONUs to the shared wavelength channel. As a result the WDM-PON is equivalent to several single-wavelength PONs (also known as “virtual PONs”), each one with the corresponding pair of wavelength channels.

A. The Uplink Model

Figure 2 depicts the access network from the uplink point of view. The nominal data rate of the upstream link from an ONU to the OLT is C_u , while the data rate of the access link from a user to an ONU is c_u . Frames arrive at ONU- i from the subscribers, which are the traffic sources of the model. A multiple-access scheme based on TDMA is operating at each OLT and therefore the generated frames must wait in a queue of finite size Q , located at ONU- i until OLT- j gives permission to transmit those frames (the thicker line). The DBA scheme arbitrates the access to the uplink wavelength channel (λ_j^u) of




Fig. 2. The uplink model.

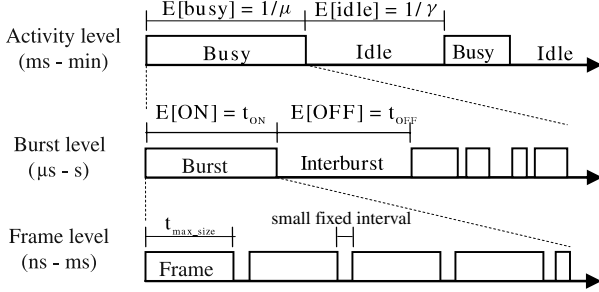


Fig. 3. The traffic generation model in the access network.

the information waiting for transmission at the queues of each ONU registered to OLT- j .

B. Traffic Modeling Levels

Figure 3 illustrates the modeling approach to the traffic generated by the sources of data (i.e., the subscribers served by the ONUs, in the uplink case). As shown, three degrees of detail are considered: the activity, burst and frame levels. The activity level ranges from some milliseconds to several minutes, and thus is important from the DRWA scheme viewpoint because it works on such timescales, identifying periods of unbalanced wavelength channel load (see Subsection III.B). At this level, traffic behaves according to busy/idle periods. Busy periods group together several trains of frames (bursts) closely related in time, which are surrounded by significant gaps (idle periods). The durations are distributed according to a general (arbitrary) distribution with means $E[\text{busy}] = 1/\mu$ and $E[\text{idle}] = 1/\gamma$, respectively. Parameters μ and γ may be summarized as the activity factor α , the fraction of time for which a traffic source is expected to be busy:

$$\alpha = \frac{E[\text{busy}]}{E[\text{busy}] + E[\text{idle}]} = \frac{\gamma}{\gamma + \mu}. \quad (1)$$

The burst level ranges from some microseconds to seconds. Note that from the uplink viewpoint, this timescale is particularly relevant because both the channel access method (TDMA) and the DBA scheme mediate between the shared medium and the data bursts. Burst and interburst durations behave according to specifically distributed ON/OFF periods with mean $E[\text{ON}] = t_{\text{ON}}$ and $E[\text{OFF}] = t_{\text{OFF}}$. Since numerous research efforts have corroborated the self-similar nature of

access networks traffic (see, e.g., the influential [9]), Pareto distributed ON/OFF periods will be assumed. The Pareto distribution is a power law probability distribution with probability density function

$$f_X(x) = \alpha \frac{x_m^\alpha}{x^{\alpha+1}}, \quad \text{for } x > x_m \quad (2)$$

and $f_X(x) = 0$ otherwise, where x_m is the scale parameter (here interpreted as the minimum possible value of burst/interburst length), and α is the shape parameter (related to the Hurst parameter $H = (3 - \alpha)/2$, which states how relatively high the probabilities of values far from x_m are). The expected value of a random variable following the Pareto distribution is

$$E(X) = \frac{\alpha x_m}{\alpha - 1} \quad (3)$$

(if $\alpha \leq 1$, $E(X)$ does not exist), while the variance is

$$\text{var}(X) = \left(\frac{x_m}{\alpha - 1}\right)^2 \frac{\alpha}{\alpha - 2} \quad (4)$$

(if $\alpha \leq 2$, $\text{var}(X)$ does not exist). Pareto distributions will characterize the burst and interburst lengths (both expressed as multiples of the transmission time $t_{\text{max_size}}$ of a frame of maximum size). The load ρ_i offered to ONU- i relative to the nominal data rate c_u is

$$\begin{aligned} \rho &= \frac{E[\text{ON}]}{E[\text{ON}] + E[\text{OFF}]} \\ &= \frac{\alpha_{\text{ON}} x_{m\text{ON}} (\alpha_{\text{OFF}} - 1)}{\alpha_{\text{OFF}} (\alpha_{\text{ON}} (x_{m\text{OFF}} + x_{m\text{ON}}) - x_{m\text{OFF}}) - \alpha_{\text{ON}} x_{m\text{ON}}}. \end{aligned} \quad (5)$$

Usually, the values of the minimum burst length $x_{m\text{ON}}$ and the scale parameters α_{ON} and α_{OFF} are motivated by real measurements (e.g., Ethernet frames), and thus for a given ρ the value of the last parameter, $x_{m\text{OFF}}$, may be deduced as

$$x_{m\text{OFF}} = \frac{\alpha_{\text{ON}} x_{m\text{ON}} (1 - \rho) (\alpha_{\text{OFF}} - 1)}{\alpha_{\text{OFF}} \rho (\alpha_{\text{ON}} - 1)}. \quad (6)$$

For comparison purposes the results obtained (see Section IV) will be judged against a classic, Poisson-based model, which assigns values to the ON/OFF periods according to the exponentially distributed distribution

$$f_X(x) = \lambda e^{-\lambda x}, \quad \text{for } x \geq 0 \quad (7)$$

and $f_X(x) = 0$ otherwise, where λ is the rate parameter (i.e., $E(X) = \lambda^{-1}$). From Eq. (3) the expected value of the exponential distribution may be equal to its Pareto counterpart

$$\lambda = \frac{\alpha - 1}{\alpha x_m}. \quad (8)$$

The frame level comprises the nanosecond to millisecond timescale. In a given ON period, frames are generated one after another with a small fixed interval to account for the interframe gap and the preamble. Frames of fixed size are assumed, with size equal to the maximum frame size (i.e., each needs $t_{\text{max_size}}$ seconds to be generated). Additionally, the last

frame has a variable size when the burst length is not a multiple of the maximum frame size.

C. Dynamic Bandwidth Allocation

The actual fraction of the nominal bandwidth experienced at wavelength channel λ_j^u by a busy ONU- i depends on the number n of current simultaneously busy ONUs that are registered with OLT- j , i.e., the fewer the ONUs potentially sharing λ_j^u , the greater the amount of bandwidth that each ONU receives. The best case is when there is only one busy ONU and there are no frames from other ONUs with transmission pending, because the ONU may send data at the nominal bit rate.

By means of the DBA, the OLT knows which ONUs need to transmit frames, and it issues time slots so that ultimately each of the n ONUs with frames pending transmission experiences a rate of C_u/n , i.e., an equal share of the nominal capacity available in the medium—however, note that when a time slot is granted, the ONU has exclusive access to the OLT and therefore during the allowed transmission window, the ONU will send data at the nominal rate C_u . The fair share of the uplink bandwidth, as assumed and implemented in the present work, is achievable in both EPON—by means of interleaved polling with adaptive cycle time (IPACT) with limited service discipline [10]—and GPON—by adapting IPACT to the fixed polling cycle nature of the uplink transmission in GPON [11].

Finally, note that from the downlink viewpoint, even if a medium access control scheme is not needed, the total bandwidth is also shared because, although the information transmitted by the OLT arrives at all ONUs, the ONUs will discard data not specifically addressed to them.

III. WAVELENGTH ASSIGNMENT SCHEMES

This section describes the basics of the two principal wavelength assignment schemes that may be put into practice when deploying an optical access network. The two schemes have in common that a procedure for assigning each ONU to a specific OLT is defined; thus the set of wavelength channel pairs is regarded as a resource that has to serve ONUs according to the specifications of each approach, with the ultimate goal of maximizing the average uplink/downlink bit rate per ONU. The static or dynamic nature of the associations established between an OLT wavelength channel pair and ONUs gives way to the following two approaches to wavelength assignment.

A. Fixed Wavelength Assignment

Under a FWA scheme, the wavelength channel pair associations established between an OLT and several ONUs (Fig. 1) are fixed and cannot be updated in real time. Assuming the case that ONUs offer and receive traffic of similar nature and intensity (homogeneous traffic), the WDM-TDMA-PON with FWA is conceptually equivalent to a set of standard TDMA-PONs operating separately (e.g., each wavelength being

supplied to a different passive star coupler that in turn feeds the fibers to an unchangeable set of ONUs). Typically, this is the result of network planning that assumes that the average bit rate requirements per ONU will be similar and stable with time, and therefore the wavelength assignment may be decided in advance and fixed during the network deployment. In turn, this determines k (the maximum number of ONUs that can be served per virtual PON while satisfying the average bit rate requirements). This approach results in a uniform segmentation of the N ONUs that divides the whole network into groups of ONUs, each group being served by a specific OLT; hence, the necessary number of wavelength channel pairs (and thus OLTs) is $M = \lceil N/k \rceil$.

To optimize a network implementing FWA, a prior forecasting of the bandwidth utilization is a requisite; however, this type of prediction is static, assumes network homogeneity, and is subject to some degree of uncertainty. Therefore, it might not hold properly during the operation of a nonstationary network scenario, where users behave unexpectedly following unaccounted traffic profiles. One alternative for coping with this problem could be designing the network with such overload situations in mind, reducing the size of ONU groups registered to the same OLT and deploying then a greater (but expensive) number of wavelength pairs; however, this approach might incur a waste of such extra capacity during periods of average or below-average traffic load.

Moreover, scalability issues arise under a FWA paradigm because it is complex to increase the relative capacity per ONU while maintaining the topology: in that case the operator would need to: first, deploy extra wavelength channels at the CO; second, route the extra wavelengths through several new star coupler ports; third, decrease the number of branching fibers at the remaining couplers; and fourth, deploy fiber to connect the new star couplers to the ONUs that were left without service in the first step. Overall, it is a complex, expensive and time-consuming task.

B. Dynamic Reconfigurable Wavelength Assignment

Under a DRWA scheme, the wavelength channel pair link (Fig. 1) between OLT- j and ONU- i may be modified conveniently in real time by altering the uplink/downlink pair; this is possible because it is assumed that any ONU may be served by an arbitrary OLT. Determining the need for wavelength reconfiguration and the new wavelength pair that fits the needs of ONU- i according to the DRWA scheme is a task computed at the CO, which in turn triggers the association of a new OLT to ONU- i .

Because of economical constraints in current deployments it might be unrealistic to allow for every possible OLT-ONU association. However, it is within the scope of this work to assess the benefits of a future mature reconfigurable scenario where ONUs are colorless (i.e., they are able to operate at any particular wavelength pair).

The DRWA approach adds a dynamic dimension to WDM in order to further increase the statistical multiplexing gain. Specifically, the DRWA scheme facilitates the reallocation of unused capacity among the virtual PONs, which are defined by those ONUs that are connected to the same OLT in a certain

instant of time. In fact, DRWA is to the wavelength domain (for all virtual PONs) what DBA is to the time domain (for a single virtual PON); what is more, both resource management strategies yield significant advantages when implemented in combination. Some significant previous contributions and related work that study the feasibility of a combined DRWA/DBA approach are [12–18]. In [12], bursty and correlated traffic was analyzed in WDM broadcast-and-select networks. The use of an adaptive TDMA protocol here highlighted the importance of the node selection and wavelength assignment time. In [13] the focus is on modern ring topologies in metropolitan networks, but proposing a dynamic access to the medium through an explicit reservation-based hybrid WDM–TDMA reconfigurable control mechanism. Buffer size and the corresponding queuing delay are put forward as main limiting factors when accessing the optical medium. More recently, focusing specifically on the upstream in WDM/TDMA EPON networks, the authors in [14] study the efficiency of dynamic scheduling and wavelength assignment for different granting schemes. Colorless ONUs are considered as well, allowing for an easy installation of the infrastructure and management of the stations.

In [15,16] WDM-IPACT is introduced, which is an extension of IPACT for MPCP in a WDM upgraded EPON scenario, which assigns not only time slots for each ONU’s uplink transmission but also wavelength channels (please refer to [17,18] for detailed studies with a similar approach). As in the present contribution, in [15] it is assumed that there is an array of fixed transmitter/receivers (T/R) at the OLT and either a fixed array of T/Rs or tunable T/Rs at the ONU. During the discovery and registration process, the OLT receives information on each ONU’s T/R type (tunable or fixed) and the tuning time. Wavelengths are supported by extending MPCP with a modified REGISTER_REQUEST message [16]. Upon receiving the uplink message, the OLT monitors the utilization of each wavelength and uses this information to decide wavelength assignment modifications. The wavelength assignment information is carried in the GATE message. Two additional messages are proposed in [15]: RECEIVER-CONFIGURATION (from OLT to each ONU to reconfigure the ONU’s filter) and RECEIVER-CONFIGURATION ACKNOWLEDGE (from each ONU to the OLT). The present contribution resorts to these message extensions of MPCP.

During the operation of the network, the main goal of the DRWA scheme is to achieve a homogeneous distribution of the traffic among all wavelength channels. If the strong time constraints of a polling cycle were not present, a straightforward (unrealistic) approach would be to run a computationally expensive optimization function where every polling cycle would decide the minimum wavelength reconfigurations in order to meet all bandwidth constraints. In this paper a practical approach is taken, adding DRWA on the much larger longer timescale of the activity level (up to once every minute). It is assumed that OLTs monitor the utilization of each wavelength pair and exploit this information by combining a “least assigned” approach (the scheme selects the wavelength pair that is managed by the OLT that has the least number of ONUs assigned to it) with a “load balancing” approach (the scheme finds an overloaded wavelength, chooses from it the assigned ONU that exceeds a certain load threshold and tries to move such an ONU to the wavelength that has the

lowest current utilization). Note that although the decisions of the DRWA require control message extensions of MPCP, the wavelength reconfiguration information carried in such messages may be used for both the uplink and downlink [15].

The key procedure consists of transferring busy ONUs from the most loaded OLT to the least loaded OLT in such a way that the numbers of busy ONUs per virtual PON are similar among all PONs, and therefore the global bandwidth is always shared evenly between busy ONUs according to a fairness criterion. In turn, this policy has a useful impact on DBA because the average number of busy ONUs per wavelength channel pair is minimized and more bandwidth is at the disposal of the associated ONUs.

In the present work, the DRWA scheme is a process that monitors the composition of the virtual PONs and knows b_j , the average number of concurrently busy ONUs that are registered to OLT- j ($j = 1, \dots, M$). Note that in the uplink, in a certain time instant, the average number of simultaneously busy ONUs depends directly on the amount of data pending transmission in the queue of each ONU. Let $B = \sum_{j=1}^M b_j$ be the total number of simultaneously busy ONUs. Ideally, to maximize the relative capacity per ONU, each of the M virtual PONs should comprise a total of B/M ONUs on average (i.e., the numbers of busy ONUs per virtual PON should be very similar). Actually, given that in general, B/M is not an integer number, some OLTs should serve an average of $\lceil B/M \rceil$ busy ONUs while the remaining OLTs should serve $\lfloor B/M \rfloor$ ONUs. Thus, in an optimized network, $\max(b_x - b_y) \leq 1$, $x, y \in \{1, \dots, M\}$, $x \neq y$, because $\lceil B/M \rceil - \lfloor B/M \rfloor = 1$. Therefore, if during the functioning of the network $(b_x - b_y) > 1$ (the threshold value) for at least one value of x and y (i.e., the numbers of busy ONUs per channel are not approximately the same), the DRWA scheme transfers a busy ONU (modifying the frequency of the emitter) from an overcrowded OLT (wavelength channel pair) to an unpopulated OLT. Moreover, fewer wavelength channel arrangements per unit of time may be expected if the operator tolerates a certain amount of asymmetry in the ratio of the number of ONUs to the number of OLT- j s and only triggers the wavelength reconfiguration when the system clearly surpass a given upper bound t , $(b_x - b_y) > t > 1$, but at the expense of a smaller reaction time for new conditions. It is especially critical to keep the reconfiguration speed requirements within a reasonable range because the faster the adjusting time, the more complex (and expensive) the optical device will be [19].

The joint operation of DBA and DRWA is fundamental. On the one hand, there is one instance of the DBA scheme running per OLT. Each DBA scheme arbitrates the access of the registered ONUs to the specific uplink wavelength channel so as to cope with the rapid, short-term traffic fluctuations that derive from the bursty nature of the information generated by applications. The order of magnitude of these variations depends on the process that issues time slots to ONUs in the DBA. For example, the smallest size of an Ethernet frame (64 bytes) requires $0.2 \mu\text{s}$ to be transmitted with GPON in a worst case scenario of one frame per time slot; this will be the fastest switching speed required. Note that this amounts to the time that the OLT has to inspect the MAC destination field and forward the frame and that in this case optical wavelength switching is not needed. On the other hand, the DRWA scheme

tracks the mean number of busy ONUs associated in a certain moment with each OLT in order to respond to variations in the medium-term and long-term bit rate performances. Modifications in such traffic patterns stem from the occasional starting or shutting down of an application, as well as other flow-related activities that affect the long-standing values of the bit rate (e.g., physical network failures, laser malfunctions, times of frequent customer connections and disconnections from the operator’s service, or the heavy and disruptive entrance of new customers). Therefore the tuning speed of the lasers at the transmitters can be of the order of tens or even a hundred milliseconds, which can be accomplished by inexpensive, slow-switching optical devices [19].

Since the decisions of the DRWA scheme must not interfere with the operation of the instances of the DBA scheme, the reconfiguration of the wavelength channel pair (i.e., transferring an ONU from one OLT to another) is performed as soon as the last byte of the current time slot is transmitted by the ONU to be reassigned.

One more advantage is that DRWA defers the need to install costly additional fiber infrastructures for the operator (provided that the equipment at the CO and ONUs for making them tunable is already present) because reusing the already deployed resources to meet the bandwidth needs of more users increases the scalability of access networks. However, the advantages of DRWA come at a cost: normally, reconfigurable optical components that allow modifying the virtual topology of the network are more expensive and more difficult to install and configure than their fixed counterparts. Thus, this trade-off must be kept in mind when assessing the convenience of dynamic wavelength reconfigurability.

IV. CONDITIONS THAT JUSTIFY REQUIRING DYNAMIC RECONFIGURABILITY

Before introducing reconfigurability on an access network, it is important to identify under which circumstances wavelength load balancing is profitable. This section presents results showing the degree of theoretical improvement in terms of bit rate that DRWA can attain over FWA, with respect to the relative popularity of several classes of ONU activity factor profiles in a WDM-PON. Reporting the performance of the system in terms of average bit rate is particularly useful for “elastic traffic,” which is the type of traffic generated by services that adapt to the available bandwidth and remain unaffected by changes in the perceived bit rate (generally, elastic traffic is produced by applications that transfer computer files without real time constraints: e.g., web-browsing, peer-to-peer file transfer). Since a minimum bandwidth prerequisite is not needed, elastic traffic will never experience blocking. In contrast, services that produce “non-elastic” traffic (like the streaming flows that result from video and audio applications) require a peak bit rate, a minimum bit rate, or both, so a certain blocking probability might appear. However, non-elastic traffic may be taken into consideration also by assuming that the network resources (available bandwidth) are enough to fulfill the QoS requirements of the aggregated demand, i.e., negligible blocking probability (e.g., less than 10^{-6}).

TABLE I
POPULATION CASES ACCORDING TO ONU PROFILES

	In-profile ONU (uplink)	Out-of-profile ONU (uplink)
Case 1		Heavy web applications $a_{out}^u = 0.3$
Case 2	Web-browsing $a_{in}^u = 0.1$	Internet video streaming $a_{out}^u = 0.6$
Case 3		Peer-to-peer file transfers $a_{out}^u = 0.9$

In this work, the access network is studied according to the three levels described in Section II. At the activity level, ONUs are described in terms of their activity factor a , as belonging to two possible types: “in-profile” and “out-of-profile” ONUs. In-profile ONUs aggregate network subscribers that are using applications which, according to the operator’s expectations, potentially result in a reasonable link load. In contrast, out-of-profile ONUs are aggregating legitimate customers which are using applications that, by their characteristics, cause unanticipated volumes of traffic.

If f is the fraction of out-of-profile ONUs, then there are $f \cdot N$ out-of-profile ONUs and the remaining $(1 - f) \cdot N$ ONUs are the majority—as expected by the operator—of in-profile ONUs. Defining the access network in accordance with the two above-mentioned profiles allows the operator to plan the deployment in proportion to a dominant user mass that makes use of the nominal capacity reasonably, while at the same time considering that there will be a marginal (but taxing) fraction of users that have the potential to overload the network, because their longer activity periods may take full advantage of the access capacity. The present work considers three population cases composed of both profiles as shown in Table I. The three cases are listed in increasing order of activity factor or, what comes to the same, decreasing order of potential statistical multiplexing. An indicative estimate a_{in}^u , a_{out}^u for the value of the uplink activity factor for in-profile and out-of-profile ONUs is given.

The number of ONUs and the magnitude of the activity factor of the traffic sources being served have an impact on the potential statistical multiplexing gain that can be achieved, e.g., the lower the activity factor per ONU, the higher the statistical multiplexing gain; this is a critical aspect in the uplink, because it results in a lower number of simultaneous ONUs with packets in their buffers pending transmission. In fact, in a perfectly optimized network, wavelength resources will be shared by the minimum number of ONUs which are busy, thus improving overall link utilization. However, precise expected activity factors may be extremely difficult to predict in advance because they depend on the aggregation of the users’ behavior, which in turn is determined by the type of services. As shown in Table I, a reasonable range of values for the offered load extends from $a = 0.1$ for web-browsing (a few user-generated requests followed by a reading time) to $a = 0.9$ for peer-to-peer file transfer (with several users establishing numerous connections to upload data simultaneously).

The main focus of studying the access network at the activity level is assessing the bit rate performance in the long term: the steady state of the system. Considering a black box approach to the system in equilibrium, one of the most important (and

straightforward) results in queuing theory, Little’s Theorem, may be applied [20], as long as nothing in the black box is creating or destroying frames. Little’s Theorem,

$$E[n] = \Lambda \cdot E[\tau], \quad (9)$$

relates three basic magnitudes that may be identified with parameters of the present study: (i) $E[n]$, the time averaged number of customers in the system (i.e., simultaneous ONUs with traffic sources in the busy state that are registered to each OLT); (ii) $E[\tau]$, the average duration that each customer spends in the system (i.e., the length of the busy intervals per ONU); and (iii) Λ , the long-term average arrival rate of customers for service (i.e., the arrival rate of idle to busy transitions). Note that Little’s Theorem is true for general arrival and service processes and therefore its estimates are insensitive to the arbitrary distribution for the busy/idle duration; however Little’s Theorem only verifies whether there is statistical equilibrium (to be precise, whether, in the long term, the aggregated amount of offered traffic from/to all ONUs does not surpass the uplink/downlink throughput; otherwise, queues will continue to increase towards infinity).

Computing the probability distribution of $E[n]$ is useful for determining the average relative bit rate performance W . For example, for the downlink, $W = \sum_{n=1}^{N/M} \frac{C_d}{n} \Pr(n)$, which is the value of the available bit rate share C_d/n weighted by its probability $\Pr(n)$. Little’s Theorem is simple, but even when assuming a FWA scheme, an exact computation of $\Pr(n)$ requires enumerating the Markov chain process with all possible system states and their transitions. Furthermore, when a DRWA is in use, the complexity of the behavior is qualitatively harder to characterize and scales badly with the number of wavelength channel pairs that are in use; consequently the present assessment resorts to simulation in order to estimate the probability distribution of n .

Note that the activity level allows describing the downlink performance, but in the uplink because of the interaction between the medium access control mechanism, the DBA scheme, and the wavelength assignment scheme, the uplink bit rate performance should be supplemented with equally important information about frame delay and frame loss caused by the finite buffer located at each ONU. Such details are provided by simulating the burst level (see Subsection IV.B).

The simulation-based performance evaluation has been carried out using the advanced modeling, analysis, and simulation environment Möbius [21], which can simulate—and under certain conditions even numerically solve—the continuous-time Markov chain processes associated with the activity level operation of the system model. The burst level is coded in C++, taking advantage of the fact that all model execution in Möbius is done exclusively in that programming language for performance reasons.

A. Average Bit Rate Performance

Access networks implementing either FWA or DRWA are simulated assuming $N = \{16, 32, 64\}$ ONUs, for the three population cases (cases 1, 2, and 3). For each value of N , several

numbers of wavelength channel pairs M (number of OLTs) are studied: (i) $N = 16$, $M = \{1, 2, 4, 8\}$; (ii) $N = 32$, $M = \{1, 2, 4, 8, 16\}$; and (iii) $N = 64$, $M = \{1, 2, 4, 8, 16, 32\}$.

The simulations address the impact of the fraction f of out-of-profile ONUs on the potential capacity improvement of DRWA with respect to FWA. The fraction of ONUs behaving as out-of-profile ones is randomly distributed among all ONUs according to a uniform distribution, so on average, a certain ONU- i has a probability f of behaving as an out-of-profile one. Using the randomly uniform distribution is a reasonable approach for estimating the typical bit rate performance of FWA when no further information about the distribution of out-of-profile ONUs among virtual PONs is available.

In the present paper, without loss of generality, the nominal bandwidth $\{C_d, C_u\}$ of each wavelength channel forming a pair is normalized to 1, independently of the symmetric or asymmetric nature of the uplink/downlink. The specific meaning of the nominal bandwidth is a matter of convenience, i.e., it may be based on an arbitrary measure, like the raw bit rate (all the bits transmitted over the physical layer, including payload data as well as protocol overhead), and the data transfer rate (the average useful data bit rate that is delivered to the application), etc. For example, for the downlink of an EPON, if the nominal bit rate is identified with the physical bit rate, then $C_u = 1.0$ conveniently summarizes 1.250 Gb/s, whereas if the nominal bit rate is taken as the maximum downlink throughput, some degradation must be introduced to take into account the data link layer protocol overhead and other factors like the number of connected ONUs, etc.

Figure 4 shows the mean relative bit rate per ONU that may be obtained with respect to f , for $N = 16$ ONUs. The performance of FWA versus DRWA is illustrated for the three population cases and for the number of wavelength channel pairs $M = \{1, 2, 4, 8\}$ —note that the performance results of a DRWA or FWA scheme with $M = 1$ are equal because, if there is only one wavelength pair, there is no room for reconfiguration. The relative bit rate per ONU is particularized for in-profile and out-of-profile ONUs; the values result from computing the average bit rate experienced when there is at least one busy ONU for each profile; additionally the bit rate per ONU when there is at least one busy ONU (regardless of its profile, denoted by “arbitrary ONU”) is also included.

During network planning, the data in Fig. 4 may be used to determine the number of wavelength channel pairs needed to achieve a specific percentage of the nominal bit rate. For example, for the population case 1 with FWA—Fig. 4(a)—if 20% of the ONUs are expected to behave as out-of-profile ones, in order to offer 80% of the OLT capacity to any ONU, $M = 4$ wavelength pairs are needed (whereas only $M = 2$ pairs are needed if DRWA is implemented—Fig. 4(b)). Furthermore, Fig. 4 may assist in calculating the bit rate performance for a varying value for the fraction of out-of-profile ONUs (as the composition of the population will be nonstationary); e.g., for the previous example with DRWA—Fig. 4(b)—if now 40% of ONUs are out-of-profile ones, the average capacity per ONU then drops to 70%.

As shown in Fig. 4, both assignment schemes improve their performance when the operator increases the number of wavelengths M . Similarly, the greater the number of out-of-profile ONUs, the less the available bit rate per

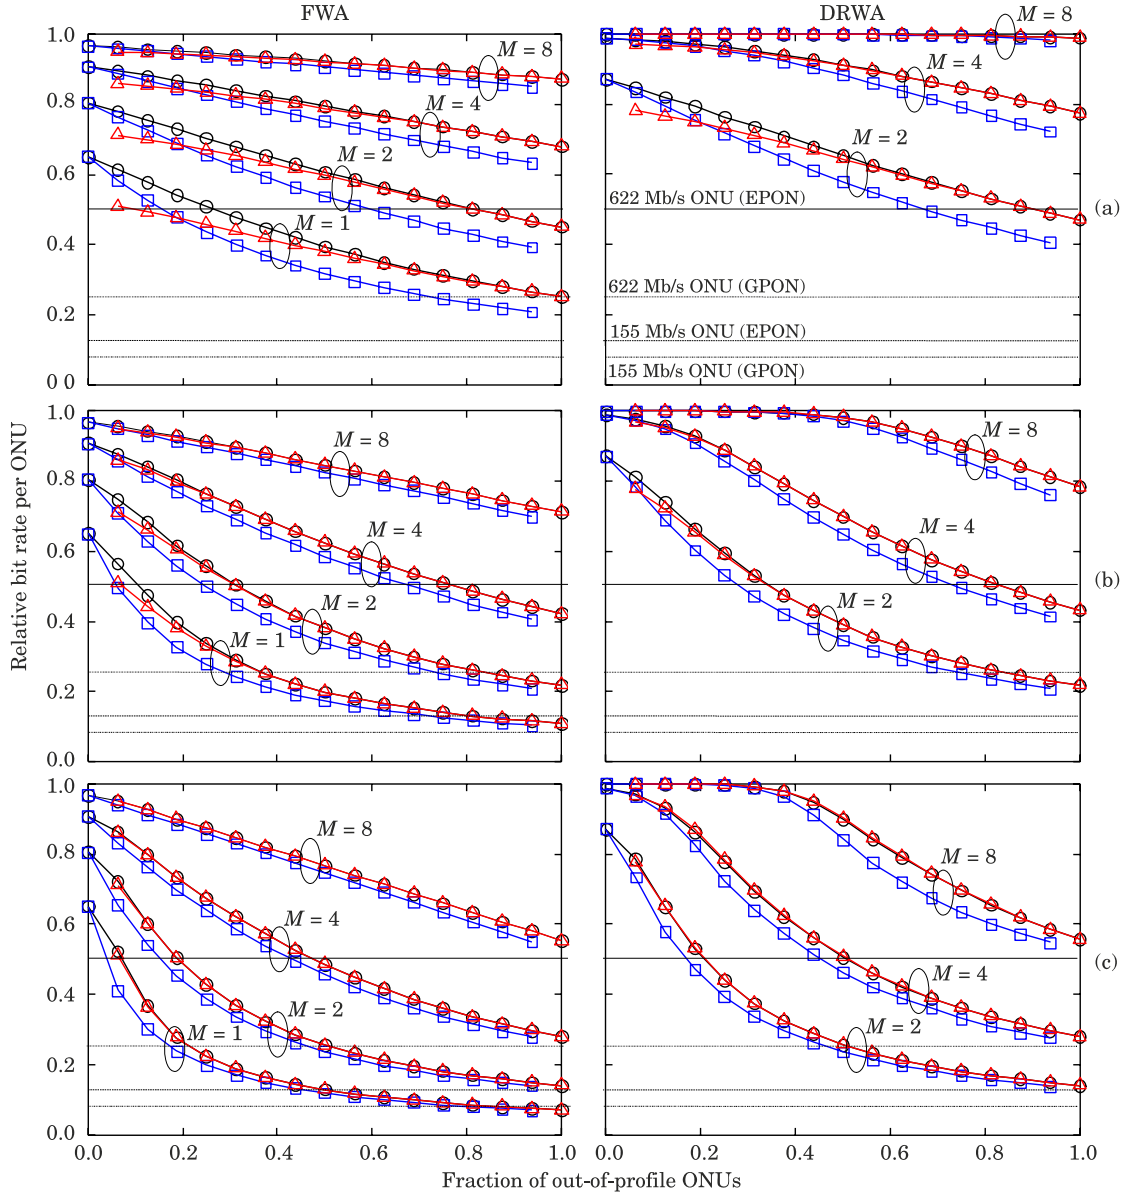


Fig. 4. (Color online) Average relative bit rates per user obtained for different numbers of wavelength channel pairs with respect to the fraction of out-of-profile users. FWA and DRWA are compared in three cases, which differ in the nature of out-of-profile users. Typical bit rates are included as a reference for EPON and GPON standards. Legend: [●] “arbitrary user” (see the text), [■] in-profile user, [▲] out-of-profile user. Row-wise: (a) case 1, (b) case 2, (c) case 3.

ONU—although note that from the operator viewpoint an unexpectedly high amount of out-of-profile ONUs corresponds to poor accuracy in predicting the users’ needs. Observe that for a concrete value of M , the DRWA scheme is always superior to its FWA counterpart. For example, if the operator desires to design an EPON that must guarantee a minimum bit rate of 622 Mb/s for an arbitrary ONU, in case 1, $M = 2$, with FWA up to 81% of out-of-profile ONUs may be supported, while DRWA supports up to 90%; in case 2, $M = 4$, with FWA up to 75% of out-of-profile ONUs are possible, while DRWA allows up to 83%.

As shown, the magnitude of the difference between DRWA and FWA depends on the potential for statistical multiplexing, being high when low aggregated load is likely when compared

with the amount of resources—i.e., when the offered loads of out-of-profile ONUs and in-profile ONUs are similar (case 1), if the fraction of out-of-profile ONUs is reasonable ($f < 0.4$), or if there are numerous wavelength pairs ($M = \{4, 8\}$). For example, note that in a GPON for case 3, $M = 8$ (a scenario where out-of-profile ONUs offer a significant load, but there are plenty of resources), by means of a DRWA the operator may serve a population which consists of up to 40% of out-of-profile ONUs with a performance acceptably close to the nominal bit rate (2.5 Mb/s), while its FWA counterpart degrades this value to 2.0 Mb/s. In contrast, an indication of network overload is the lack of statistical multiplexing that may be deduced when the bit rates obtainable with DRWA and FWA are similar. Since in those situations the number of wavelength channel

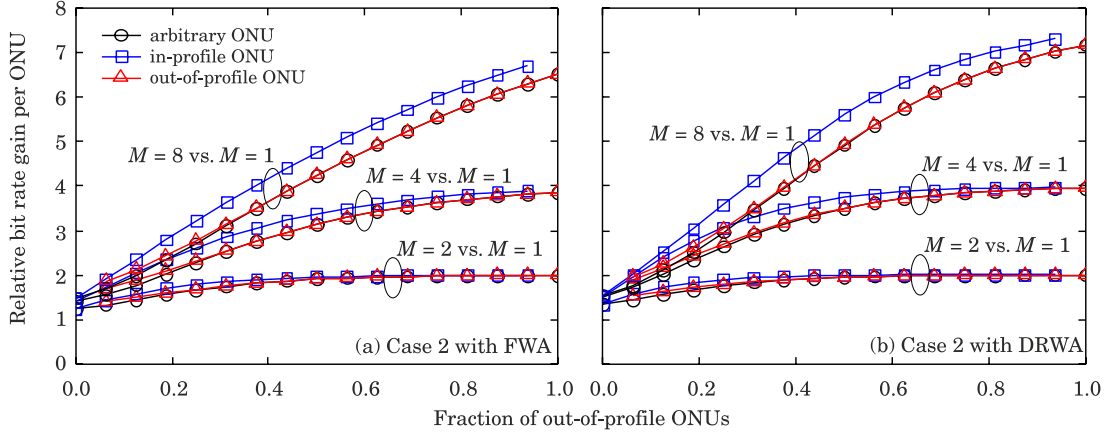


Fig. 5. (Color online) Average relative bit rate gains obtained for different numbers of wavelength channel pairs compared to having a single wavelength channel. FWA and DRWA are compared for case 2.

pairs is insufficient, the operator needs to add extra pairs or deploy additional infrastructure to do that.

Results for $N = \{32, 64\}$ (not shown due to space limitations) indicate that the general trend is that as N doubles, doubling the number of available wavelength channel pairs always increases the amount of bit rate improvement per ONU when using a DRWA scheme over the same scenario with FWA; additionally, the difference of the bit rate performance perceived by in-profile, out-of-profile and arbitrary ONUs is reduced until they are almost the same.

Figure 5 shows the mean bit rate gain—relative to the scenario with a single wavelength pair—experienced by ONUs with respect to the fraction of out-of-profile ONUs, for a network with $N = 16$ ONUs. The performance of FWA—Fig. 5(a)—and DRWA—Fig. 5(b)—is depicted for the population case 2 and a number of wavelength channel pairs $M = \{2, 4, 8\}$. As represented in Fig. 5, the performance gain of DRWA is superior to that of FWA, the difference being significant for high incidences of out-of-profile ONUs: e.g., for 40% of out-of-profile ONUs, in-profile ONUs get almost five times more capacity with DRWA, in comparison to the less than four times using FWA. Note also that, in general, in-profile ONUs—which could be a target group for being protected for bit rate by the network—obtain higher gains than out-of-profile ONUs. Nevertheless out-of-profile ONUs are not to be harmed because of bad operator planning if they aggregate rightful customers; thus reconfiguration should also address protection of individual paying customers which demand full use of their subscribed bandwidth.

Figure 6 shows the mean bit rate improvement of DRWA relative to FWA for a given number of wavelength pairs, with respect to the fraction of out-of-profile ONUs. Population case 1—Fig. 6(a)—and case 2—Fig. 6(b)—are depicted, and again for $N = 16$ ONUs. The first conclusion is that reconfigurability is helpful under all conditions (the bit rate gain is always higher than 1). However the specific magnitude of the improvement is a subject of further analysis. As shown, the curves exhibit global maxima, e.g., 17.99% gain for $M = 4$ when 18.75% of the ONUs are behaving as out-of-profile ones—aggregating flows of the Internet video streaming

category; see Fig. 6(b). Maxima occur at different values of f for different values M of wavelength pairs.

As a general rule, for a negligibly low presence ($f \approx 0$) of out-of-profile ONUs there is a value of M that maximizes the advantage of DRWA over FWA, i.e., lower values of M are not enough to produce a sizable surplus of capacity which could be usable by the DRWA, while higher values of M are difficult to exploit because of the low offered load of the predominant in-profile ONUs and the excess of capacity. For example, note case 1: DRWA with $M = 8$, while being better than FWA with $M = 4$ —see Figs. 6(a) and 4—is otherwise only barely able to improve the performance of both DRWA with $M = 4$ or FWA with $M = 8$ —Fig. 6(a). Nevertheless, as the amount of out-of-profile ONUs increases, the aggregated load rises, and while the previously appropriate value of M decreases its performance, larger values of M become effective, mitigating the impact of overload and obtaining bit rate gain even in the presence of an unexpectedly large fraction of out-of-profile ONUs. The non-intuitive behavior of the bit rate gain as f increases with M is a consequence of the DRWA scheme approach. When the value of f is low, the DRWA scheme, by transferring a busy ONU from an overcrowded OLT to an unpopulated OLT, is in fact isolating out-of-profile ONUs, assigning an exclusive wavelength channel pair that is shared with some in-profile ONUs. As f rises, there is a value of f_0 —e.g., $f_0 = 0.2$ in Fig. 6(b) with $M = 4$ —where the mean number of out-of-profile ONUs per wavelength channel pair is exactly 1 (the DRWA achieves peak gain) and values above f_0 cannot increase further the bit rate gain because more than one out-of-profile ONU will be irremediably sharing the same wavelength channel pair.

Note that for each value of M there is an upper bound for the range of values of f where in-profile ONUs will perceive more gain than out-of-profile ONUs, e.g., $f = \{0.200, 0.375, 0.650\}$ for $M = \{2, 4, 8\}$, as shown in Fig. 6(b), which for practical purposes may be used to define intervals of usefulness of DRWA for each value of M . Furthermore, the relative importance of the maximum depends directly on the offered load of the out-of-profile ONUs relative to the in-profile ONUs, e.g., when compared to $a_{in} = 0.1$ the maxima are more noticeable for $a_{out} = 0.6$ —Fig. 6(b)—than for $a_{out} = 0.3$ —Fig. 6(a); conversely,

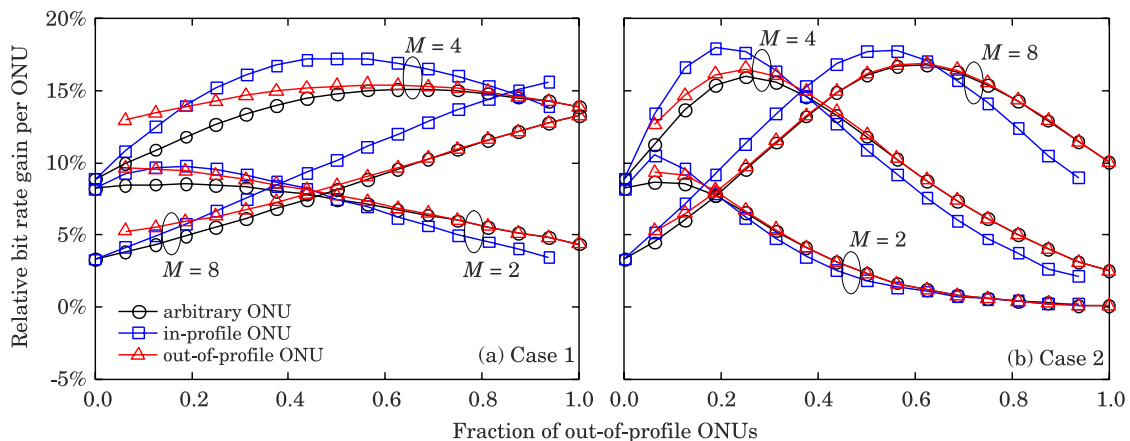


Fig. 6. (Color online) Average relative bit rate gains expressed as percentage obtained with DRWA compared to FWA for different numbers of wavelength channels. Cases 1 and 2 are compared.

in situations where overload is probable, mainly because either f or a_{out} (or both) have been underestimated, the gain obtained by DRWA is only competitive if a suitable number of wavelength channel pairs may be deployed with ease; e.g., case 2—Fig. 6(b)—when $f = 0.62$. If the operator undervalues the amount of necessary resources deploying $M = 2$ channel pairs, when using a DRWA scheme, an arbitrary ONU will obtain only a 2% gain over the same situation with FWA; however, if now a higher number ($M = 8$) is deployed, the gain amounts to 16.9%.

Figure 7 helps in identifying the trends for the relative bit rate for $N = \{16, 32, 64\}$ ONUs, for both DRWA and FWA. Population cases 1 and 3 are studied, and half the ONUs are out-of-profile ones. As shown, for a fixed M , the average bit rate decreases with N , as potentially more busy ONUs share a wavelength pair; conversely, for a fixed N , the average bit rate increases with M , because more resources are available. Note also that, as N increases, the difference between the bit rate perceived by in-profile and out-of-profile ONUs diminishes—e.g., it becomes nearly indistinguishable for $N = 64$. The reason for this is that when the value of N is substantial, the state dependent arrival process that models the ONUs tends to be a process with constant intensity. As a result, the performance obtained is the same regardless of the profile.

Figure 7 also illustrates that given a population case and a fraction of out-of-profile ONUs, there is an M that maximizes the bit rate gain obtainable with DRWA in comparison to FWA (i.e., causes maximum relative difference between the DRWA and FWA plots)—e.g., for Fig. 7(a) when $N = 64$, the value $M = 16$ is the one where the ratio between relative bit rates (0.96 for DRWA and 0.75 for FWA) achieves a maximum (28.5% of improvement). Moreover, the bit rate performance is more than proportional to the M/N ratio: rather than maintaining the bit rate, doubling the number of ONUs and OLTs (N and M , respectively) always increases this performance parameter. For example, in Fig. 7(b), for an arbitrary ONU, when $N = 16$ and $M = 8$, the DRWA/FWA gain is $0.90/0.77 = 17\%$; doubling the previous values ($N = 32$ and $M = 16$) the gain is $0.91/0.76 = 20\%$; doubling again ($N = 64$ and $M = 32$) the gain is $0.93/0.75 = 23\%$. This is the typical “trunking efficiency” effect that should be expected when scaling the network.

Up to now, it has been assumed that the behavior of the ONUs is randomly distributed according to a uniform distribution. However, the performance of FWA is sensitive to the distribution, so a similar performance cannot be expected if the fraction of out-of-profile ONUs follows a different distribution. Indeed, two distributions produce the most radically different performances: (i) when out-of-profile ONUs are exactly uniformly distributed among virtual PONs—which is a highly improbable situation—the performances of the FWA and DRWA schemes are practically identical; (ii) conversely, a maximally unbalanced situation appears when out-of-profile ONUs are distributed in such manner that all out-of-profile ONUs are clustered together in a static group of virtual PONs that serves them exclusively (note that this is a deterministic distribution) and then the performance of a FWA from the out-of-profile ONU’s viewpoint is the worst possible (and hence the bit rate gain offered by a DRWA scheme in such a situation will be the maximum possible). Note that the distribution of out-of-profile ONUs only affects the bit rate performance of a FWA scheme, because when a DRWA scheme is in use, load balancing is enforced and ONUs get the same bit rate regardless of the distribution of out-of-profile ONUs.

Figure 8 shows average values for the relative bit rate gain of DRWA with respect to FWA for the “typical” (randomly uniform) and “maximum” (maximally unbalanced) distributions for case 3.

As shown, under a maximally unbalanced distribution, out-of-profile ONUs must expect higher bit rate gains when DRWA is introduced; e.g., with $M = 8$ wavelengths—Fig. 8(b)—when the fraction of out-of-profile ONUs is 31.25%, for a randomly distributed situation the bit rate gain is 17.30% while the maximum potential gain is 78%; even more significant is the bit rate when $M = 4$ —Fig. 8(a)—e.g., when the fraction of out-of-profile ONUs is 25%, for a randomly distributed scenario the gain is 12% while the maximum achievable gain is more than 175%. Note that in-profile ONUs obtain relative bit rate gains that are near or below 1 because a situation where out-of-profile ONUs are grouped at certain virtual PONs under a FWA scheme is advantageous for in-profile ONUs, and a DRWA will only degrade the relative bit rate. Finally, the manifest zigzag pattern is due to the fact that the number of

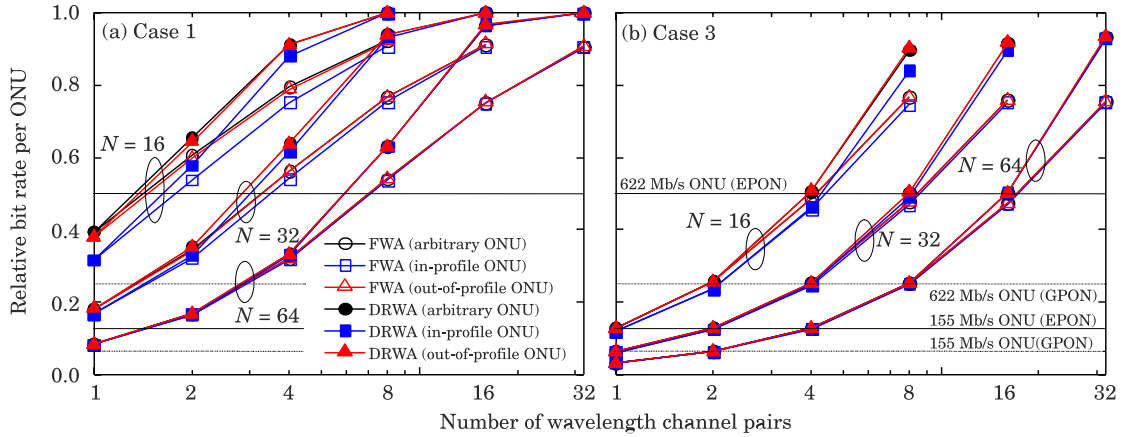


Fig. 7. (Color online) Average relative bit rates obtained with FWA and DRWA when the fraction of out-of-profile users equals $f = 0.5$ with respect to the number of wavelength channels. Cases 1 and 3 are compared. Typical bit rates are included as a reference for EPON and GPON standards.

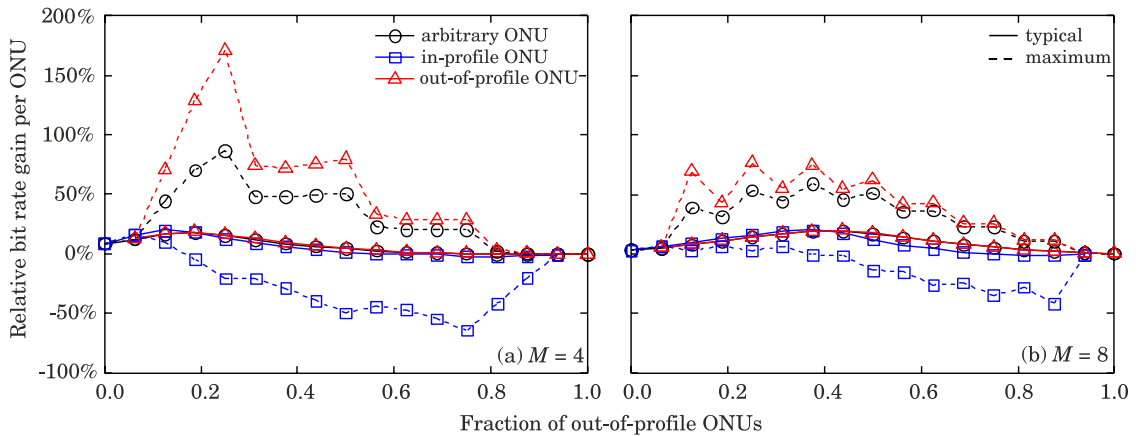


Fig. 8. (Color online) Average and maximum relative bit rate gains obtained with DRWA compared to those with FWA. The results correspond to case 3 for $M = \{4, 8\}$ wavelength channel pairs.

out-of-profile ONUs is not always a multiple of the number of ONUs per virtual PON, $k = M/N$. Actually, the sharp drops in the curve happen when the number of out-of-profile ONUs is such that, without one of them, the number of out-of-profile ONUs would be a multiple of k (i.e., one virtual PON is serving a group of in-profile ONUs and one out-of-profile ONU).

B. Uplink Frame Delay and Frame Loss

The uplink at the lower burst level (Fig. 2) is simulated considering an uplink data rate of $C_u = 1.25$ Gb/s and 1.244 Gb/s for EPON and GPON standards, respectively. A Gigabit Ethernet link ($c_u = 1.00$ Gb/s) connects the subscribers to the ONU. In the uplink, WDM extended [15,16] IPACT [10] for MPCP with “fixed” discipline (adapted to the fixed polling cycle nature of the uplink transmission [11]) is implemented for GPON, while “limited” and “gated” grant scheduling services [10] are implemented for EPON. Values for the remaining parameters are established according to [10]. The buffer size of the ONUs is $Q = 10$ MByte. For fixed

and limited disciplines, the fixed and maximum polling cycle time for IPACT implementations is $T_{\max} = 10t_{\max_size}$. Pareto distributions for the ON/OFF periods are assumed, with parameters $\alpha_{\text{ON}} = 1.2, \alpha_{\text{OFF}} = 1.4$ and $x_{m\text{ON}} = 1$, and compared to exponential distributions with the same expected values for the average duration of the ON/OFF periods. All results for Poisson/(Pareto) traffic are presented with a 99%/(90%) confidence interval with a 1%/(10%) confidence half-width using the batch means process provided by Möbius (the necessarily lower values for Pareto derive from the higher variance inherent to a long range dependent probability distribution).

Figures 9(a) and 9(b) show the mean frame delay and frame loss, respectively, with respect to the relative bit rate per ONU. As shown, the three scheduling disciplines are compared against the two possible burst distributions. Note that a small value of the relative bit rate per ONU means that a greater number of ONUs are trying to transmit their frames at the same time, and vice versa. Also note that the maximum average bit rate of 0.8 relative to the uplink rate of 1.25 Gb/s arises because the Gigabit Ethernet link




Fig. 9. (Color online) Mean frame delay (a) and mean frame loss (b) with respect to the relative bit rate per ONU for Pareto and exponentially distributed durations of the ON/OFF periods. A fixed grant discipline implementation of IPACT [14] is adapted for GPON, while limited and gated disciplines are assumed for EPON.

is the limiting bottleneck (i.e., $1.00/1.25 = 0.8$). As expected, exponential distributions yield optimistic results for the mean frame delay (almost always one order of magnitude below Pareto estimations and sometimes up to two orders), while the reported frame loss for the exponential case is negligible. The gated discipline has the worst performance within a wider range of relative bit rates because it is not limited by a value of T_{\max} and allows transmitting the entire buffer contents in one batch, which increases cycle times and, thus, latency. The fixed polling cycle establishes a T_{\max} , but ignores the amount of data pending transmission in the ONU, and therefore slightly improves on the previous results. In comparison, the limited discipline, which uses cycle times that adapt to the amount of queued data, performs significantly better. Figure 9, finally, shows that the improvement in the relative bit rate achieved by a DRWA scheme always implies a reduction of frame loss and frame delay (in the last case, more significant when the relative bit rate rises above 0.5).

V. DISCUSSION

From the load balancing viewpoint, DRWA and DBA complement each other. While a well-designed DBA scheme is particularly successful at arbitrating the access to the wavelength channel resource (but only when the aggregated offered load of all ONUs is within a reasonable interval), an optimum DRWA organizes ONUs such that those of them serving users with high bit rate demands are effectively isolated at a wavelength channel and those of them serving users with low offered loads are aggregated together for the benefit of the DBA scheme.

Multi-dwelling units in a residential environment have been considered in order to assist in the interpretation of the results. In such scenarios the traffic per ONU is more aggregated and therefore less subject to variability, which decreases the performance of DBA; however, since each ONU could be serving radically different groups of subscribers, traffic requirements may be geographically irregular and

difficult to predict. Uncertainty in the traffic demands leads to characterization inaccuracies during network planning; thus wavelength reconfigurability opens the possibility of postponing such studies in favor of addressing the issue satisfactorily during network operation.

Aside from load balancing, introducing wavelength reconfigurability is also valuable as a way of adding or reserving extra capacity (on demand) for certain high-profile users (e.g., corporate subscribers using high-end applications) at identified locations of the access network. Actually, the scenario described by the maximally unbalanced case of Fig. 8 could be read as a situation where the location of the out-of-profile ONUs (providing connectivity to legitimate and very valuable subscribers) is predictable, and hence significant performance gains should be expected (e.g., in order to meet the exceptional peak demands in mobile cellular networks that increasingly appear at sporting and musical events, as well as major festivals and celebrations).

VI. CONCLUSIONS

Since the utilization of a reconfigurable wavelength assignment scheme in a WDM-TDMA-PON leads to significant economic and complexity costs, it is essential to substantiate the deployment of such schemes in comparison to an inexpensive fixed wavelength assignment scheme. Load balancing is one of the key motivations for rearranging wavelength channel pairs.

In this paper, a procedure for addressing this issue at the network modeling stage, during the planning by the operator, is introduced and then evaluated in a realistic, nonstationary approach for the traffic conditions during the operation of the network. The results of the assessment by means of simulation (typical and maximum obtainable bit rate expressed in terms relative to the nominal capacity of a wavelength channel) show that the benefits achieved in terms of network performance depend on the number of wavelength channels as well as the fraction of out-of-profile ONUs and their behavior relative to in-profile ONUs.

The conclusions of the present study help us to define the conditions for which a dynamic reconfigurable wavelength assignment scheme excels at exploiting the existing network resources in the most cost-effective way. DRWA outperforms FWA, but the magnitude of the improvement depends on the magnitude of the aggregated load with respect to the number of wavelength channel pairs and the number of ONUs—being substantial only if the fraction f of ONUs that offers more traffic than expected does not exceed a reasonable margin (e.g., $f < 0.4$). Also, for each situation there is a specific number of wavelength channel pairs that maximizes the benefits of DRWA over FWA, balancing the benefits of a useful excess of capacity with the potential for statistical multiplexing of the superposition of ONUs. A threshold for f exists also, which defines two ranges of values of f : values where in-profile ONUs perceive more gain than out-of-profile ONUs and vice versa. As the number of wavelength channel pairs increases, the trend leads to a better than proportional improvement in bit rate performance, reducing the difference in the bit rate experienced by in-profile and out-of-profile ONUs. Finally, the performance of a FWA is sensitive to the distribution of out-of-profile ONUs among virtual PONs; therefore in situations where an unbalanced distribution is stable, out-of-profile ONUs will obtain extra bit rate gains when a DRWA scheme is introduced. Remarkably, although a reconfigurable scheme in a modest deployment with 16 ONUs obtains typical peak gains of 10%–20%, maximum gains of up to 175% are actually possible in practical scenarios with exceptionally unbalanced traffic demands.

Finally, frame delay and frame loss for uplink transmission are studied for Pareto and exponentially distributed lengths of the ON/OFF transmission periods in EPON and GPON scenarios, providing a linkage to the bit rate performance findings previously presented. The results show that for Pareto traffic, a DRWA scheme always provides noticeable frame loss improvements, while reduced frame delay improvements are more significant when the relative bit rate per ONU is higher than 0.5.

ACKNOWLEDGMENTS

This work was supported by the EC 7th Framework Program: Architectures for Flexible Photonic Home and Access networks (ALPHA), under contract ICT CP-IP 212 352, from the Generalitat of Valencia under contract ACOMP/2010/196. The authors thank the Performability Engineering Research Group (PERFORM) at the University of Illinois at Urbana-Champaign for developing the software tool Möbius.

REFERENCES

- [1] 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.*, vol. 23, no. 1, pp. 277–286, Jan. 2005.
- [2] M. Maier, M. Herzog, and M. Reisslein, "STARGATE: the next evolutionary step toward unleashing the potential of WDM EPONs," *IEEE Commun. Mag.*, vol. 45, no. 5, pp. 50–56, May 2007.
- [3] P. J. Urban, B. Huiszoon, R. Roy, M. M. de Laat, F. M. Huijskens, E. J. Klein, G. D. Khoe, A. M. J. Koonen, and H. de Waardt, "High-bit-rate dynamically reconfigurable WDM-TDM access

- network," *J. Opt. Commun. Netw.*, vol. 1, no. 2, pp. A143–A151, June 2009.
- [4] R. Glatty, P. Guignard, and P. Chanclou, "Fair resource distribution within the flexible WDMA/TDMA optical access network based on GPON infrastructure," *J. Opt. Commun. Netw.*, vol. 1, no. 2, pp. 17–24, June 2009.
- [5] R. Roy, G. Manhoudt, and W. van Etten, "Optical-router-based dynamically reconfigurable photonic access network," *J. Opt. Netw.*, vol. 8, no. 1, pp. 51–76, Jan. 2009.
- [6] T. Koonen, K. Steenbergen, F. Janssen, and J. Wellen, "Flexibly reconfigurable fiber-wireless network using wavelength routing techniques: the ACTS project AC349 PRISMA," *Photonic Network Commun.*, vol. 3, no. 3, pp. 297–306, July 2001.
- [7] J. Homa and K. Bala, "ROADM architectures and their enabling WSS technology," *IEEE Commun. Mag.*, vol. 46, no. 7, pp. 150–154, June 2008.
- [8] T. A. Strasser and J. Taylor, "ROADMS unlock the edge of the network," *IEEE Commun. Mag.*, vol. 46, no. 7, pp. 146–149, July 2008.
- [9] W. E. Leland, M. S. Taqqu, W. Willinger, and D. V. Wilson, "On the self-similar nature of Ethernet traffic (extended version)," *IEEE/ACM Trans. Netw.*, vol. 2, no. 1, pp. 1–15, Feb. 1993.
- [10] G. Kramer, B. Mukherjee, and G. Pesavento, "Interleaved polling with adaptive cycle time (IPACT): a dynamic bandwidth distribution scheme in an optical access network," *Photonic Network Commun.*, vol. 4, no. 1, pp. 89–107, Jan. 2002.
- [11] B. Skubic, J. Chen, J. Ahmed, L. Wosinka, and B. Mukherjee, "A comparison of dynamic bandwidth allocation for EPON, GPON, and next-generation TDM PON," *IEEE Commun. Mag.*, vol. 47, no. 3, pp. 40–48, Mar. 2009.
- [12] G. I. Papadimitriou and A. S. Pomportsis, "Self-adaptive TDMA protocols for WDM star networks: a learning-automata-based approach," *IEEE Photon. Technol. Lett.*, vol. 11, no. 10, pp. 1322–1324, 1999.
- [13] C. Linardakis, H. C. Leligou, A. Stavdas, and J. D. Angelopoulos, "Using explicit reservations to arbitrate access to a metropolitan system of slotted interconnected rings combining TDMA and WDMA," *J. Lightwave Technol.*, vol. 23, no. 4, pp. 1576–1585, 2005.
- [14] K. Kanonakis and T. Tomkos, "Improving the efficiency of on-line upstream scheduling and wavelength assignment in hybrid WDM/TDMA EPON networks," *IEEE J. Sel. Areas Commun.*, vol. 28, no. 6, pp. 838–848, 2010.
- [15] M. McGarry, M. Maier, and M. Reisslein, "An evolutionary WDM upgrade for EPONs." Arizona State University, Tech. Report, 2005. Available: <http://mre.faculty.asu.edu/EPONupgrade.pdf>.
- [16] M. McGarry, M. Maier, and M. Reisslein, "WDM Ethernet passive optical networks (EPONs)," *IEEE Commun. Mag.*, vol. 44, no. 2, pp. 15–22, Feb. 2006.
- [17] A. R. Dhaini, C. M. Assi, M. Maier, and A. Shami, "Dynamic wavelength and bandwidth allocation in hybrid TDM/WDM EPON networks," *J. Lightwave Technol.*, vol. 25, no. 1, pp. 277–286, Jan. 2007.
- [18] A. R. Dhaini, C. M. Assi, and A. Shami, "Dynamic bandwidth allocation schemes in hybrid TDM/WDM passive optical networks," in *3rd IEEE CCNC*, 2006, vol. 1, pp. 30–34.
- [19] X. Ma and G. S. Kuo, "Optical switching technology comparison: optical MEMS vs. other technologies," *IEEE Commun. Mag.*, vol. 41, no. 11, pp. S16–S23, Nov. 2003.
- [20] L. Kleinrock, *Queueing Systems*. Wiley-Interscience, 1975, Vol. 1, Chap. 2.
- [21] *Möbius Manual*, University of Illinois at Urbana-Champaign, 2009. Available: www.mobius.illinois.edu/manual/MobiusManual.pdf.