A Highly Flexible and Efficient Passive Optical Network Employing Dynamic Wavelength Allocation

Yu-Li Hsueh, *Student Member, IEEE*, Matthew S. Rogge, *Student Member, IEEE*,
Shu Yamamoto, *Senior Member, IEEE*, and Leonid G. Kazovsky, *Fellow, IEEE*

***Abstract*—A novel and high-performance passive optical network (PON), the SUCCESS-DWA PON, employs dynamic wavelength allocation to provide bandwidth sharing across multiple physical PONs. In the downstream, tunable lasers, an arrayed waveguide grating, and coarse/fine filtering combine to create a flexible new optical access solution. In the upstream, several distributed and centralized schemes are proposed and investigated. The network performance is compared to conventional TDM-PONs under different traffic models, including the self-similar traffic model and the transaction-oriented model. Broadcast support and deployment issues are addressed. The network's excellent scalability can bridge the gap between conventional TDM-PONs and WDM-PONs. The powerful architecture is a promising candidate for next generation optical access networks.**

***Index Terms*—Dynamic wavelength allocation (DWA), optical access network.**

I. INTRODUCTION

THE exponential growth of Internet traffic volume directly relates to the increase in high-speed connections to end users. Still, the demand for ever higher data rates remains strong. In the backbone networks, capacities greater than 1 Tb/s over a single fiber have been achieved by means of wavelength-division-multiplexing (WDM) technology with hundreds of channels. The emergence of low-cost and high-speed Ethernet-based local area networks (LANs) is accelerating the demands for high-speed connections as well. To break the bottleneck between the LAN and the ultrahigh-capacity backbone networks, cost-effective and high-performance access solutions are desirable. Passive optical networks (PONs) have been identified as one of the promising access solutions, as shown by the many research and standardization efforts focusing on this area.

Currently, the broadly accepted optical access solution is the time-division multiplexed (TDM) PON. TDM-PONs combine the high capacity afforded by optical fiber with the low installation and maintenance cost of a passive infrastructure. The broadcast transmission in the downstream and the time-sharing transmissions in the upstream limit the bandwidth of individual users, but the resulting low transceiver count minimizes the cost enough to justify the tradeoff.

One straightforward approach to increase the capacity is to assign a set of wavelengths to each user for down/upstream transmissions, which leads to the WDM-PON [1]–[3]. WDM-PONs create point-to-point links between the central office (CO) and each user, so no sharing is needed. However, high performance is not without high cost—whenever the user shuts down his connection, the corresponding transceiver in the CO is idle and cannot support other users, in which case the network resource is left unused. WDM-PONs are high-performance but luxurious access solutions today. However, as bandwidth demands increase and optical component costs decrease, WDM-PONs will become more practical and significant.

Unfortunately, the migration path from TDM to WDM is not trivial. A full WDM-PON has dedicated transmitters/receivers for each end user as well as wavelength-routing devices in the infrastructure. The cost required to add these new wavelengths and modify the infrastructure to support WDM constitute a major roadblock to adoption. Ideal solutions provide cost-effective and smooth service upgrades with minimal impact on the existing TDM infrastructure.

In [4] we initiated the concept of a novel optical access network architecture, the SUCCESS-DWA PON, developed under the Stanford University Access networking project at the Photonics and Networking Research Laboratory (PNRL). It employs dynamic wavelength allocation (DWA) to achieve a flexible, cost-effective, and high-performance PON architecture. In this paper, we describe the architecture in detail, investigate the performance, and comment on several implementation issues. One significant feature of the SUCCESS-DWA PON is that its excellent scalability can easily bridge the large gap between TDM and WDM PONs. In addition, the architecture provides excellent cost efficiency and network performance by sharing bandwidth across multiple physical PONs. Existing arbitrary field-deployed PON infrastructures remain in tact when brought together into a SUCCESS-DWA PON, making the architecture an ideal candidate for upgrading existing PONs. The inherent flexibility and high performance allow the architecture to serve equally well as high-end access networks serving large businesses and campuses.

The material and contributions of this paper are organized as follows: Section II gives the description of the proposed network; analysis and performance comparisons are shown in Section III; the support for broadcast, the wavelength plan, and implementation issues are discussed in Section IV; the conclusion is drawn in Section V.

Manuscript received December 16, 2003; revised September 21, 2004.

Y.-L. Hsueh, M. S. Rogge, and L. G. Kazovsky are with the Electrical Engineering Department, Stanford University, Stanford, CA 94305 USA.

S. Yamamoto is with KDDI R&D Laboratories, Inc., Saitama 356-8502, Japan.

Digital Object Identifier 10.1109/JLT.2004.838811

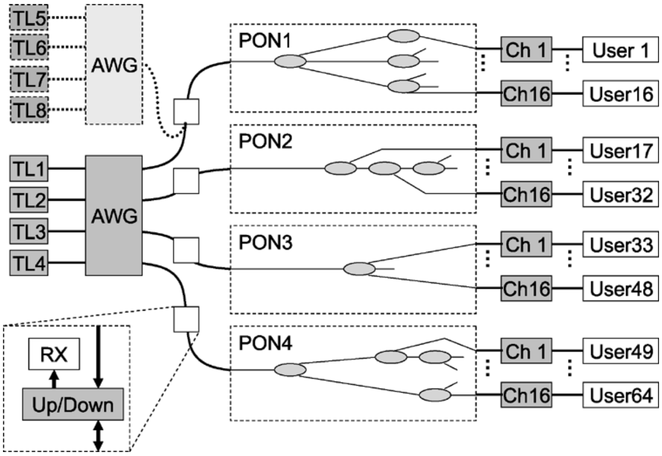

Fig. 1. Basic architecture of SUCCESS-DWA PON in the downstream. The grayed-out AWG and TLs demonstrate the system upgrade path.

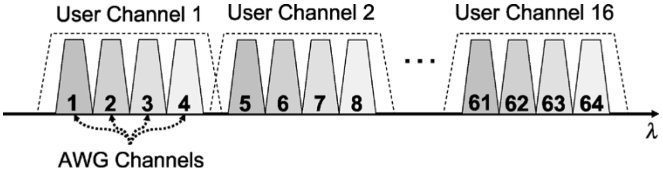

Fig. 2. Wavelength bands for AWG channels and thin-film WDM filter channels.

II. Network Architecture

A. Downstream

Fig. 1 shows the network architecture. Tunable lasers (TLs), the arrayed waveguide grating (AWG), and thin-film WDM filters (Ch1–Ch16 in Fig. 1) constitute the key parts of this network. TLs and the AWG reside in the CO, while the WDM filters are within the Optical Network Units (ONUs). The AWG which couples the TLs together is cyclic or near-cyclic. Notice that the field-deployed infrastructure is left untouched. This illustrative network includes four TLs and sixty-four end users located across four physical PONs. The architecture can be extended in numerous ways, e.g., to encompass more physical PONs or to provide more TLs per PON, which is addressed later. The dotted PON boxes correspond to arbitrary passive fiber distribution networks of existing PONs. Traditional PONs may be physically deployed with different distribution techniques, including multistage splitting (e.g. PON1 in Fig. 1), bus tapping (e.g., PON2), single wide-band splitters (e.g., PON3), or a hybrid combination (e.g., PON4). Regardless of the physical configuration, the design requires only that all wavelengths from the optical line terminal (OLT) reach all ONUs. The upstream traffic is separated from the downstream traffic with a wide-band WDM filter between the AWG and the PON.

In the basic architecture, each ONU within a single PON contains a unique fixed-wavelength filter and a burst-mode receiver. The key lies in the fact that the passband of the ONU filter encompasses the free spectral range of the AWG. For example, 200 GHz ONU filters would work with a 4×4 cyclic 50 GHz AWG. The relative filter shapes are illustrated in Fig. 2.

The seemingly complicated architecture provides a very powerful capability—any tunable laser can individually address any ONU across separate physical PONs at any given time. If TL1 were to tune his laser through consecutive AWG channels, the AWG would route channels $\{1,5,9,\ldots\}$ toward PON1, assuming channels $\{1,5,9,\ldots\}$ are the "straight through" channels of the AWG. Similarly, channels $\{2,6,10,\ldots\}$ would be routed to PON2, and likewise for PON3 and PON4. If a TL wishes to communicate with a user on a particular PON, he must determine what wavelength (a) falls in the passband of the user and (b) exits the AWG destined for the correct PON—for each TL-user combination, there exists exactly one wavelength for which both (a) and (b) are satisfied.

For example, consider TL1 and User 18 (user channel 2, PON2). If TL1 wishes to communicate with User 18, he compares the AWG channels acceptable by User 18, which are $\{5,6,7,8\}$, and those that reach PON2 from TL1, which are $\{2,6,10,\ldots\}$. The common wavelength in the two sets determines that TL1 should transmit on AWG channel 6. Generally, for a TL numbered L communicating with the user channel number U, the corresponding AWG channel can be expressed as

$$\lambda = ((U-1) \bmod U_{\mathrm{MAX}}) \times L_{\mathrm{MAX}} - \left(L - 1 - \left\lfloor \frac{(U-1)}{U_{\mathrm{MAX}}} \right\rfloor\right) \bmod L_{\mathrm{MAX}} + 1 \qquad (1)$$

in which U_{MAX} and L_{MAX} denote the total number of user channels, and the number of input AWG ports (the maximum number of TLs which can be coupled via the AWG), respectively.

Since each ONU contains only one photodetector (PD), two TLs must not access the same ONU simultaneously and a suitable media access control protocol is necessary. Notice that the AWG *does allow* all TLs to simultaneously transmit on the same wavelength. In this case, each TL will be routed to a different PON. For each downstream frame, the TLs tune to appropriate wavelengths and transmit the data to corresponding end users. The transmission durations for end users are globally managed to achieve optimal performance. This is very useful especially when accommodating very bursty Internet traffic. All TLs share the load, shifting bandwidth back and forth across the separate physical PONs as necessary—we call this technique dynamic wavelength allocation.

To illustrate the flexibility of this architecture, compare an initial deployment of four TDM PONs to a SUCCESS-DWA PON that spans four physical PONs. The first several subscribers would likely occupy physical locations across more than one PON, and in the worst case they may be spread across all four PONs. In the four TDM PON case, then, all four OLTs (lasers) must be activated, despite the fact that some OLTs may only be serving a few subscribers. With the SUCCESS-DWA PON, on the other hand, only one TL and AWG are initially added to the central office, and the subscribers across multiple PONs are all serviced by the single TL. As demand grows, additional TLs can be added to the AWG. When the subscription rate is high enough, the two scenarios seem to converge—both have four transmitters serving all subscribers. However, the SUCCESS-DWA PON enjoys the benefit of statistical multiplexing over a larger customer base, so its performance will exceed that of the four TDM PONs.

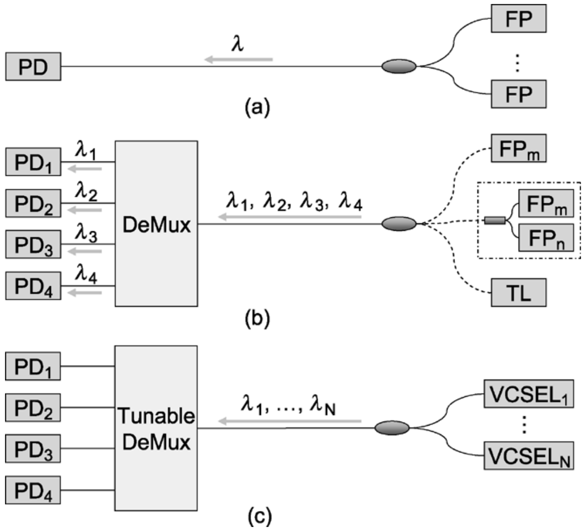

Fig. 3. Upstream scenarios: (a) conventional TDM-PON; (b) SUCCESS-DWA with distributed upstream schemes; and (c) SUCCESS-DWA with centralized upstream schemes.

When demand dictates, the SUCCESS-DWA PON can be scaled far beyond the conventional TDM PON. The grayed-out AWG in Fig. 1 illustrates the concept. Additional AWGs can be added to shift from four TLs serving four PONs to four TLs serving each single PON. If 8×8 AWGs are utilized, sixteen end users can be served by eight TLs, which results in a very high-performance network close to a WDM-PON. The excellent scalability provides a smooth and graceful upgrade from a TDM-PON all the way toward a WDM PON. The step-by-step system upgrade easily tracks user demands, and the initial overhead can be even lower than that of conventional TDM-PONs. In addition, the TLs provide protection for each other, maintaining service to all physical PONs in the event of a failure.

B. Upstream

Depending on the performance requirements, there is a wide range of possible scenarios for upstream transmission. In this subsection, several different upstream schemes are investigated and reported.

One inexpensive upstream scenario requires only a fixed-wavelength transmitter in each ONU—a cost-effective 1310 nm Fabry–Pérot (FP) laser can serve as the light source. Between the AWG and the physical PON, the upstream wavelength is extracted and a single PD terminates all the traffic for that PON. Within each physical PON, then, bandwidth is shared amongst all users on that PON—this is precisely the upstream scheme in a traditional TDM-PON, as illustrated in Fig. 3(a). The performance is limited, of course, but the cost is minimized. In [5], however, it was reported that the measured upstream traffic rates were highly related to the downstream traffic rates. Users enjoying high-speed downloads also expect a commensurate data rate in the upstream. This suggests that a high-performance upstream scheme is appropriate.

Ideally, the concept of DWA could also be employed in the upstream, allowing similar features in terms of cost, scalability, and performance. To realize DWA in the upstream, there must be tunability in the network. Fig. 3(b) shows the scenario with tunable devices in the ONU. In this example, a demultiplexer (DeMux) in the OLT routes the upstream wavelengths to the four corresponding PDs. The light source at each ONU can be i) a fixed-wavelength FP; ii) two or more FPs at different wavelengths; or iii) a TL. For i), the ONU is virtually grouped with other ONUs at the same wavelength, and shares the same PD in the time domain. For ii), the ONU can choose one of the FPs for transmission. For iii), the ONU can use any of the upstream wavelengths.

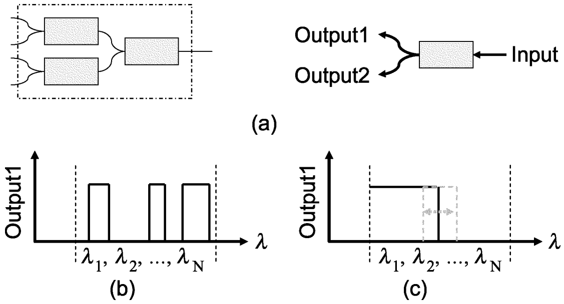

Fig. 4. Tunable demultiplexer: (a) cascade implementation composed of three single-stage 1×2 wavelength demultiplexers with complementary outputs; (b) arbitrary wavelength demultiplexing; and (c) consecutive wavelength demultiplexing.

The tunablility can also be in the OLT. In Fig. 3(c), the DeMux at the OLT is tunable, which may potentially be realized by MEMS [6], tunable thin-film filters [7], tunable fiber couplers [8], or other technologies. Each end user has a specific upstream wavelength λ_{m}, and cost-effective VCSELs would be ideal light sources.

Conceptually, a 1×4 optical DeMux could be realized by cascading 1×2 DeMux as shown in Fig. 4(a). The function of the 1×2 DeMux is partitioning the incoming wavelengths into two output groups, such that the two output ports would have complementary output spectra. Tunable DeMux with different wavelength grouping abilities could affect the network performance. Fig. 4(b) illustrates the spectrum of a tunable DeMux with arbitrary wavelength grouping ability, while Fig. 4(c) illustrates that with consecutive wavelength grouping ability, which shows a relatively easier but limited wavelength demultiplexing function.

Based on the illustrated setups in Fig. 3, we investigate five upstream schemes and evaluate each in terms of cost, performance, and scalability.

Scheme A: Fixed grouping of end users with FPs at several different wavelengths.

Scheme B: Some users are equipped with dual FPs at different wavelengths, and the rest are equipped with a single FP.

Scheme C: Some users are equipped with TLs, and the others are equipped with a single FP.

Scheme D: A tunable DeMux in the OLT that performs consecutive wavelength demultiplexing.

Scheme E: A tunable DeMux in the OLT that performs arbitrary wavelength demultiplexing.

For Scheme A, in the case of 16 users and four wavelengths, the users are divided into four groups. Users within a group transmit on the same upstream wavelength and time-share the same PD in the OLT. This simple grouping among users by

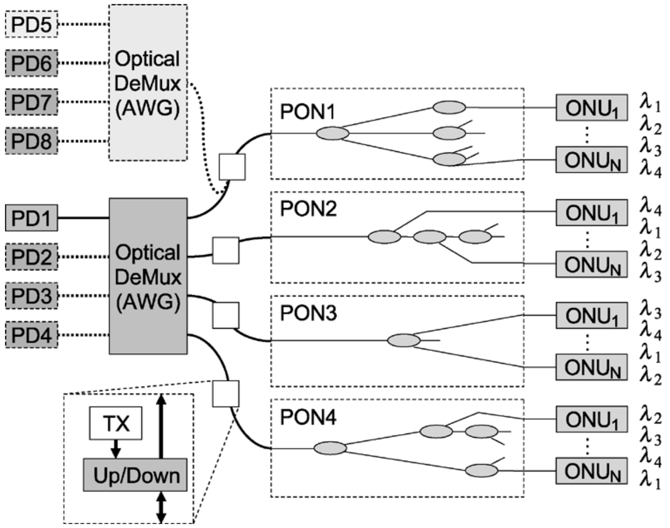

Fig. 5. The basic SUCCESS-DWA upstream architecture for users across multiple PONs. The grayed-out PDs and AWG demonstrate the system upgrade path.

wavelengths can be viewed as a straightforward upgrade from the conventional TDM-PON. It is static in wavelength allocation and provides a baseline for comparison with the DWA schemes. Schemes B and C are distributed DWA schemes, in which the tunability is spread across the ONUs. The distributed schemes allow more flexibility since the deployment of tunable devices can be judged by user demands. On the other hand, Schemes D and E are centralized DWA schemes, in which the tunability resides in the OLT.

Qualitatively, full tunability results in the best performance. Scheme E exhibits full-tunability, but at the cost of an expensive, centralized tunable DeMux. Scheme C can provide equal performance if all users are equipped with TLs. While the centralized schemes subject all ONUs to the high cost of the tunable device, the distributed schemes require only those ONUs which demand high performance to be upgraded. The added design flexibility of distributed schemes makes them more preferable.

Ideally, the upstream scenario should have a strong and smooth migration path similar to the downstream scheme. Initially, each user is equipped with a fixed-wavelength transmitter that corresponds to the upstream group to which he has been assigned. If the groups are created carefully, one can still achieve scalability similar to the downstream SUCCESS-DWA scheme. A complete upstream SUCCESS-DWA PON architecture is shown in Fig. 5. Four physical PONs are connected to the OLT, and a 4×4 AWG functions as the DeMux in Fig. 3 which is responsible for routing the incoming wavelengths to the corresponding PDs. It is worth noting that this AWG does not require cyclic features, therefore a thin-film based AWG could be adopted, and lower insertion losses and costs can be expected. Additionally, upstream and downstream AWGs pass completely different wavelengths and require different channel spacings. Therefore, separate AWGs are necessary for up/downstream. More details about the wavelength plan are described in Section IV.

Similar to the downstream, only one PD and its corresponding receiver module are activated in the initial deployment. To cover all end users located on different physical PONs, the first several subscribers from different PONs are assigned different upstream wavelengths, i.e., the first several subscribers from PON1 are assigned λ_1, the first several subscribers from PON2 are assigned λ_4, and so on, as shown in Fig. 5. When the number of users increases, a second PD and its corresponding receiver module can be installed in the OLT. For new subscribers to be served by the second PD, they are assigned λ_2 on PON1, λ_1 on PON2, and so on, as in Fig. 5. In general, with a $K \times K$ AWG, the first assigned wavelength for subscribers from the Pth PON is $\lambda = (1 - P) \bmod K + 1$.

This architecture can scale far beyond its TDM-PON counterpart, which is illustrated by the grayed-out AWG in Fig. 5. When demand dictates, similar AWGs can be added to shift from four PDs serving four PONs to four PDs serving each PON. For the bandwidth-demanding users, tunable devices can replace the fixed-wavelength transmitters, as in Schemes B and C. Unlike the SUCCESS-DWA scheme in the downstream which makes use of fast TLs with $\sim$10 ns tuning times to achieve DWA on the order of μs, the upstream DWA schemes perform relatively slower wavelength reallocations on the order of ms, due to the more involved communications between the OLT and ONUs for the upstream transmission scheduling.

III. ANALYSIS AND RESULTS

In this section, we compare the performance of the SUCCESS-DWA PON with the aggregate performance of a group of conventional TDM PONs. Together, the conventional TDM-PONs serve the same number of users with the same number of transmitters as the SUCCESS-DWA PON. Under these constraints, the performance gain is quantified for many different traffic scenarios. Using the same constraints to compare a conventional WDM PON to the SUCCESS-DWA PON is not very informative, as a fully upgraded SUCCESS-DWA PON with the same number of TX/RX as end users perfectly emulates a conventional WDM-PON.

It is well known that the arrival of Internet packets can exhibit extreme rate variations in multiple time scales. The main reason for rate fluctuations at time scales greater than a few hundred milliseconds is due to extreme variability in the flow sizes, while the variability which occurs at smaller time scales is due to the burstiness induced by TCP [9], [10]. In the first set of simulations, we employ the α-stable self-similar traffic model developed in [11] to capture the burstiness and long-range correlations of realistic Ethernet packets. The characteristic exponent α is chosen to be 1.63 and the self-similarity parameter $\mathrm{H} = 0.80$, which were extracted from an actual file transmission over Ethernet in [11]. The traffic dispersion c_1 and the mean traffic c_2 are chosen at different values for different traffic loads while keeping a constant ratio of $\mathrm{c}_1/\mathrm{c}_2 = 2.5$.

The first simulations compare downstream performance of a four-TL SUCCESS-DWA PON to four TDM-PONs. In both cases, sixty-four users are evenly distributed across the four physical PONs. Fig. 6(a) and (b) plot the latency and queue depth characteristics, respectively, for the SUCCESS-DWA and TDM-PON architectures. Clearly, the SUCCESS-DWA PON outperforms the TDM-PONs by a sizeable margin. Note that both TDM and SUCCESS-DWA PONs are subject to the exact same traffic patterns in any given simulation run. Due to the

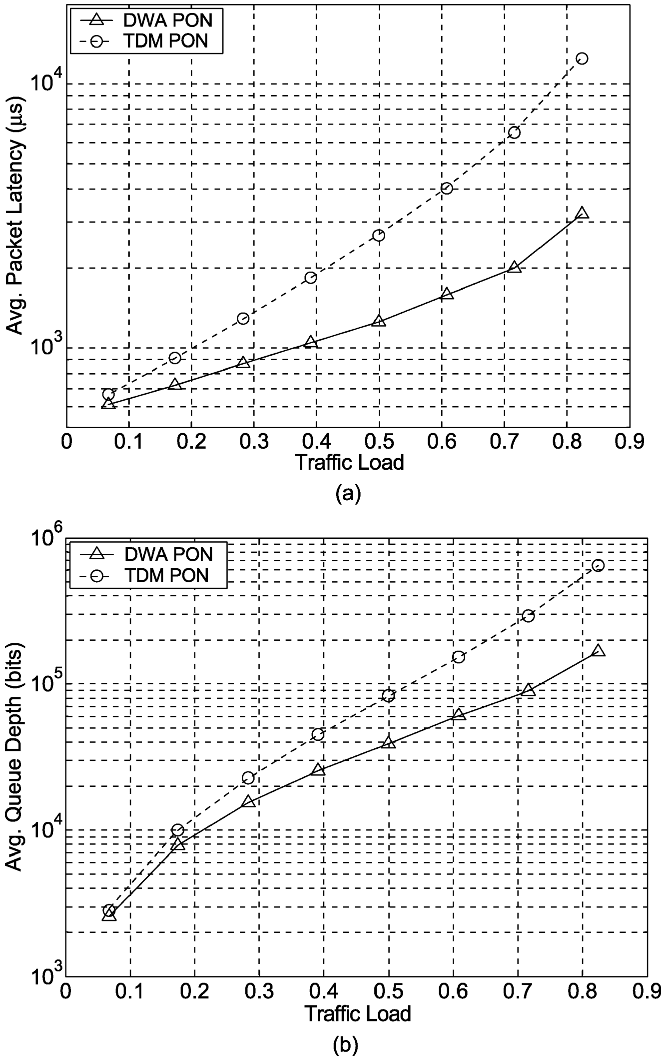

(a)

(b)

Fig. 6. Downstream characteristics (a) average packet latency (b) average queue depth.

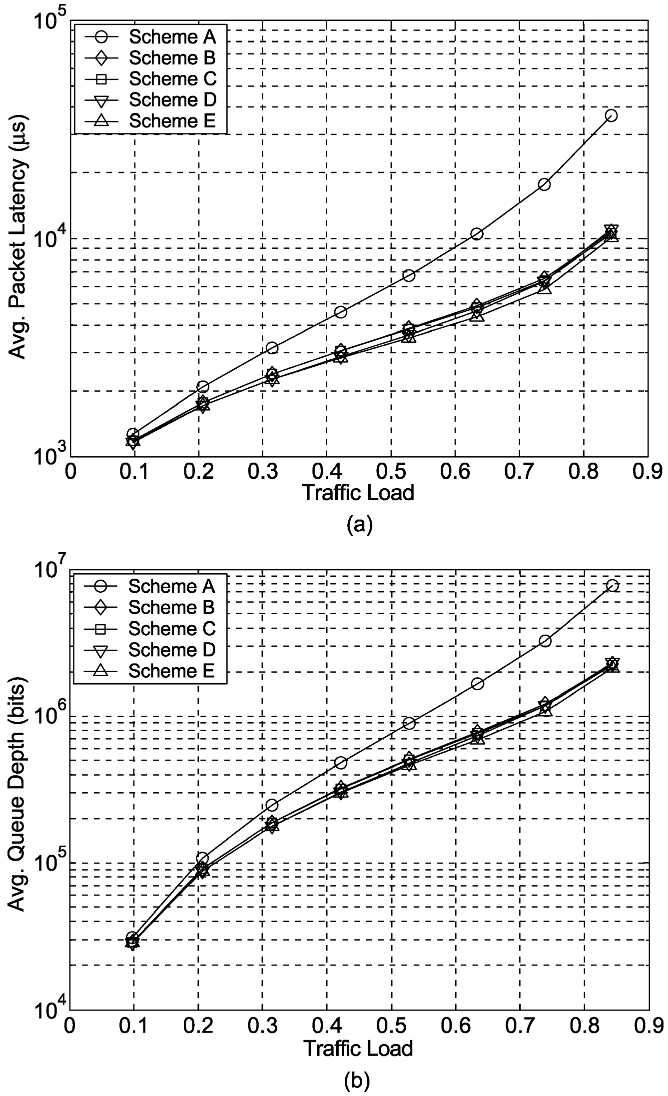

(a)

(b)

Fig. 7. Upstream characteristics: (a) average packet latency and (b) average queue depth.

fixed c_1/c_2 ratio, the variance of the queue lengths increases with the traffic load. This results in increased statistical multiplexing gain for the SUCCESS-DWA PON as the traffic load increases. If more TLs were coupled together, as in an 8×8 SUCCESS-DWA PON, the performance would be even better, since more transmitters can shift their wavelengths back and forth across multiple PONs when necessary.

For the upstream, Fig. 7(a) and (b) show the latency and queue depth characteristics versus the traffic load for the five schemes. For distributed Schemes B and C, half of the ONUs are equipped with tunable transmitters. Sixteen ONUs served by four PDs are considered in this analysis. Clearly, all the SUCCESS-DWA schemes B-E outperform the fixed-grouping scheme A. There is no significant difference among the upstream DWA schemes, even though schemes B and C are only partially populated with tunable devices.

Much of the work on traffic modeling has been derived from measurements taken either on the aggregated links of the Internet, or at the servers or campus networks attached to the Internet—the α-stable self-similar traffic model used up to now is one example. With the α-stable traffic model, every end user is assumed continuously downloading/uploading data with self-similar packet arrivals. Optical access networks, however, sit at the far edge of the network—close to the users, not the servers. Additionally, they have not been widely deployed and characterized.

The Markov process has been widely employed to model Internet traffic [12]–[15] and variable bit rate video traffic. A large number of ON/OFF sources with application specific transaction lengths or data volume distributions can be superposed to constitute a batch process [12]–[15] to emulate Internet traffic. In an attempt to approximate a sequence of transactions as initiated by an actual end user, we adopted a transaction-oriented Markov traffic model with specific transaction lengths corresponding to the popular user applications. In this approach, aggregation of the traffic from or to multiple end users results in a superposition of Markovian processes at the OLT, in which the traffic buffering and dispatching behaviors are of interest. Specifically, we consider three different families of transactions roughly corresponding to (a) web content [100 kB–300 kB], (b) larger files [3–4 MB] (i.e., MP3), and (c) significant data transfers [100 MB–200 MB] (i.e. multimedia content). The bursty nature of user transactions is modeled with discrete-time Markov chains.

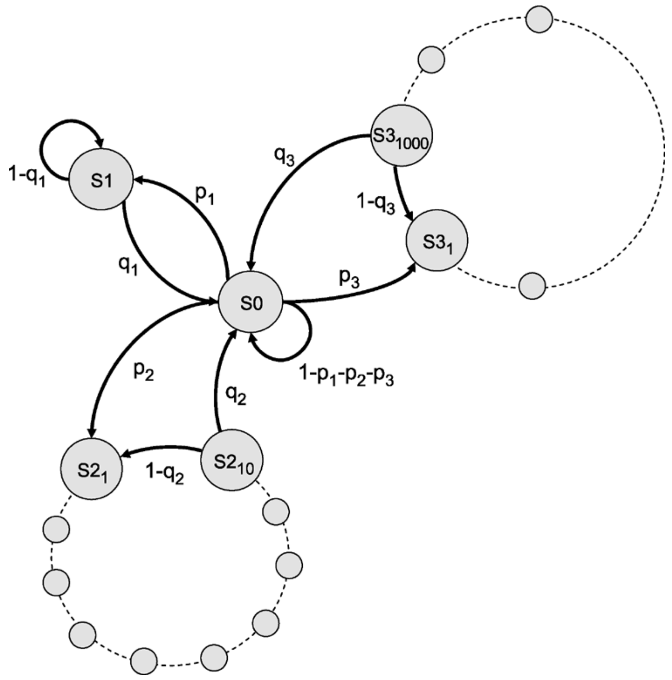

Fig. 8. Transaction-oriented traffic model based on Markov chains.

In Fig. 8, S1, S2, S3, and S0 correspond to transactions (a), (b), (c), and the idle state, respectively. The parameters are related by the equations

$$p_1 = \frac{q_1}{(\eta^{-1}-1)(1+10\alpha+1000\beta)} \tag{2}$$

$$p_2 = \alpha\frac{q_2}{q_1}p_1 \tag{3}$$

$$p_3 = \beta\frac{q_3}{q_1}p_1 \tag{4}$$

where η is the average traffic load of the user; α and β denote the ratios of transaction types (b) to (a) and (c) to (a), respectively.

For each time slot in the idle state, a user can initiate any one of the three transaction types or remain in the idle state. Once a transaction has been initiated, the users demand remains constant until the transaction is complete. By varying the probabilities of each type of transaction and its degree of burstiness, we are able to examine several different traffic scenarios. Different offered loads and subscriber scenarios are considered as well. Since the optimal total buffer size depends heavily on both the traffic and the practical limitations of the design, we chose to use infinite buffer sizes and focus on latency as a measure of performance.

Two different sets of traffic parameters were chosen: traffic scenario 1 (TS1) provides an approximately equal mix of transaction types; TS2 emphasizes the amount of small (web) transactions. In TS1, the transaction duration for web content varies considerably while larger files and significant transfers have relatively constant durations. In TS2, the transaction duration variation for web content is substantially reduced. Fig. 9 plots the latency characteristics for TDM and SUCCESS-DWA PON architectures for both TS1 and TS2. The SUCCESS-DWA PON significantly outperforms the TDM PONs. Under this traffic model, the distribution of active users tends to be more uneven when the traffic load is low, resulting in more substantial gains at lower traffic loads. Several other traffic scenarios were simulated—all showed appreciable performance advantages for the SUCCESS-DWA PON. Note that both TDM and SUCCESS-DWA PONs are subject to the exact same traffic patterns in any given simulation run.

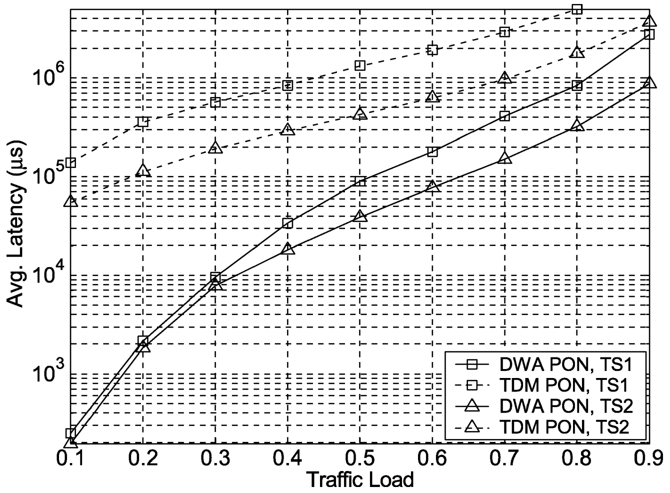

Fig. 9. Latency comparison between SUCCESS-DWA PON and TDM PONs for different traffic scenarios. For TS1, $\alpha = \beta = 1$, $q_1 = 0.5$, $q_2 = 0.9$, and $q_3 = 0.6$. For TS2, $\alpha = \beta = 0.1$, $q_1 = 0.8$, $q_2 = 1$, and $q_3 = 0.95$.

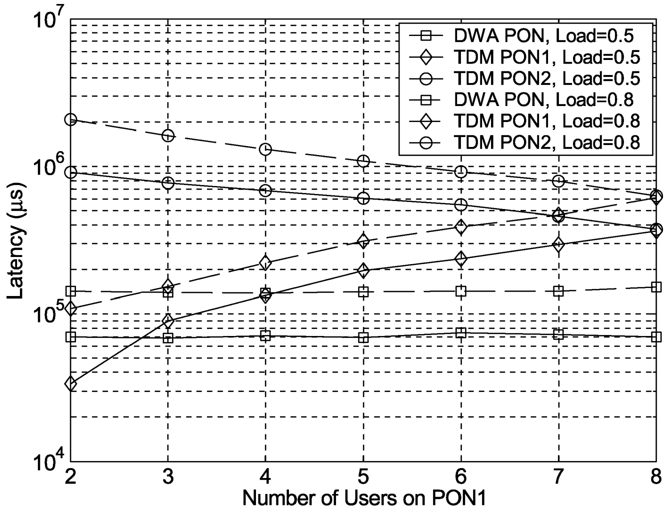

Fig. 10. Latency versus user distribution shows consistently good performance for the SUCCESS-DWA PON and strong dependence on user distribution for TDM-PON. Total number of users on PON1 and PON2 is 16.

Another key benefit of the SUCCESS-DWA PON is its indifference to user distribution. Consider four TDM-PONs where the subscribers are unevenly distributed—most of the active subscribers are attached to PON1, for example. Since the PONs are disjoint, the available resources from PONs 2, 3, and 4 cannot be used to improve the performance of PON1. The SUCCESS-DWA PON treats all users fairly, regardless of their location across the physical PONs. To illustrate this feature, simulations with varying user distributions were performed. In these simulations, two physical PONs support a total of sixteen users, i.e., an average subscription rate of 50%. In Fig. 10, the x-axis represents different possible distributions of users across the two PONs. As Fig. 10 indicates, not only do the TDM PONs perform poorly in general, but the TDM-PON with the most

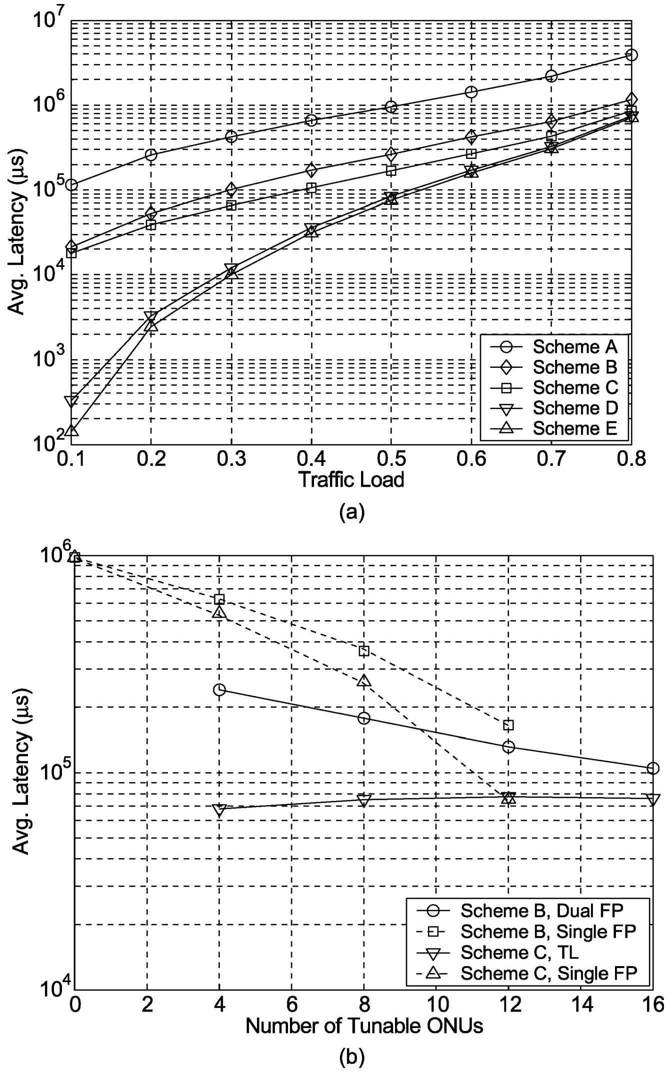

Fig. 11. (a) Average packet latency versus traffic load for all ONUs in Schemes A–E; half of the ONUs are equipped with tunable transmitters in Schemes B and C. (b) Average packet latency for fixed and tunable ONUs in Schemes B and C, when the number of tunable ONUs increase from 0 to 16. Traffic load is 0.5.

users suffers considerably. As expected, the SUCCESS-DWA PON performs equally well in all cases.

This model is also applied to the upstream simulations. Fig. 11(a) shows the latency characteristics for each of the five upstream schemes described in Section II. Scheme A exhibits the largest latencies due to the least flexibility, while the centralized tuning Schemes D and E show the smallest latencies because of full wavelength tunability. The distributed tuning Schemes B and C significantly outperform Scheme A when only half of the users are equipped with tunable devices. Fig. 11(b) illustrates the dependence of tunable and fixed ONU performances on the number of tunable devices in the network. Interestingly, as more users upgrade to the tunable devices, even the fixed ONUs benefit from the added tunability in the network.

IV. PRACTICAL ISSUES

A. Broadcast Support

The large-scale deployment of full-service optical access networks will come to reality if its cost can be low enough, it can enhance revenues, and reduce operating costs [16]. Delivery of digital video or multimedia services to the end users could be one of the driving forces for PONs. As high-speed access technologies become available to most end users, it is also expected that high-speed multicast applications will evolve quickly in most network environments. Therefore, in the physical layer architecture, it will be beneficial to incorporate broadcast or multicast abilities.

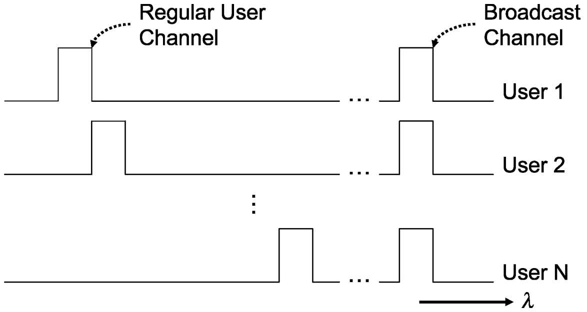

Fig. 12. User filter passband spectra with regular user channels and the broadcast channel.

TDM-PONs broadcast data downstream by design, and thus it is straightforward to realize broadcast when necessary. On the contrary, a full WDM-PON does not inherently support multicast or broadcast in the physical layer. For the SUCCESS-DWA PON, broadcast can be realized by having two passbands on the user WDM filters at the ONUs, as shown in Fig. 12. One passband is the unique user channel for receiving regular downstream traffic, and the other passband, common to all users, is a broadcast channel (BC).

For the illustrative architecture in Fig. 1, whenever the broadcast is activated on a certain PON, e.g., TL1 is transmitting on AWG channel 65, which is routed to PON1 and falls into the BC, all users on PON1 receive this signal while users on the other PONs are not affected. Therefore, we can realize broadcast in a PON-by-PON sense. Note that whenever broadcast is activated on a certain PON, all users on that PON are subject to the broadcast signal and cannot receive the regular downstream traffic. This is because an extremely cost-sensitive ONU would utilize only one photodetector—the normal and broadcast transmissions would have to share in the time domain. If instead, ONUs were fitted with two receivers, one dedicated to broadcast traffic, this constraint could be lifted.

B. Wavelength Plan

The TLs in the OLT can be accurately controlled in wavelength to allow DWDM in the downstream transmission. For the SUCCESS-DWA PON, consider a 4×4 AWG with 100 GHz channels and 16 user WDM channels. Together with the BC, the total wavelength range can be estimated as $0.8\ \mathrm{nm} \times 4 \times (16 + 1) = 54.4\ \mathrm{nm}$. Commercially available [17] fast TLs for telecommunications have output wavelengths falling within the C-band (1525–1562 nm) and L-band (1570–1615 nm), with a tuning range of $>\sim$60 nm. There are many possible ways to allow more user channels or more TLs, as shown in Fig. 13. With a finer spacing of the AWG channels, more user channels or more TLs can be added into the system. A tradeoff lies between the number of channels and the maximum modulation

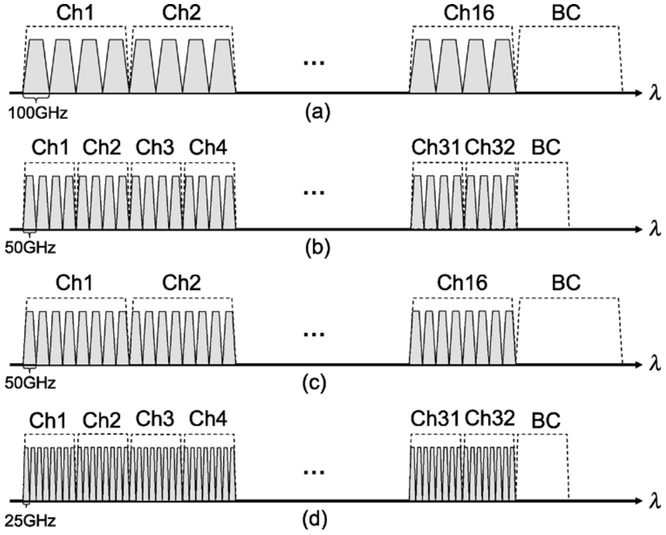

Fig. 13. Wavelength plan in the downstream (a) $4 \times 16 \times 100$ GHz, (b) $4 \times 32 \times 50$ GHz, (c) $8 \times 16 \times 50$ GHz, and (d) $8 \times 32 \times 25$ GHz.

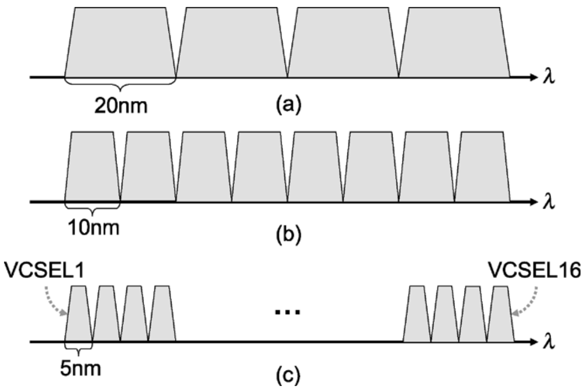

Fig. 14. Wavelength plan in the upstream: (a) distributed upstream schemes with four wavelengths falling on ITU CWDM channels, (b) distributed upstream schemes with eight wavelengths, and (c) centralized upstream scheme with 16 wavelengths.

bandwidth. In addition, the quality and corresponding cost of the required filters may make some scenarios suboptimal.

In the SUCCESS-DWA PON, only one set of user WDM channels is defined, which is reused in every physical PON. It is worth noting that the number of user channels is not a hard limit for the number of end users. A system with sixteen user channels can accommodate more than 16 end users on each physical PON if necessary. In this case, some users may be assigned user channels that are already occupied by other users on the same PON and must therefore time-share the bandwidth on that channel. The scheduling algorithm in the OLT is responsible for managing the traffic flow for fair access of end users.

The laser wavelength control in the ONUs is more challenging than in the OLT, and CWDM is more suitable for upstream transmission. For the distributed SUCCESS-DWA upstream schemes with four PDs in Fig. 3(b), four wavelengths corresponding to the ITU CWDM channels in the O-band (1260–1360 nm) can be utilized, which have a channel spacing of 20 nm as shown in Fig. 14(a). If eight PDs and eight corresponding wavelengths are of interest, one trivial approach employs 10 nm-spaced channels with each occupying half of the ITU CWDM channel as in Fig. 14(b). Nonstandard light sources and WDM filters may be required. For the centralized SUCCESS-DWA upstream schemes, since each user needs a unique upstream wavelength, cost-effective and narrow-linewidth VCSELs can serve as the upstream light sources. However, wavelength control may be necessary to achieve the required accuracy illustrated in Fig. 14(c). This can impose extra costs, favoring the distributed SUCCESS-DWA upstream schemes over the centralized ones.

C. Implementation Issues

Implementation of the SUCCESS-DWA PON protocols requires additional considerations including: laser tuning time, burst-mode preamble sequence, and additional guard time. The penalty resulting from this overhead heavily depends on the design of the TL, the burst-mode subsystems, and the specific implementation of the OLT queuing manager. For example, concatenating the packets into larger groups before transmitting causes only marginal penalties in latency but easily reduces the overhead to 1% or less.

The scheduling algorithm of the SUCCESS-DWA PON needs extra consideration to handle multiple transmitters. In the downstream, at most one TL can address a particular user at any instant in time. The scarce/shared resource in this architecture is the TL, since the number of TLs is typically less than the number of end users. The scheduler must maintain fairness and avoid collisions while maximizing the use of the TLs. One possible approach employs virtual output queues to avoid head-of-line blocking, and longest-queue-first scheme or other more sophisticated schemes to determine the order in which to service the queues. Since many TLs service the same set of queues, a suitable arbiter must be chosen to ensure that no two TLs attempt to service the same queue. Virtual output queues inherently avoid infrastructure collisions in this architecture. However, scalability can become a significant issue when the number of users is large. Another approach employs queuing at the transmitters. The incoming packets are assigned to the transmitter with the shortest queue to minimize delay, for example. To ensure that no two packets headed to the same end user overlap across the different TL queues, the scheduler can keep track of the destinations of the last packets in the queues and assign the incoming packet to the TL queue if they have the same destinations. Quality of service can also be achieved by prioritizing the packets when assigning to the queues. In general, a multiple transmitter system allows more flexibility to tailor the network performance than the conventional TDM-PON.

In the upstream, the scheduling algorithm should also prevent packet collisions in each of the upstream wavelengths. As in conventional TDM PONs, the ranging process is necessary to measure the propagation delays from the ONUs to the OLT, and the OLT allocates nonoverlapping time slots to the ONUs on a per-wavelength basis. Dynamic bandwidth allocation (DBA) schemes can be employed to maximize the throughput. If multiple lasers in a single ONU (see upstream scheme B) are permitted to transmit at the same time, the ONU can run at full speed when all its lasers are turned on. This allows the OLT more flexibility to allocate bandwidth among ONUs thereby accommodate very bursty Internet traffic.

TLs have been considered as light sources in optical access networks for several years [18]. As technology progresses, fast TLs with tuning times within 20–30 ns [19] have been experimentally demonstrated and can serve as the downstream light sources in SUCCESS-DWA PONs. On the other hand, active research efforts are focusing on tunable VCSELs [20], [21] whose potential low cost and high reliability would make them promising candidates for deployment in ONUs.

High-speed burst-mode receivers are one of the crucial components in optical packet-switched networks. Many research efforts continue in this field and promising achievements have been demonstrated [22]–[24]. Motivated by packet-switched applications, one can expect monolithic burst-mode receivers to be commercially available with costs low enough for large-scale deployment. Recall that the existing TDM-PONs require burst-mode receivers in the upstream.

In its basic form, the SUCCESS-DWA PON requires approximately 6–8 dB more additional power (AWG + up/downfilter + user filter) than a traditional TDM PON in the downstream. This additional power can be supported by lasers with higher output powers, optical amplifiers, or forward-error-correction (FEC) codes. In the upstream, the extra power loss is due to the insertion loss of the DeMux, for which the same techniques can be applied.

V. CONCLUSION

We propose a novel PON that employs dynamic wavelength allocation to efficiently provide services across several physical PONs. The design is extremely flexible in terms of capacity and immune to uneven user distributions across PONs. Compared to WDM-PONs, the SUCCESS-DWA PON enjoys flexibility, resource reallocation, and potentially lower costs. The SUCCESS-DWA PON also provides a gradual and cost-effective means to scale capacity as demand increases. The SUCCESS-DWA PON is designed to move bandwidth freely among multiple physical PONs, potentially greatly improving performance. The network can be scaled from one TL per K PONs all the way toward K TLs per one PON, where K is the number of AWG input/output ports. In addition, multiple physical PONs enjoy shared protection against equipment failures at the OLT. Judging from user demands, the fixed transmitters at the ONUs can be upgraded to tunable transmitters for better performance. It is demonstrated that even the fixed ONUs benefit from the upgrade. With excellent scalability, the SUCCESS-DWA PON can be configured to span the range of capacities between conventional TDM-PONs and full WDM PONs. Broadcast can be supported with the broadcast channel in the WDM filters. DWDM in the downstream and CWDM in the upstream constitute a reasonable compromise between performance and cost. The field-deployed physical infrastructure is kept untouched and need not be a specific topology. The powerful architecture is a promising candidate for next generation optical access networks. A testbed is currently on-going for the experimental demonstration of the SUCCESS-DWA PON.

REFERENCES

[1] S.-J. Park, S. Kim, K.-H. Song, and J.-R. Lee, "DWDM-based FTTC access network," *J. Lightw. Technol.*, vol. 19, pp. 1851–1855, Dec. 2001.

[2] G. Maier, M. Martinelli, A. Pattavina, and E. Salvadori, "Design and cost performance of the multistage WDM-PON access networks," *J. Lightw. Technol.*, vol. 18, pp. 125–143, Feb. 2000.

[3] R. D. Feldman, E. E. Harstead, S. Jiang, T. H. Wood, and M. Zirngibl, "An evaluation of architectures incorporating wavelength division multiplexing for broad-band fiber access," *J. Lightw. Technol.*, vol. 16, pp. 1546–1559, Sep. 1998.

[4] M. S. Rogge, Y.-L. Hsueh, and L. G. Kazovsky, "A novel passive optical network with dynamic wavelength allocation," in *Proc. OFC Tech. Dig.*, 2004, Paper FG1.

[5] N. Vicari and S. Köhler, "Measuring Internet user traffic behavior dependent on access speed," in *Proc. ITC Specialist Seminar on IP Traffic Measurement, Modeling and Management*, 2000, pp. 5-1–5-7.

[6] K. Yu and O. Solgaard, "MEMS optical wavelength deinterleaver with continuously variable channel spacing and center wavelength," *IEEE Photon. Technol. Lett.*, vol. 15, pp. 425–427, Mar. 2003.

[7] L. Domash, E. Ma, N. Nemchuk, A. Payne, and M. Wu, "Tunable thin film filters," in *OFC Tech. Dig.*, 2003, pp. 522–524.

[8] W. Shin and K. Oh, "The broadband wavelength selective cross-connect switch based on a fused optical fiber coupler," in *Proc. CLEO Tech. Dig.*, 2002, pp. 7–9.

[9] S. Floyd and V. Paxson, "Difficulties in simulating the Internet," *IEEE/ACM Trans. Netw.*, vol. 9, pp. 392–403, Aug. 2001.

[10] J. W. Roberts, "Traffic theory and the Internet," *IEEE Commun. Mag.*, vol. 39, pp. 94–99, Jan. 2001.

[11] A. Karasaridis and D. Hatzinakos, "Network heavy traffic modeling using α-stable self-similar processes," *IEEE Trans. Commun.*, vol. 49, pp. 1203–1214, Jul. 2001.

[12] L. Muscariello, M. Mellia, M. Meo, R. L. Cigno, and M. A. Marsan, A Simple Markovian Approach to Model Internet Traffic at Edge Routers.

[13] S. H. Kang and D. K. Sung, "Real time cell loss estimation for ATM multiplexers with heterogeneous ON/OFF sources," in *Proc. ICC'97*, Jun. 1997, pp. 1679–1683.

[14] W. Willinger, M. S. Taqqu, R. Sherman, and D. V. Wilson, "Self-similar through high-variability: statistical analysis of Ethernet LAN traffic at the source level," *IEEE/ACM Trans. Networking*, vol. 5, no. 1, pp. 71–86, Feb. 1997.

[15] A. Klemm, C. Lindemann, and M. Lohmann, "Traffic modeling of IP networks using the batch Markovian arrival process," in *TOOLS2002*, London, U.K., Apr. 2002.

[16] N. Chand, P. D. Magill, S. V. Swaminathan, and T. H. Daugherty, "Delivery of digital video and other multimedia services (>1 Gb/s bandwidth) in passband above the 155 Mb/s baseband services on a FTTx full service access network," *J. Lightw. Technol.*, vol. 17, pp. 2449–2460, Dec. 1999.

[17] B. Mason, J. Barton, G. A. Fish, L. A. Coldren, and S. P. DenBaars, "Design of sampled grating DBR lasers with integrated semiconductor optical amplifiers," *IEEE Photon. Technol. Lett.*, vol. 12, pp. 762–764, Jul. 2000.

[18] N. J. Frigo, P. P. Iannono, P. D. Magill, T. E. Darcie, M. M. Downs, B. N. Desai, U. Koren, T. L. Koch, C. Dragone, H. M. Presby, and G. E. Bodeep, "A wavelength-division multiplexed passive optical network with cost-shared components," *IEEE Photon. Technol. Lett.*, vol. 6, pp. 1365–1367, Nov. 1994.

[19] K. Shrikhande, I. M. White, M. S. Rogge, F.-T. An, E. S. Hu, S. S.-H. Yam, and L. G. Kazovsky, "Performance demonstration of a fast-tunable transmitter and burst-mode packet receiver for HORNET," in *Proc. OFC Tech. Dig.*, vol. 4, 2001, pp. ThG2-1–ThG2-3.

[20] L. Chrostowski, C.-H. Chang, R. Stone, and C. J. Chang-Hasnain, "Uncooled injection-locked 1.55 μm tunable VCSEL as WDM transmitter," in *Proc. OFC Tech. Dig.*, 2003, pp. 753–754.

[21] D. Vakhshoori, J.-H. Zhou, M. Jiang, M. Azimi, K. McCallion, C.-C. Lu, K. J. Knopp, J. Cai, P. D. Wang, P. Tayebati, H. Zhu, and P. Chen, "C-band tunable 6 mW vertical-cavity surface-emitting lasers," in *Proc. OFC Tech. Dig.*, 2000, pp. PD13-1–PD13-3.

[22] Y. Yamada, S. Mino, and K. Habara, "Ultra-fast clock recovery for burst-mode optical packet communication," in *Proc. OFC Tech. Dig.*, 1999, pp. 114–116.

[23] K. Murata and T. Otsuji, "A novel clock recovery circuit for fully monolithic integration," *IEEE Trans. Microwave Theory Tech.*, vol. 47, pp. 2528–2533, Dec. 1999.

[24] S. Brigati, P. Colombara, L. D'Ascoli, U. Gatti, T. Kerekes, and P. Malcovati, "A SiGe BiCMOS burst-mode 155-Mb/s receiver for PON," *IEEE J. Solid-State Circuits*, vol. 37, pp. 887–894, Jul. 2002.

Yu-Li Hsueh (S'03) received the B.S. degree in electrical engineering and the M.S. degree in electrooptical engineering from National Taiwan University, Taipei, R.O.C., in 1996 and 1998, respectively. He is currently pursuing the Ph.D. degree in electrical engineering at Stanford University, Stanford, CA.

His research interests include optical access networks, photonic crystal fibers, and optical waveguide modeling.

Matthew S. Rogge (S'03) received the B.S. degree in electrical and computer engineering from the University of Missouri at Columbia in December 1997. Since then, he has been at Stanford University, Stanford, CA, working toward the M.S. and Ph.D. degrees in electrical engineering.

He has worked on optical networking projects in the metropolitan and access spaces. He contributed to the HORNET project, a metropolitan area network that employs tunable transmitters to achieve efficient, scalable packet transport and is currently working on optical access networks.

Shu Yamamoto (SM'95) received the B.S., M.E., and Ph.D. degrees in electronics engineering from the University of Tokyo, Tokyo, Japan, in 1977, 1979, and 1989, respectively, and the M.S.E.E. degree from California Institute of Technology, Pasadena, CA, in 1983.

He joined Kokusai Denshin Denwa (KDD) Company, Limited (currently KDDI Corporation), Tokyo, Japan, in 1979. He has worked on optical fiber submarine cable systems and optical networking at KDDI Research and Development Laboratories. From 2000 to 2004, he was with KDDI Laboratories USA, Inc., Palo Alto, CA. In 2004, he returned to the KDDI R&D Laboratories, Inc., Saitama, Japan.

Leonid G. Kazovsky (M'80–SM'83–F'91) received the M.S. and Ph.D. degrees in electrical engineering from the Leningrad Electrotechnical Institute of Communications in 1969 and 1972, respectively.

From 1974 to 1984, he taught and was engaged in research at Israeli and United States universities. From 1984 to 1990, he was with Bellcore, Red Bank, NJ, researching coherent and WDM optical fiber communications systems. In 1990, he joined Stanford University as Professor of Electrical Engineering and founded the Photonics and Networking Research Laboratory.