Taking Turns with Adaptive Cycle Time

An Upstream Media Access Scheme for Extended-Reach FTTx

Ahmed H. Helmy and Habib Fathallah
Electrical Engineering Department,
and Prince Sultan's Advanced Technologies Research Institute
King Saud University, Riyadh, Saudi Arabia
ahelmy1@ksu.edu.sa and hfathallah@ksu.edu.sa

***Abstract*—Several centralized algorithms have been proposed for upstream bandwidth-allocation in passive optical networks (PONs); making the optical line terminal (OLT) in the central office the intelligent device that arbitrates time-division access to the shared upstream channel. When the distance between the OLT and optical network units (ONUs) is extended from 20 km to beyond 100 km, as suggested by next generation long-reach PONs, it becomes difficult for centralized algorithms to support service differentiation required for real-time applications. This is because these algorithms are based on bandwidth negotiation messages frequently exchanged between the OLT and ONUs, which become seriously delayed, when the network is extended, causing the performance to degrade. In this paper, we propose a distributed scheme for the emerging LR-PON while also allowing centralized control over the access network. Simulation results show that the average upstream packet delay can be significantly decreased below that of centralized algorithms while maintaining a high throughput.**

Keywords-distributed media access; Ethernet; fiber-to-the-x (FTTx); long-reach passive optical network (LR-PON); optical access network.

I. Introduction

Passive optical networks (PONs) seem to be a promising solution to the "last mile" bandwidth bottleneck that is exacerbating from the increasing demand over emerging and expanding Internet services. These access networks have basically no active elements in the signals' path from source to destination, but only employ passive optical components, such as optical fibers and splitters. Compared to other access technologies, they provide higher bandwidth, allow longer distance between customers and the central office (CO), and also allow easy upgrades to higher bit rates [1].

Recent advances in optical communications allowed for longer transmission and distribution distances for PONs. Such extended access networks were initially referred to as super-PONs, but are now more often called Long-Reach PONs [2]. However, since active components are incorporated, a more accurate term would be Long-Reach Optical Access Networks. Long-Reach access networks extend the feeder length from the traditional 20 km up to a 100 km and beyond by exploiting both optical amplifiers and WDM technologies. As a consequence, a LR-PON can combine access and metro optical networks into a single integrated system [2]. It incorporates a significantly greater splitting ratio substantially increasing the number of users served, thereby enabling full usage of optical capacity. However, this technology also comes with new research challenges [2], one of which is the increased round-trip time (RTT).

In a long reach optical access network, a single optical line terminal (OLT) resides in a main CO and connects the access network with the core network. The OLT implements layer 2 and layer 3 functions, such as resource allocation and service management. A DWDM ring connects the CO with multiple local exchanges residing in the end users' areas, with a dedicated wavelength for each local exchange, as illustrated in Fig. 1. After a local exchange, the fiber is split and connected to multiple optical network units (ONUs) within a drop section of 10 km [2]. Optical amplifiers are usually placed within local exchanges to compensate for the power loss due to both the long transmission distance and the high splitting ratio.

The logical connection between the OLT and ONUs in a LR-PON is illustrated in Fig. 2, with disregarding DWDM wavelength conversions and the local exchange, since both are transparent from the data transmission's perspective. Just as in a traditional PON, all downstream transmissions (from the OLT to ONUs) are done in broadcast-basis, since it is a point-to-multipoint network. In the upstream direction, the network is a multipoint-to-point network; multiple ONUs all transmit toward one OLT through a common passive combiner. Data from an ONU can only reach the OLT but not other ONUs, due to the directional property of the combiner. Downstream and upstream transmission channels are conventionally separated by assigning each a different wavelength. The main concern in PONs lies in the upstream direction, on which some media access control (MAC) mechanism should be employed to fairly coordinate the multiple ONUs' transmissions and avoid data collisions. This is known as the upstream bandwidth-allocation problem, which becomes very challenging when the distance between the OLT and ONUs is extended as in LR-PONs.

The rest of this paper is organized as follows. Section II discusses centralized bandwidth allocation (e.g., IPACT), explaining how its performance is highly affected by extending the network span. Section III introduces the principles of our proposed distributed approach and how quality of service can be managed by appending centralized control. Section IV presents numerical results and Section V concludes the study.

This research has been supported by the NPST program of King Saud University, Project Number 09-ELE667-02.

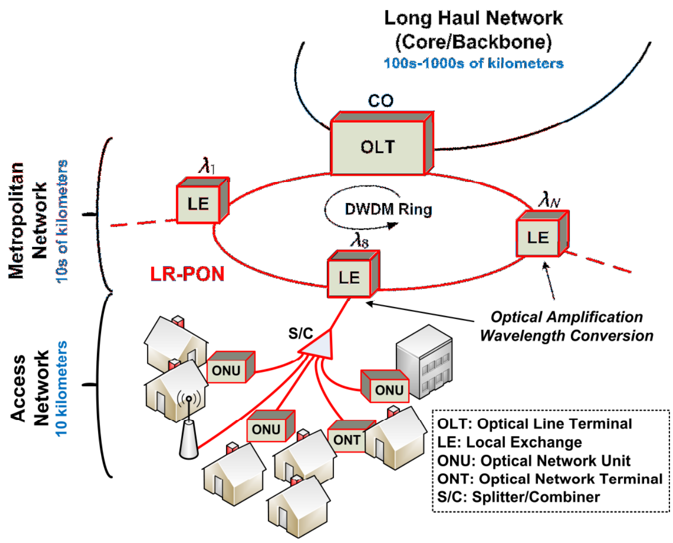

Figure 1. LR-PON combines metro and access networks into a single integrated system.

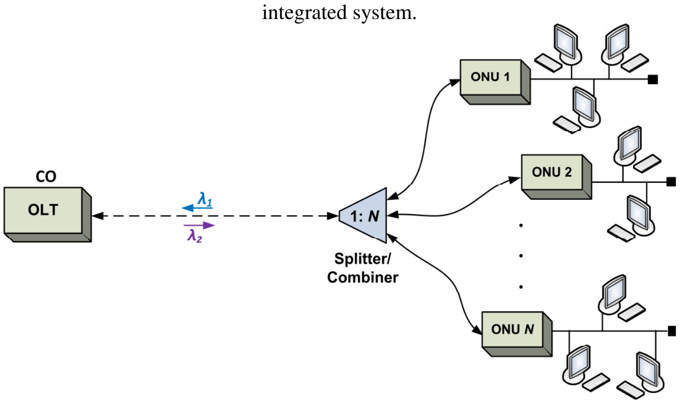

Figure 2. Logical connection between OLT and ONUs in a LR-PON.

II. Centralized Dynamic Bandwidth Allocation

Centralized arbitration schemes have been adopted in PON standards, in which the OLT arbitrates time-division access to the shared channel. In order to make accurate timeslot assignments in such arbitration schemes, the OLT needs to know the exact buffer state of a given ONU. Conventionally, polling schemes are used, which are based on report and grant messages defined in IEEE 802.3ah. In a polling scheme, the OLT keeps a polling table with an entry for each ONU to record both its round-trip time (RTT) and its buffer status. Each ONU sends a report message informing the OLT of the amount of buffered data it needs to send. The OLT continuously updates its polling table, processes all requests, and polls ONUs in a round-robin fashion granting them transmission windows (timeslots) via grant messages. The granted transmission window for each ONU corresponds to its reported queue status.

Therefore, in polling schemes, the OLT is able to know the buffer load of each ONU and allocate upstream bandwidth according to its bandwidth demand. The ONUs do not need to monitor the network state nor negotiate new parameters. However, the ability of centralized algorithms to maintain a small average delay is affected by the distance between the OLT and ONUs. This is because the OLT does not have real-time network-state information. Instead, it has a delayed version due to the long propagation delay. Moreover, after reporting its buffer status, each ONU must wait at least one RTT before granted media access. Therefore, the longer the distance is extended, the more the grant delays are significant.

Polling protocols typically operate in cyclic basis. Each ONU is polled once in each polling cycle and allocated a transmission window based on its demand. The maximum polling cycle ($C_{\max}$) allowed by the OLT is related to the maximum transmission windows of the N ONUs by:

$$C_{\max} = \sum_{i=1}^{N} W_{i,\max} + NT_g \qquad (1)$$

where $W_{i,max}$ is the maximum allowed window for ONU$_i$ and T_g is the guard interval between successive ONU transmissions. The polling cycle duration changes with the dynamics of the upstream traffic. Under light traffic loads, the cycle duration usually reduces to much less than $C_{\max}$, which has the effect of giving ONUs more bandwidth than the guaranteed, and thereby reducing the average packet delay. However, the polling cycle cannot be less than the maximum RTT. This is simply because each polling table entry should be updated within the OLT before issuing a grant to the corresponding ONU [3]. Under heavy traffic, the effective cycle duration reaches its maximum giving ONUs their minimum guaranteed bandwidths. We therefore define the effective cycle duration (C_{eff}) as the sum of the granted transmission windows within a given cycle;

$$\max(RTT) \le C_{eff} = \sum_{i=1}^{N} W_i + NT_g \le C_{\max} \qquad (2)$$

where W_i is the granted transmission window for ONU$_i$ corresponding to its buffer status reported to the OLT. Thus the polling cycle adapts to the network load, but it is lower-bounded by the maximum RTT. This makes centralized polling delays distance-dependent and greatly affected by extending the OLT-ONU distance.

Interleaved polling with adaptive cycle time (IPACT) has been a pioneer centralized algorithm for dynamic bandwidth allocation (DBA) in Ethernet PON [3]. It employs a pipelined timeslot-assignment, allowing the OLT to send a grant message to the next ONU before data and piggy-backed report message(s) arrive from the previously polled ONU(s). This is feasible since the upstream and downstream channels are separated, and since the OLT maintains relevant information about each ONU in the polling table. The OLT employs service level agreements (SLAs) of end users to upper-bound the allocated bandwidth (window size) of each ONU. To do so, several schemes were investigated in [3], of which the limited service discipline was found to exhibit the best performance. In limited service, the OLT grants an ONU the number of bytes requested, but no more than a certain maximum;

$$W_i = \min(R_i, W_{i,\max}) \qquad (3)$$

where R_i is the requested transmission window for ONU$_i$ and W_i is the granted one. Such a scheme has the shortest cycle and is also the most conservative as it assumes that no packets arrive after an ONU sends its request. To date, Research continues on adding to and further improving IPACT [4]. In this work, we compare the performance of our distributed approach with that of IPACT.

III. PROPOSED DISTRIBUTED SCHEME

Distributed schemes may show significant improvement over centralized schemes as ONUs are located further from the CO. In a distributed scheme, ONUs do not wait for grants in order to start transmission, but they themselves decide when to send data and for how long. An ONU that is scheduled next should time its transmission to arrive at the OLT right after the previous transmission. This requires additional connectivity between ONUs, which suggests some modification of PON. However, a pure distributed scheme does not support centralized control that is necessary for managing bandwidth allocation and QoS according to user SLAs. In this work, we propose a modified distributed scheme that maintains some form of centralized operator control, which we call Taking Turns with Adaptive Cycle Time (TTACT).

A. Principle of TTACT

As the distance between the OLT and ONUs is extended, we believe it is no longer appropriate for ONUs to report their buffer status and wait for a grant to transmit. To achieve better performance in an LR-PON, ONUs should be allowed to communicate together to manage media access. They will take turns in transmitting according to some sequence, with respect to certain maximum transmission windows set by the OLT. Such a scheme operates on a cyclic basis with adaptive cycle duration proportionate to the size of upstream transmissions. This new scenario however introduces two main challenges; obtaining means of communication between ONUs with no further fiber deployment, since it is both costly and unpractical, and maintaining centralized operator control over the network.

B. Aquiring Communication between ONUs

In a distributed approach, an ONU that is scheduled next in the sequence must time its transmission such that it arrives at the OLT right after the transmission of the previous ONU. This can be done in two ways; either by monitoring the previous ONU data transmission and detecting when the link becomes idle (e.g. carrier sensing), or by receiving a message declaring how many bytes the previous ONU will send and thereby knowing when to start its transmission. The former can be accomplished using the optical loop-back technique [5], in which the 1: N combiner is replaced with a 3: N and connecting two ports together. As illustrated in Fig. 3, this causes portion of the upstream transmission power to reflect back to ONUs. To sense each other's upstream transmissions, additional receivers operating at the upstream wavelength are placed within the ONUs. This technique however has two drawbacks; the upstream power loss due to reflection, which may limit the maximum OLT-ONU distance or the splitting ratio, and the varying idle periods between ONU transmissions due to different ONU distances from the combiner, as shown in Fig. 3. The idle period between two ONU transmissions equals the splitter-round-trip time of the one about to transmit;

$$T_{idle} = (2L_{SC,ONU_{i+1}})/S \tag{4}$$

where $L_{SC,ONUi}$ is the i^{th} ONU's distance from the splitter and S is the speed of light in the fiber. Unless subsequent ONUs are kept within 1 km propagation distance from each other, these idle periods will exceed the intended guard intervals causing poor utilization and slightly increasing packet delays.

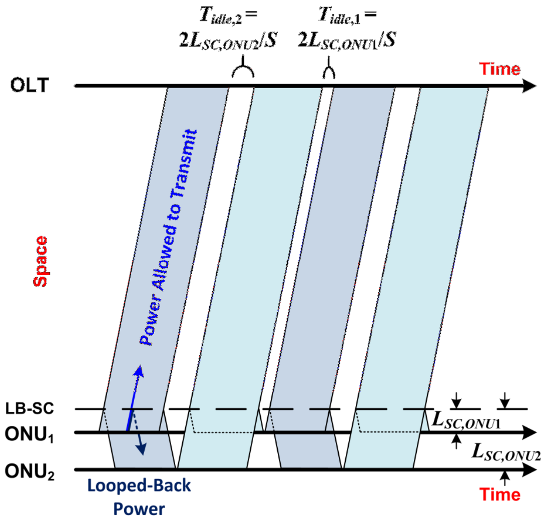

Figure 3. Varying idle-periods in the loop-back technique

To overcome the drawbacks of the optical loop-back technique, we propose using an additional wavelength for communication between ONUs to manage media access. As shown in Fig. 4(a), a fiber Bragg grating (FBG) is inserted close to the combiner in order to reflect the out-of-band control wavelength facilitating a multipoint-to-multipoint network. As illustrated in Fig. 4(b), each ONU will send a very short time-stamped frame at the beginning of its transmission, announcing how many bytes it intends to send, regarding the fact that ONUs have a common time reference. Upon receiving the frame, the following ONU will schedule its transmission after the previous one leaving a small dictated guard interval. This reduces inter-transmission idle periods significantly, since the time it takes for the frame to reach the following ONU will be during current ONU transmission. There may however still be some idle periods when a current ONU has nothing to transmit.

Unlike centralized schemes, the cycle duration of the distributed scheme is independent of the maximum round-trip time, making its delays (excluding propagation delays) independent of the distances between the OLT and ONUs. The cycle duration here is mainly dependent upon ONUs distances from the FBG. It cannot be less than the time it takes a packet to circle through all ONUs, which is the time it takes the packet to travel from each ONU to the FBG and reflect back;

$$C_{\min} = \sum_{i=1}^{N} \frac{2L_{SC,ONU_i}}{S} + N\frac{l}{R} \tag{5}$$

where l is the packet length, R is the transmission rate and l/R is the packet transmission delay. Thus, the distributed scheme will show significant improvement in the delay performance over a centralized scheme when the sum of all the distances between the FBG and each ONU is less than the maximum ONU distance from the OLT. In other words, it will show improvement when its minimum cycle given in (5) is less than the maximum round-trip time (centralized minimum cycle).

C. Centralized Operator Control

Centralized operator control is necessary to manage bandwidth allocation and quality of service (QoS) according to user SLAs. To maintain centralized control over the proposed

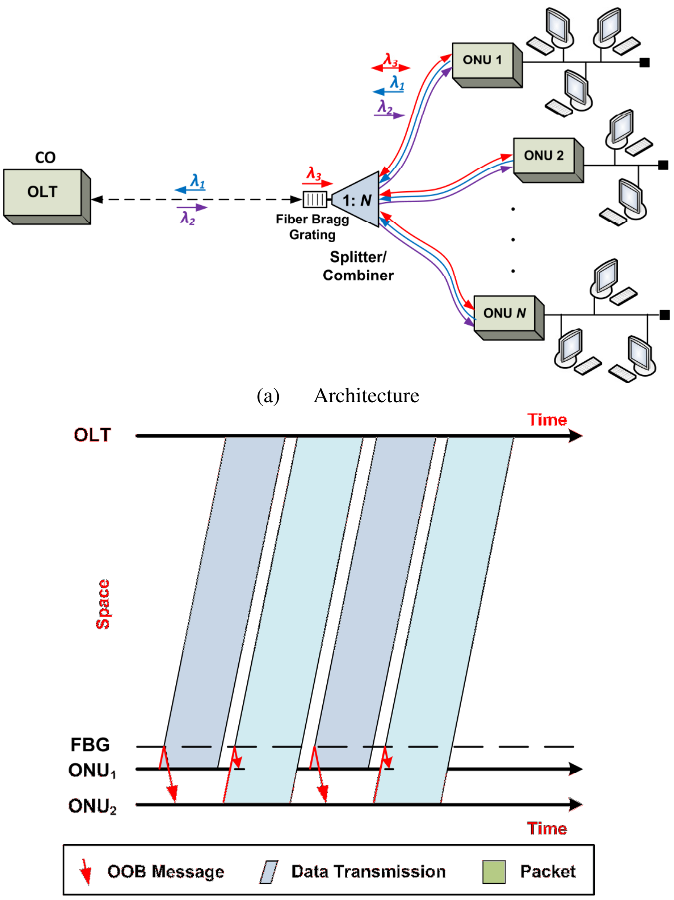

(a) Architecture

(b) Space-Time diagram

Figure 4. OOB loop-back technique using a fiber Bragg grating.

scheme, we propose that ONUs occasionally receive control parameters from the OLT. These parameters control how ONUs take turns transmitting by specifying their transmission sequence, the maximum allowable transmission window for each ONU, and its QoS parameters. The key feature of our scheme here is that ONUs do not have to wait for these parameters to transmit. These control parameters merely supervise and optimize the upstream transmission process. Some of these parameters are only changed according to SLAs or operator managements, whereas others change with response to network conditions (Table 1). The ONUs could be designed to operate on default parameter values or the last received ones. That way, the OLT needs only to send control frames when a parameter is required to change.

Service differentiation can be supported by two ONU scheduling paradigms: inter-ONU scheduling and intra-ONU scheduling [1]. The former is responsible for arbitrating transmissions of different ONUs, whereas the latter is for arbitrating transmissions of different priority queues in each ONU. In centralized approaches, the two paradigms can be implemented by either making the OLT perform both, or by setting the OLT to perform inter-ONU scheduling while each ONU performs its intra-ONU scheduling [1]. In our proposed distributed scheme, ONUs can perform both inter-ONU scheduling and intra-ONU scheduling, thus supporting real-time QoS scheduling. The OLT still maintains control via its

TABLE I. CENTRALIZED CONTROL PARAMETERS

Control Parameter	Parameter Arguments	
	Change Provoking Conditions	***Rate of Change (min: max)***
Tranmission sequence	Initialization	RTT: *x*RTT
	Faulty ONU detection	
	Idle ONU	
Maximum window for each ONU (W_i)	Assigned by operator (According to SLAs)	RTT/2: ∞
No. of priority queues for each ONU (P_i)	Assigned by operator (According to SLAs)	RTT/2: ∞

control frames by specifying the maximum transmission window for each ONU, and the number of priority queues to use, according to users SLAs.

D. *Transmission Sequence and Fault Management*

The ONUs transmission sequence can be set or even occasionally rearranged to improve the network performance in terms of both average packet delay and throughput. This is well demonstrated when the distances between the splitter and ONUs vary widely. For example, we assume two distant ONUs are in sequence. If the first does not have data to send, a long upstream idle period will take place until the control message reaches the next ONU. This leads to further delaying of packets waiting to be transmitted. Instead, to minimize idle periods and achieve better performance, the transmission sequence should be set so that each distant ONU follows a close ONU. With all ONUs active, an optimum sequence can be set corresponding to the ONU distances from the splitter.

Another advantage of having centralized operator control is monitoring and detecting faulty network equipment. The OLT can take a faulty ONU out of the sequence and update ONUs with the new sequence. For accurate monitoring and fault detection, ONUs control messages should be allowed to reach the OLT, which could be done by designing the FBG to pass portion of the control wavelength to the OLT. Allowing the OLT to receive copies of the control frames exchanged between ONUs can be very useful not only in fault detection. In fact, the OLT will also be capable of adapting the transmission sequence to network traffic. For example, if a given ONU continues to have nothing to transmit for consecutive cycles, the OLT can either move it in the sequence right before the closest ONU to the grating, or take it out of the sequence for a number of cycles. This has the capability of significantly reducing idle periods under certain network traffic conditions and improving the upstream channel utilization.

IV. SIMULATION RESULTS

In this study, we consider a 100 km LR-PON access network consisting of an OLT and 16 ONUs as in [6]. The FBG is assumed to be attached to the splitter/combiner and located 90 km away from the OLT. This is the worst case scenario for our proposed approach, since the grating is at the beginning of the drop section, making the furthest ONU 10 km away. The ONU distances are placed randomly from the local exchange to the maximum ONU distance. The ONUs share the same uplink wavelength of 1Gbps, whereas from the access

side end-users have an access rate of 100 Mbps. The system throughput is therefore less than the peak aggregated load from all ONUs. The traffic model used is Ethernet exponential traffic. Each ONU has a finite memory buffer of 10 Mbytes. The max cycle duration C_{max} is set to 5ms whereas the inter-transmission guard interval T_g is set as 5μs.

A comparison of average packet delays of IPACT with those of the distributed scheme TTACT is shown in Fig. 5. The distributed scheme shows significant improvement over IPACT mainly due to eliminating the grant delay component associated with centralized DBA. The throughput of IPACT is also compared with that of TTACT in Fig. 6. IPACT's maximum throughput is 97%, whereas TTACT achieves a maximum of 98.1%. The reason for the significant difference under heavy load is because an ONU in IPACT may be granted a timeslot smaller than requested more often, and with the OLT lacking any knowledge about its queue's composition, the timeslot may not exactly fit a number of frames. Since Ethernet frames cannot be fragmented, an unfitting frame will be differed to the next timeslot leaving an unused remainder in the current slot and slightly degrading utilization. This does not happen in TTACT since an ONU declares to the following ONU exactly how many bytes it will send.

Fig. 7 shows the effect of extending the network span on packet delays of both IPACT and TTACT when ONUs are not further than 10 km from the FBG. Extending the feeder shows to have less effect on the delays of the distributed scheme. Centralized IPACT delays are shown to be distance-dependent with a minimum of about 2RTT, whereas TTACT's delays go below 1RTT as the feeder distance is extended beyond 85 km.

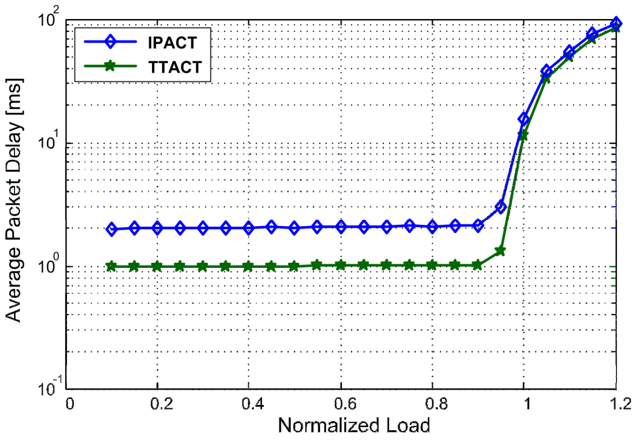

Figure 5. Average packet delay of both schemes in a 100 km LR-PON.

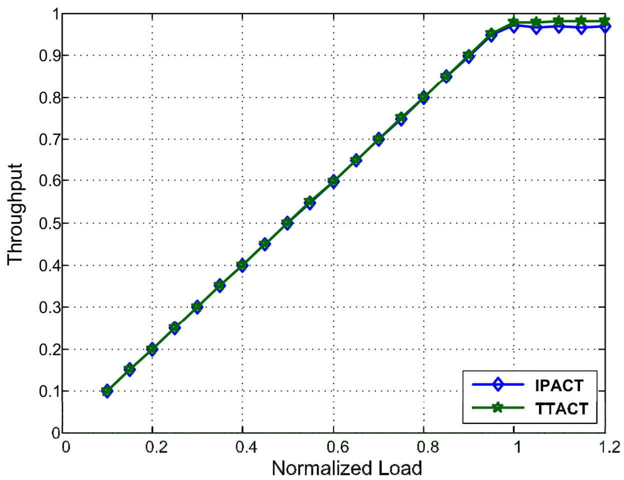

Figure 6. Throughput of both schemes in a 100 km LR-PON.

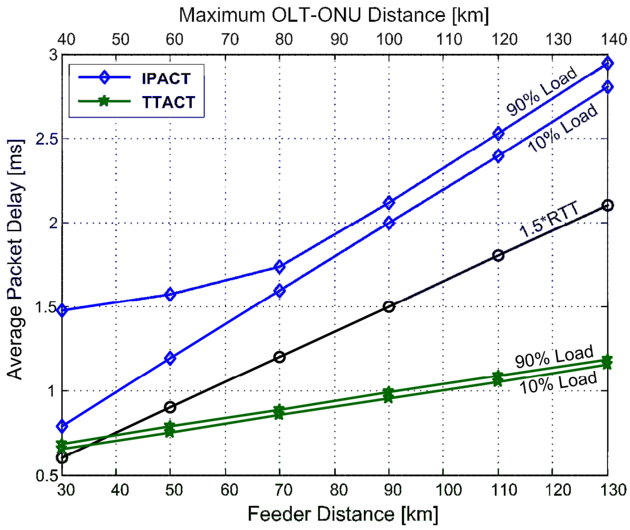

Figure 7. Distance effect on average packet delays of both schemes.

V. Conclusion

In this paper, we addressed the problem of bandwidth allocation in LR-PON. We proposed a novel distributed scheme we call TTACT to remedy the effect of the long CO-to-users control loop by suggesting out-of-band communication that allows ONUs to manage media access in a LR-PON, at the expense of placing additional transceivers. The improvement in the delay performance of this proposed scheme is bounded by the sum of distances between the FBG and each ONU being less than the maximum ONU distance from the OLT. Its performance gains more benefit over centralized approaches as the feeder distance is extended. Moreover, we proposed maintaining some form of centralized operator control to manage bandwidth distribution and QoS. Numerical results show that the average upstream packet delay can be decreased beyond that of a centralized scheme, while increasing the network throughput, since no bandwidth negotiation messages are exchanged between the OLT and ONUs, nor do timeslot remainders exist, as with centralized polling schemes.

References

[1] G. Kramer, "Ethernet Passive Optical Networks," McGraw Hill: New York, 2005.

[2] H. Song, B. Kim, and B. Mukherjee, "Long-Reach optical access networks: A survey," *IEEE Comm. Surveys and Tutorials*, vol. 12, no. 1, pp. 112-123, Feb. 2010

[3] G. Kramer, B. Mukherjee, and G. Pesavento, "IPACT: A dynamic protocol for EPON," *IEEE Comm. Mag.*, vol.40, no. 2, pp. 74-80, 2002.

[4] C. Sue and H. Cheng, "A fitting report position scheme for the gated IPACT", *ACM Trans. on Networking*, vol. 18, No. 2, pp. 624-637, 2010.

[5] B. Nesai et al., "An optical implementation of a packet-based (Ethernet) MAC in a WDM PON overlay," in Proc. of *OFC*, vol. 3, 2001.

[6] H. Song, B. Kim, and B. Mukherjee, "Multi-Thread polling," *IEEE J. on Selected Areas in Comm.*, vol. 27, no. 2, pp. 134-142, Feb. 2009.