Contention-Based Operation in TDM-Type Passive Optical Networks

Denis A. Khotimsky
Verizon
Burliungton, MA, USA
denis.khotimsky@verizon.com

Eugene B. DeFrancisci
Verizon
Wantagh, NY, USA
gene.b.defrancisci@verizon.com

Gregory K.Sherrill
Verizon
Ashburn, VA, USA
gregory.k.sherrill@verizon.com

Rajesh Yadav
Verizon
Boston, MA, USA
rajesh.yadav@verizon.com

***Abstract*— The fundamental operational principle of a Passive Optical Network (PON), which is based on a point-to-multipoint fiber infrastructure, is the unicast non-overlapping bandwidth allocations to the ranged Optical Network Units (ONUs). This paper details the modification to this principle, which has been proposed and advocated by Verizon, that introduces a broadcast bandwidth allocation to which eligible ranged ONUs respond only if necessary. While broadcast allocations in PON open door for the detectable collisions, they lead to notable bandwidth efficiency gains without requiring a nominal line rate increase.**

Keywords—passive optical networks, TDM/TDMA PON, contention-based operation

I. INTRODUCTION

The term Passive Optical Network (PON) refers to the optical access technology which employs a passively branched point-to-multipoint fiber infrastructure, known as an Optical Distribution Network (ODN), to provide connectivity between the central office equipment – Optical Line Termination (OLT) – and the customer equipment – Optical Network Units (ONUs). As in the upstream transmission direction the point-to-multipoint ODN effectively represents a shared medium, a large class of time-division multiplexing (TDM) PON systems employs Time-Division Multiple Access (TDMA) to control the upstream access to that medium. Each ONU is provided an allocation for an upstream transmission opportunity within a tightly controlled time interval.

From the day of its inception, the fundamental operational principle of a PON system has been the unicast non-overlapping nature of bandwidth allocations to the ranged Optical Network Units (ONUs) [1]. Within a sequence of sequentially transmitting ONUs, each upstream transmission, or burst, involves an overhead that is necessary for the OLT receiver to achieve fine adjustment to the mean optical power, clock rate and phase of the upstream optical signal. To ensure proper operation of a PON system the OLT should provide each ONU with an upstream allocation at a certain minimum frequency regardless of whether or not the ONU has data to send. Such operation and maintenance (OAM) allocations incur the same burst overhead as the data allocations. With the increasing line rates and a large number of ONUs that may remain idle, the inefficiency of contention-free allocation method becomes noticeable and substantial [2]. An alternative to the conventional contention-free allocation is to allow broadcast allocations for the OAM purposes that the ONU can use on an as-needed basis, and to return thus freed digital bandwidth to the pool that is subject to the Dynamic Bandwidth Assignment (DBA) to the actively transmitting ONUs.

The contention-based operation of the TDM PON systems has been proposed and advocated by Verizon and is presently included into the Transmission Convergence (TC) layer specifications of NG-PON2 and HSP PON systems.

II. THE CONCEPT OF CONTENTION-BASED PON OPERATION

A large class of passive optical network (PON) systems use time-division multiplexing (TDM) in the continuous-mode downstream transmission and time-division multiple access (TDMA) in the burst-mode upstream transmission. Termed TDM/TDMA PON (or TDM PON, for short), these systems constitute the dominant portion of all PON systems standardized via both ITU-T and IEEE and deployed to date. From the inception of TDM PON systems over 30 years ago, their main operating principle has been avoiding burst data collision in the upstream by accurately measuring the round-trip time from the Optical Line Termination (OLT) to each individual Optical Network Unit (ONU), establishing an equalization delay for each individual ONU, and scheduling individual burst data transmissions in a non-overlapping fashion, as seen by the OLT. Thus, with exception allowed for the ONU activation process only, the conventional operation of the TDM-type PON systems is contention-free: it is based on precisely timed directed upstream transmission allocations, or bandwidth grants, to individual ONUs.

The consecutive upstream transmission bursts originated by different ONUs are separated by the guard time and each carry a preamble and a delimiter to allow effective detection by the OLT receiver (see Fig. 1). The guard time, preamble and delimiter combined constitute burst-mode overhead that reduces the overall available upstream capacity and is viewed as a natural aspect of the TDM-type PON system. The burst overhead size is fixed: it does not depend on the burst payload size.

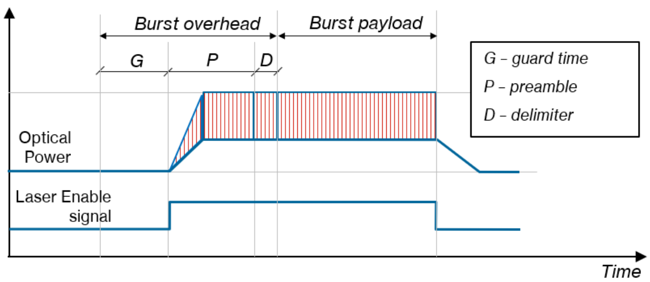

Figure 1— Physical layer burst structure

An upstream bandwidth scheduler in a TDM PON system resides at the OLT. Typically, it operates in a cyclic fashion allocating at least one upstream transmission opportunity to each ONU within a time interval known as Dynamic Bandwidth Assignment (DBA) cycle, T_{DBA}. The size of T_{DBA} typically equals one to four milliseconds. Even if an ONU remains idle (has no data traffic to transmit upstream), it requires a single bandwidth grant per T_{DBA} to report its bandwidth demand that may change suddenly. The useful payload of such an OAM grant is quite small and, therefore, its bandwidth efficiency is low. With the line rate increase, the number of ONUs per PON system may get larger, as may the size of the overhead due to larger preambles needed. Consequently, the overall fraction of time lost to burst mode overhead per each T_{DBA} interval grows bigger. Besides OAM grants, the scheduler may provide data grants to an ONU, according to the configured service. An idle ONU fills the directed data allocations with idle transmissions which are discarded by the OLT.

While a competent DBA algorithm would reassign most of the bandwidth assigned to an idle ONU to other ONUs, and would keep allocations to the protocol-exempt ONUs to a minimum, the use of directed allocations still incurs a notable overhead waste. Typically, in a PON-based access network, only a few ONUs are busy, that is, continuously have traffic to send upstream. The large majority of the ONUs are idle: with no upstream traffic, they follow the protocol requirements responding to the direct bandwidth allocations with idle transmissions. A significant potential for upstream bandwidth efficiency improvement is associated with reducing the idle upstream transmissions in a PON system.

Another opportunity to improve bandwidth efficiency is associated with the directed allocations that may remain empty on the protocol-based grounds. One example is wavelength protection in a time-and-wavelength division multiplexed (TWDM) PON system, wherein a set of ONUs that are operational in one wavelength channel are offered time-critical protection in another wavelength channel. The regular directed allocations provided in the protection channel would remain unused, except for the protection switching events. Another example is the protocol-based ONU power management operation, wherein an ONU is allowed to go asleep and ignore its directed allocations.

Being concerned with the efficiency improvement in PON systems without increasing the actual line rate, the authors have introduced and advocated the idea of contention-based operation of the TDM-type PONs. The core of the proposal is to provide a protocol mechanism replacing the directed allocations that are likely to procure an idle response or to remain unanswered with the contention-based broadcast allocations which multiple activated and ranged ONUs can use, but would actually do so relatively rarely.

Figure 2 illustrates the concept of contention-based broadcast allocations, while contrasting it with the conventional directed allocations per ONU. Out of the four ONUs which are present on the PON: A, B, C, and D, only ONU B has data to transmit upstream, whereas ONUs A, C, and D are idle. Under the conventional approach, each of four ONUs is provided with a directed allocation, three of which are filled with idle transmissions (I).

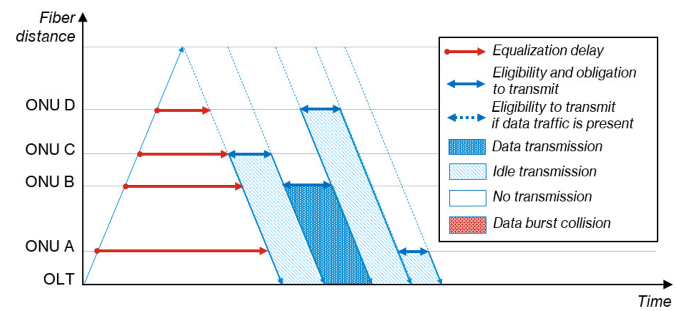

(I) Conventional directed allocations per ONU

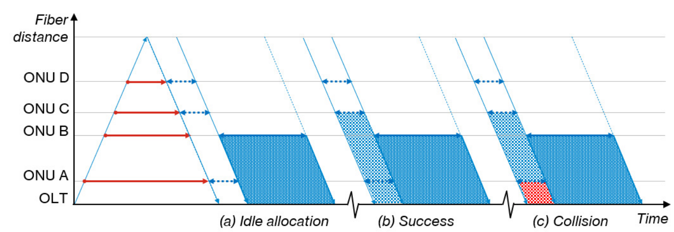

(II) Broadcast contention-based allocations

Figure 2 — Application of contention-based allocations

If the contention-based operation is enacted, one broadcast allocation replaces multiple directed allocations, freeing upstream bandwidth for the use by active backlogged ONUs. Under the protocol, an ONU can only use the broadcast allocation assigned for a specific contention-based function when a specific event occurs: a previously idle ONU receives data for upstream transmission; a previously sleeping ONU wakes up on a local stimulus; or protection switching takes place. As such trigger events are generally rare, most of the broadcast allocations remain idle (IIa). When a trigger event occurs at a single ONU, that ONU uses a broadcast allocation to transmit a burst containing an OAM message but no data, and the OLT upon receiving the burst restores directed allocations for that specific ONU (IIb). In rare cases, two or more ONUs may have a trigger event occurring in close proximity in time, in which case multiple ONUs may attempt to use the same broadcast allocation causing a collision (IIc). When the OLT detects a collision in a time interval corresponding to a broadcast allocation, it temporarily restores directed allocations to all ONUs that can potentially be involved in a collision to identify those ONUs requiring service, and later again withdraws directed allocations to those ONUs that do not need them.

III. Details of Contention-Based Operation

The components of the contention-based operation include: specification of the set of supported contention-based functions, a protocol for assignment of broadcast allocations; the collision detection method; and the collision resolution mechanism involving the transitions between the conventional directed allocations and the contention-based allocations.

A. *Examples of contention-based functions*

1) *Idle ONU Support*

The OLT employs a contention-based broadcast allocation for the ONUs that become active after remaining idle for some period of time. If an ONU consistently sends idle transmissions in response to directed data allocations, the OLT may withdraw directed allocations for such an ONU and include the ONU into the eligible list. An ONU that has not been receiving directed allocations and that has new data arriving for upstream transmission can use a contention-based allocation to request that directed data allocations be restored.

2) *Protection switching in a TWDM PON system*

The OLT employs a contention-based broadcast allocation in a pre-configured protection wavelength channel. All the protected ONUs are included into the eligible list. An ONU that tunes into its pre-configured protection channel upon a protection switching event can use a contention-based allocation to announce its arrival into the protection channel and request directed allocations from this point on.

3) *ONU power conservation support*

The OLT may employ a contention-based allocation to support autonomous wakeup for the ONUs that are conserving power by exercising a sleep mode behavior. Any ONU that has obtained OLT's consent to enter a sleep period is eligible to use such an allocation.

B. *Contention-based allocations*

The OLT associates an allocation identifier with each supported contention-based function. Unlike conventional allocations that are directed to a specific ONU, the contention-based allocations are broadcast, that is, they can be used by any eligible ONU on the PON. The OLT announces the assignment of broadcast contention-based allocations using the same OAM message that is used for the conventional allocation identifier assignment, repeating the announcement after each ONU activation.

C. *Collision detection*

On a high level, the OLT detects a collision within a transmission interval associated with a contention-based allocation, if it observes optical power within the transmission interval but fails to delineate a well-formed burst carrying an OAM message. More specifically, a successful transmission in a time interval associated with a broadcast allocation is indicated by a succession of three events: de-assertion of the Loss of Optical Signal (LOS) by the optical transceiver module; detection of burst delimiter by the Medium Access Control (MAC) logic, and validation of a well-formed burst, also by the MAC logic (see Fig. 3). The OLT detects a collision, if LOS has been de-asserted, but either of the two subsequently expected MAC events does not happen.

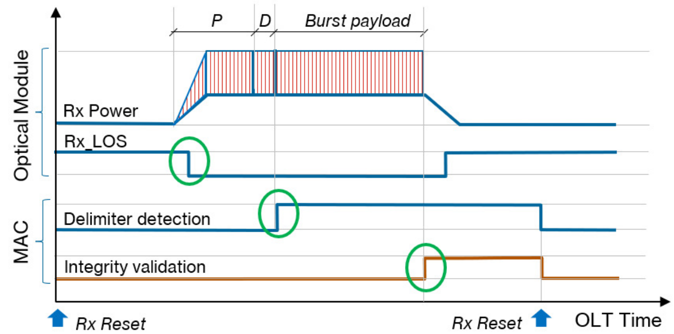

Figure 3 — Well-formed burst carrying an OAM message

D. *Collision resolution*

For each type of contention-based allocation, the OLT maintains a list of the ONUs that are presently eligible to use the allocation. Once a collision is detected, in the next available DBA cycle, the OLT replaces the contention-based broadcast allocation with the directed allocations to all eligible ONUs. Those ONUs whose response indicates that they are no longer eligible for the contention-based service continue to receive directed allocations for the time being and are removed from the eligible list. For the other ONUs, the OLT reconfirms their eligibility and withdraws directed allocations for the contention-based allocations.

IV. Performance Evaluation

The performance evaluation is performed for the idle ONU support use case, with a numeric PON throughput estimator that has been verified and shown accurate for multiple practical scenarios. A single-channel 9953.28 Gb/s symmetric system is considered with variable number of subtending ONUs. Only four ONUs are active whereas the rest represent the background load. The DBA cycle is set to 1 ms, with each active ONU being offered an upstream transmission opportunity once per 125 µs, whereas each background ONU transmits once per DBA cycle. When offered a directed allocation, each background ONU transmits a 132-byte burst composed of the OAM information, with the burst overhead consistent with the current state of transceiver technology (N2 optical path loss class). A background ONU randomly becomes active with the mean rate of once per second (that is, once per 1000 DBA cycles). A collision occurs if two or more ONUs become active within the same DBA cycle. To mitigate a collision, directed allocations are restored for the duration of two DBA cycles. The target function is the aggregate upstream throughput of the active ONUs which is determined with respect to the XGEM encapsulated user traffic (XGEM header is considered an overhead).

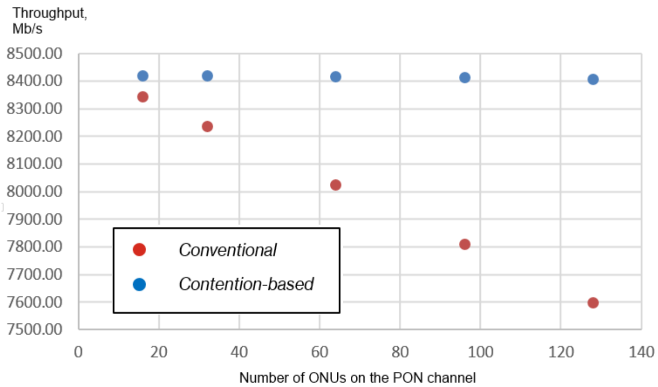

Figure 4 — Upstream service throughput comparison for the two allocation models

Figure 2 contrasts the aggregate upstream service throughput of the new contention-based allocation model with that of the conventional contention-free allocation model, considering them as functions of the number of subtending ONUs on the PON, which varies from 16 to 128. As the number of subtending ONUs increases, the aggregate upstream throughput of the PON system using conventional directed allocations falls nearly linearly whereas the throughput of the PON system using contention-based broadcast allocations remains essentially constant. The difference between the two reaches as much as 0.8 Gb/s for 128 ONUs.

With the use of the contention-based allocation model, the service throughput gain is achieved as a trade-off with the delay to provide the full assigned rate service to the newly active ONU. The additional delay typically equals two DBA cycles.

V. Summary

The concept of contention-based operation of TDM/TDMA PON systems has been recently introduced and analyzed. It is presently under the standardization approval process by the ITU-T. The contention-based concept represents a major step forward in 30-year operational history of PON by allowing an activated and ranged ONU to use broadcast allocations, which are also available to other eligible ONUs, in order to reduce the number of directed allocations that are likely to remain unused or idle. This approach leads to notable improvement in the PON system upstream bandwidth efficiency without changing the line rate. While primarily targeted at the upcoming generation of Higher-Speed PON systems, the contention-based operation feature set has been retrofitted into NG-PON2 specifications and can also benefit other standard-based PON systems. Its implementation for an ongoing NG-PON2 deployment has already started.

References

[1] Cedric F. Lam. Passive Optical Networks: Principles and Practice. Elsevier, 2007.

[2] Dave Hood, Elmar Trojer. Gigabit-capable Passive Optical Networks. Wiley, 2012.

[3] Recommendation ITU-T G.989.3 (2021) Amd. 1 (06/2023). 40-Gigabit-capable passive optical networks (NG-PON2): Transmission convergence layer specification. Annex J.

[4] Recommendation ITU-T G.9804.2 (2021). Higher speed passive optical networks – Common transmission convergence layer specification.