Optical Switching and Networking 13 (2014) 17–33

Contents lists available at ScienceDirect

Optical Switching and Networking

journal homepage: www.elsevier.com/locate/osn

TCP performance in multi-EPON access networks under different optical core switching paradigms

Maurizio Casoni [a], Walter Cerroni [b,*], Matteo Fiorani [a]

[a] DIEF, University of Modena and Reggio Emilia, Via Vignolese 905, Modena, Italy
[b] DEI, University of Bologna, Via Venezia 52, 47521 Cesena (FC), Italy

ARTICLE INFO

Article history:
Received 4 June 2012
Received in revised form
27 January 2014
Accepted 27 January 2014
Available online 13 February 2014

Keywords:
Ethernet passive optical networks
Optical circuit switching
Optical burst switching
Optical packet switching
TCP performance
Long-range EPONs

ABSTRACT

In this paper the end-to-end TCP performance of a hybrid network composed of multiple Ethernet Passive Optical Networks (EPONs) in the access segment connected to the same edge node of a core network is evaluated. Three possible core network paradigms are considered: Optical Circuit Switching (OCS), Optical Burst Switching (OBS) and Optical Packet Switching (OPS). TCP performance is evaluated through simulations with *ns*-2. The hybrid network scalability is assessed by varying the number of EPONs connected to the same edge node and the number of Optical Network Units (ONUs) per EPON. Interoperability issues between access and core are investigated and the impact of the most important network parameters is highlighted. The three transfer modes considered for the core network are properly compared under different input conditions, discussing the related trade-offs that may lead to the most suitable choice based on the specific application scenario. Finally, some design issues are investigated with reference to emerging Long-Range EPON (LR-EPON) solutions.

1. Introduction

New Internet applications, such as high-definition video delivery, reliable storage, cloud services and voice-over IP (VoIP), are pushing for higher capacities available at the user locations. To keep up with this trend, both access and core networks are experiencing substantial improvements.

As regards the access, passive optical networks (PONs) represent the most promising technology for providing the bandwidth required by the future Internet. PONs are capable of delivering bandwidth-intensive integrated voice, data and video services at distances up to 20 km. A PON consists of a point-to-multipoint (P2MP) optical network with no active components along the path between the central office and the user location. A concentrator is placed close to the end users with a passive optical coupler used to interconnect them. All the transmissions in a PON are performed between an Optical Line Terminal (OLT) and an Optical Network Unit (ONU). The OLT is located at the central office and it is directly connected to the backbone network, whereas the ONU is located either at the curb (namely Fiber To The Curb (FTTC) solution) or at the user location (namely Fiber To The Building (FTTB) or Fiber To The Home (FTTH) solutions).

In a PON the downstream transmissions are broadcast by the OLT to the ONUs through the 1:N passive coupler, where N typically varies between 4 and 64 depending on the signal power. Packets are extracted by their respective destination ONUs based on their physical address. In the upstream direction (i.e., from ONU to OLT), the traffic flows must share the same channel and thus a multiple access control scheme has to be implemented in order to avoid collisions. Based on the scheme used for the upstream

* Corresponding author. Tel.: +39 0547 339209.
E-mail addresses: maurizio.casoni@unimore.it (M. Casoni), walter.cerroni@unibo.it (W. Cerroni), matteo.fiorani@unimore.it (M. Fiorani).

http://dx.doi.org/10.1016/j.osn.2014.01.013
1573-4277

transmission, PONs can be divided into Wavelength Division Multiplexed (WDM) PONs and Time Division Multiplexed (TDM) PONs. In a WDM PON each ONU uses a different wavelength to transmit data to the OLT. Although this solution is very effective, it introduces very high costs due to the large number of channels needed, which makes it not very suitable for an access network in the short term. Therefore, TDM is the most commonly adopted solution so far.

In a TDM PON, only two wavelengths are used for upstream and downstream communications. In the upstream direction each ONU is assigned a time-slot and two successive slots are separated by a guard time. During each time-slot multiple frames can be transmitted. If N is the number of ONUs, each set of N slots and related guard-times is called a cycle. Two major standards for TDM PONs have emerged, Ethernet PON (EPON) [1] standardized by IEEE, and Gigabit PON (GPON) [2] standardized by ITU-T. Following the increasing demand for bandwidth in the access segment, both IEEE and ITU-T have recently ratified two new standards [3,4] for increasing the downstream transmission rate in PONs to 10 Gbps. These standards are characterized by the ability to coexist with traditional PON implementations. As observed by the authors in [5], so far the deployment of 10 Gbps-PONs is impeded by cost and implementation feasibility issues, and for this reason these technologies will represent an attractive solution for access networks only in the mid-term period. As a consequence, in the following we will take into account traditional PON implementations only. Due to its high level of development and recent advances, Ethernet appears to be the most suitable technology to be adopted in PONs. An Ethernet Passive Optical Network (EPON) is a PON that carries all data encapsulated into Ethernet frames and provides a symmetrical capacity of 1 Gbps.

As for core networks, Wavelength Division Multiplexing (WDM) technology is able to exploit the huge capacity offered by optical fibers by dividing the available bandwidth in multiple independent channels, each of which capable of carrying high data rate signals. A WDM network usually comprises edge nodes and core nodes. Edge nodes are located at the periphery of the network and are used to interconnect the core and the access segments. Core nodes are the internal nodes of the optical network and are used to route data from ingress to egress edge nodes, typically across long-haul links. Three different optical switching paradigms have been proposed to transparently forward data in WDM networks: Optical Circuit Switching (OCS), Optical Burst Switching (OBS) and Optical Packet Switching (OPS).

In OCS networks, a dedicated path (called a lightpath) connecting source and destination must be created across the network before the actual data transmission starts. A lightpath reserves a specific wavelength in each link along the path toward the destination and it is established using a two-way reservation mechanism. Data are then transmitted transparently along the lightpath. The main advantage of OCS is that, as long as the lightpath setup is successful, it provides a contention-free transmission. However, when the offered traffic is bursty in nature, the bandwidth available in a lightpath is underutilized, leading to poor network efficiency. Enabling technologies for OCS are well-established, making it a networking solution feasible in the short-term.

In OBS networks [6], data flows coming from access networks are gathered at the edge nodes and assembled into bursts according to a proper algorithm. Before a burst is sent, a control packet is generated and sent toward the destination to make a one-way resource reservation which is limited to the burst transmission time. The burst itself is sent after a fixed delay (called offset time) which is equal to or greater than the total processing time encountered by the control packet. Upon its arrival at each core node along the path, the burst size and arrival time are read from the control packet and the burst is scheduled in advance to an appropriate outgoing wavelength, thus avoiding the need for optical buffers. Due to the statistical multiplexing approach adopted, OBS provides higher bandwidth utilization than OCS in the presence of bursty traffic, but it may introduce high burst loss probabilities (due to contentions in the core nodes) and thus low achievable throughput. Enabling technologies for OBS are mostly available, making it a suitable mid-term solution.

In OPS networks [7], optical packets are sent without any resource reservation in advance. Control information is encapsulated into the packet header, which is separated from the packet payload by a time-guard. At each core node along the path the packet is stored into an optical buffer for the time required by the header processing. Wavelength converters and additional optical buffers must be implemented in core nodes in order to solve packet contentions. OPS provides very high bandwidth utilization and low packet losses, but introduces high costs and power consumption. This is mainly due to immature and expensive fast-switching technologies as well as optical buffering solutions based on Fiber Delay Lines (FDLs), which have large footprint and introduce high power losses.

The aim of this paper is to understand how Internet-based applications can be affected by the coexistence of optical access and core network technologies. This requires an end-to-end performance assessment study that must consider the entire path between two communication endpoints, including both the access and core network segments. Therefore, in this paper we evaluate the performance of Transmission Control Protocol (TCP), as the most widespread transport-layer protocol providing congestion control mechanisms, under different hybrid access/core scenarios with multiple EPONs on the access side and different switching paradigms in the core segment. This approach intends to provide the basis to discuss about different trade-offs that may lead to the most appropriate technological choices which satisfy specific application requirements. It represents a fundamental first step toward an optimal end-to-end network design and may help network operators in choosing the core switching technology that suits best to their requirements. Since we are interested in a high level analysis, we do not take into account the design details of the core network (e.g., network topology, burst/packet scheduling algorithms, burst/packet contention resolution mechanisms, etc.), which will be considered in future works on the topic.