Experimental Demonstration of DPDK Optimised VNF Implementation of Virtual DBA in a Multi-Tenant PON

Frank Slyne[(1)], Jasvinder Singh[(2)], Amr Elrasad[(1)],
Robin Giller[(2)], and Marco Ruffini[(1)]

[(1)] CONNECT Centre, Trinity College Dublin, {fslyne,elrasada,ruffinm}@tcd.ie
[(2)] Intel Corporation, Shannon, {jasvinder.singh,robin.giller}@intel.com

Abstract *We demonstrate a VNF implementation of a sliceable PON architecture which has been optimised using DPDK data plane acceleration techniques. This gives Virtual Network Operators optimal control over capacity scheduling in a large scale multi-tenant PON environment.*

Introduction

vDBA or Virtualisation of the Dynamic Bandwidth Assignment (DBA) function is a stepping stone towards providing full multi-tenancy in a PON environment[1]. vDBA gives Virtual Network Operators (VNOs) the ability to control capacity scheduling, as though they had dedicated and had full control of the physical OLT devices[2]. When PONs are used as bearers for other services such as LTE and future 5G services, VNOs require a more rigorous guarantee over levels of quality of service. By virtualising the DBA functions, VNOs get in-depth control, at the individual physical frame level, of the upstream capacity allocation in PON. This allows the enforcement of strict QoS parameters.

Experiment

This demonstration was produced as part of a cooperative project between Trinity College Dublin (TCD) and Intel Ireland. In our experiment, we implement a virtual OLT (vOLT) on standard server platform that uses open source technologies such as OVS-DPDK and DPDK for high throughput. Building on our previous experiment[3], this demonstration makes in-depth use of Intel DPDK libraries, providing enhanced scalability as a larger number of Virtual Machines for vDBA Virtual Network Functions (VNFs) can be integrated into the system through the virtual Switch (vSwitch). In our previous experiment, we had demonstrated a real time implementation of a sliceable Passive Optical Networks (PON) so as to enable multi-tenancy, based on the principle of virtual Dynamic Bandwidth Allocation (vDBA). We also demonstrated the feasibility of virtualising DBA functionality which was introduced in[1]. We moved the DBA functionality from vendor specific physical OLT to a software-based appliance running on General Purpose Processor. We identified a number of benefits arising from this, such as making the DBA algorithm more agile, more flexible and increasing market competition.

However, the performance of our previous demonstration was limited due to the overhead of packet handling and processing across the physical and virtual environment. The transmission of a large volume of small traffic status report packets upstream from the physical OLT to the virtualisation platform generates latency issues. In fact, the handling of such large volumes of small packets, in a General Purpose Processing environment, can cause significant interruption of shared resources such as CPU, memory and I/O. Such latency issues become problematic for the calculation of the DBA Bandwidth Map, which needs to be completed within the strict time limits of the frame interval (125 uSec), irrespective of the number of ONU's and T-CONTs.

Our new approach, developed in our current experiment, overcomes these constraints by optimising the data transmission and data processing of the vDBA scheme through minimising the I/O and packet processing cost. State-of-the-art DPDK data plane acceleration drivers and libraries are used through all application submodules. Our previous and current projects allow us to conduct a comparison between non-optimised and optimised scenarios, using performance indicators such as real-time latency, jitter, and throughput of each Virtual Network Operator (VNO) for each T-CONT in different traffic loading scenarios. Our experiment extends the vDBA implementation initially demonstrated in[1] and the concepts exposed in Fixed Access Network Sharing[2] and Cloud Central Office[4] frameworks. We demonstrate the Frame Level Sharing (FLS)[5] aspect of vDBA which allows the sharing of upstream frames among multiple VNOs to maximize bandwidth utilization, minimize latency, and provide a high level of service isolation among the VNOs sharing the PON, without increasing scheduling delay compared to traditional (non-virtualized) PONs.

Experimental Setup

Our demonstrator (see figure 1), implements a shared PON scenario with a number of real and virtualised ONUs, each with 3 T-CONTs. The main components within the testbed are: a physical PON, a set of emulated ONUs, a traffic generator and a multi-access edge computing node. The physical PON is based on one OLT and two ONUs (with the ONUs multiplexed into the same physical board), implemented on FPGA development boards[13] offering 10Gb/s symmetric capacity. The emulated ONUs, running in software, are used to increase the number of users, and generate typical self-similar traffic. The traffic generator produces both real-time sensitive and best effort traffic flows (such as file transfer and video streaming) through the physical PON. Traffic flows are VLAN-tagged which are then mapped to specific TCONTs at the ONUs. Openstack runs the Network Function Virtualisation (NFV) implementation of the PON, running the virtual DBA and the merging engine. The Merging Engine is the element that merges all virtual bandwidth maps from the different VNOs generating one physical bandwidth Map allocation and the SDN control plane. The virtualisation node is logically composed of the Virtual Network Functions(VNFs), an Openstack virtualization platform, a DPDK Data Plane Acceleration toolset and an Orchestration and Control layer.

We have implemented the Merging Engine (ME) and the vDBA functions for the Virtual Network Operators (VNOs) as Virtual Network Functions (VNFs), allowing these functions to be instantiated and scaled independently. The virtualized infrastructure, shown in figure 1, leverages Single Root Input/Output Virtualization (SRI-OV) technology[?] and Open vSwitch[6] with Data Plane Development Kit (DPDK) enhancements[7]. The DPDK offers a set of lightweight software libraries and optimized drivers to accelerate packet processing. It utilizes polling threads, huge pages, numa locality, zero copy packet handling, lockless queue and multi core processing to achieve low latencies and a high packet processing rate. Thus, all VNFs leverage the DPDK drivers and libraries to minimize the I/O and packet processing cost. The PCI Special Interest Group[8] on I/O Virtualization proposed the Single Root I/O Virtualization standard for scalable device assignment. PCI devices supporting the SRIOV standard present themselves to host software as multiple virtual PCI devices, thus introduce the idea of physical functions (PFs) and virtual functions (VFs). The PFs are the full-featured PCIe functions and represent the physical hardware ports; VFs are the lightweight functions that can be assigned to VMs. The userspace VF driver for the merging engine VNF helps VM to directly access the FPGA interface, thus, provides near line-rate packet I/O performance. The OVS-DPDK replaces the standard OVS kernel data-path with a DPDK-based data-path, creating a user-space vSwitch on the host for faster connectivity between VMs. The OVS-DPDK ports have vHost user interfaces which allow user to fetch/put packets from/to the VMs. Furthermore, all the VNFs in different VMs employs para-virtualized interface that utilizes the DPDK userspace virtio poll mode driver to accelerate packets I/O from OVS-DPDK. Each of the VNFs used for VNOs implements vDBA mechanism, thus, have identical functionality in terms of packet processing. The VM running merging engine VNF has two interfaces - VF interface for packets I/O with FPGA interface, and the second one is virtio interface to communicate packets with OVS-DPDK switch. There are two directions in which traffic flows in this virtualized system: North/South and East-/West. In the North/South flow pattern, traffic is received from the network through FPGA interface and sent back out to the network. In the East-/West flow pattern, traffic is processed by a VNF and sent to another VNF through OVS-DPDK for further processing.

Novelty

The virtualized solution is built upon high-performance software technologies, such as DPDK and OVS-DPDK, that reduce I/O virtualization overhead and accelerate packet I/O as well as packet processing performance. A user space device driver provides virtual machines (with vDBA VNF) with a fast access to the network interface card to send and receive packets from the network. Furthermore, to meet high performance target, the entire packet I/O functionality relies on DPDK user space drivers and libraries that offer I/O path optimization techniques

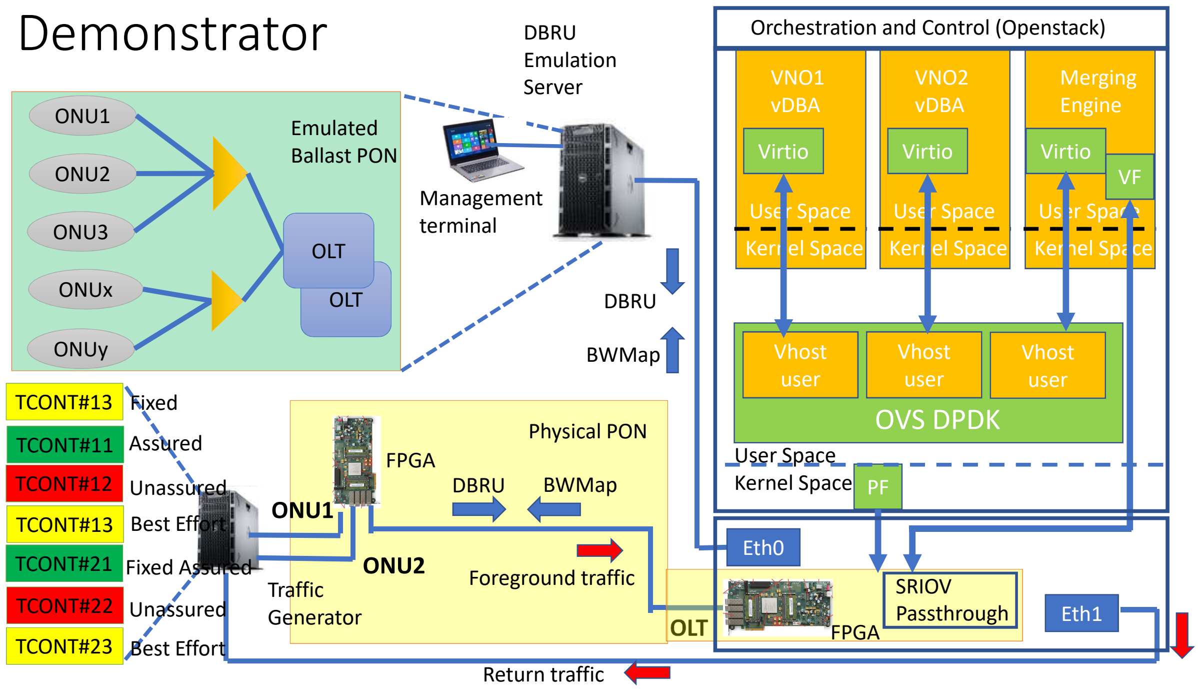

Fig. 1: Experimental setup

such as zero copy, batch processing of the packets, buffer allocations, interrupt-less I/O, etc. For inter-VM communication, OVS-DPDK offers high speed switching in user space, thus, avoiding kernel space data path and eliminating costly context switches between kernel space and user-space.

ECOC Relevance

This demonstration is highly relevant to the ECOC demonstration event as it shows an SDN/NFV implementation of a fully sliceable PON. The demonstration was carried out through the Sicence Foundation Ireland (SFI) projevct O'SHARE between TCD and the collaborating partner Intel Ireland. This topic is timely and highly relevant to industry, which is currently standardising the framework for Fixed Access Network Sharing (FANS)[2] and cloud Central Office (cloud-CO)[4]. The demonstration we propose is of high relevance to stakeholders (vendors, operators and service providers) interested in developing the next generation of access network infrastructure in support of 5G.

Acknowledgements

Financial support from Science Foundation Ireland (SFI) grants 14/IA/2527 (O'SHARE) and 13/RC/2077 (CONNECT) is gratefully acknowledged.

References

[1] Elrasad et. al, "Virtual Dynamic Bandwidth Allocation Enabling True PON Multi-Tenancy," OFC, 2017.

[2] Broadband Forum WT-370, 2017, http://www.broadband-forum.org.

[3] Slyne et al, "Demonstration of Real Time VNF Implementation of OLT with Virtual DBA for Sliceable Multi-Tenant PONs". OFC 2018.

[4] Broadband Forum WT-384, 2017, http://www.broadband-forum.org.

[5] Elrasad et al. "Frame Level Sharing for DBA virtualization in multi-tenant PONs." ONDM, 2017.

[6] Open vSwitch [Available online], http://www.openvswitch.org/features/ (Accessed on 17 May, 2018)

[7] Data plane Development Kit [Available Online], https://dpdk.org (Accessed on 17 May, 2018)

[8] PCI SIG. Single Root I/O Virtualization and Sharing 1.1 Specification, 2010.

[9] Ruffini, Marco, et al. "DISCUS: End-to-end network design for ubiquitous high speed broadband services." Transparent Optical Networks (ICTON), 2013.

[10] Cheng, Ning, et al. "Flexible TWDM PON system with pluggable optical transceiver modules." Optics express 22.2 (2014): 2078-2091.

[11] Alvarez et al. "Evaluating dynamic bandwidth allocation of virtualized passive optical networks over mobile traffic traces." JOCN 8.3 (2016): 129-136.

[12] Chanclou, Philippe, et al. "Optical fiber solution for mobile fronthaul to achieve cloud radio access network." Future Network and Mobile Summit, IEEE, 2013.

[13] McGettrick, et. al, "Experimental End-to-End Demonstration of Shared N:M Dual-Homed Protection in SDN-Controlled Long-Reach PON and Pan-European Core,"JLT, Vol 34, No. 8, Sept. 2016.