Deterministically Scheduled PON for Industrial Applications

Konstantinos Christodoulopoulos*, Sarvesh Bidkar, Thomas Pfeiffer, Rene Bonk
Nokia Bell Labs, Stuttgart, Germany
** now with the Department of Informatics and Telecommunications, University of Athens, Greece*
kchristodou@di.uoa.gr

Abstract: We propose a deterministic scheduling scheme for TDM-PON upstream bursts to achieve low latency and jitter with high throughput efficiency. We demonstrate co-scheduling of TDM-PON with an Ethernet Time Sensitive Network to serve industrial applications. © 2022 The Author(s)

1. Introduction

Passive optical networks (PON) are a mature, standardized and low-cost technology to provide broadband access to residential users in a point-to-multipoint (P2MP) topology. In recent years, new use cases for PON have been proposed and demonstrated, including the Passive Optical LAN (POL), the transport of 5G data in various functional-split configurations [1][2], and for industrial applications [2][3].

Several of the emerging applications are time-critical: they require low latency, and some require low peak-to-peak latency variation (jitter). For example, the most stringent class of industrial applications, the "isochronous real-time", require a maximum of 1 μs jitter [4]. To achieve cost savings and co-optimization, the target is to converge multiple time-critical applications together with best effort applications on the same network [5]. For such a converged network, Ethernet IEEE 802.1 standard was extended with a set of time sensitive networking (TSN) features that provide various means for temporal control of traffic flows [6]. For example, for 5G fronthaul a relevant TSN standard is IEEE 802.1Qbu, implementing frame pre-emption, while for industrial applications a relevant standard is IEEE 802.1Qbv, implementing scheduled gates at the switch ports. Many applications with strict latency requirements have known traffic patterns. For example, industrial applications typically follow a cyclic communication pattern between the controller and endpoints: fixed size packets transmitted in a specific order, repeated at a specific period. Field buses and networks (e.g. Profinet) and TSN 802.1Qbv [6][7] isolate and follow each time-critical flow cycle to control its latency through scheduled reservations, while also serving best-effort traffic at the non-scheduled time periods.

To reduce the latency in a TDM-PON we can provision a traffic container (TCONT) with committed information rate (CIR) equal to application's bandwidth (BW) requirement and delay tolerance (T_{BDT}) (see G 989.3 Amd3 [8]) equal to application's period. However, the placement of upstream (US) bursts depends on other TCONTs and their priorities. Thus, PON ensures correct bandwidth and burst rate on average but does not exactly follow each packet. This results in relatively high latency and jitter (Fig. 1a). To further reduce the latency, we can configure a small burst period, offer additional bursts, but sacrifice BW efficiency (Fig. 1b). If the burst period is a divisor of the PON frame duration we also have the advantage to repeat the same bandwidth map (BWmap) at each frame. Still, a mismatch between flow and burst periods create jitter, which can disrupt the time-critical application. For cases in which this jitter is unacceptable, the PON can be enhanced with jitter compensators with a small increase in latency [3].

In this paper we demonstrate an alternative solution to [3], a deterministic DBA (detDBA) that controls the placement of bursts at each frame for selected TCONTs, and provides them with deterministic performance (lossless transmission and controlled latency). For a periodic application, the PON is synchronized with the application and follows its period (Fig. 1c) to achieve the lowest possible (scheduling) latency, low jitter and consumes only the required BW. The

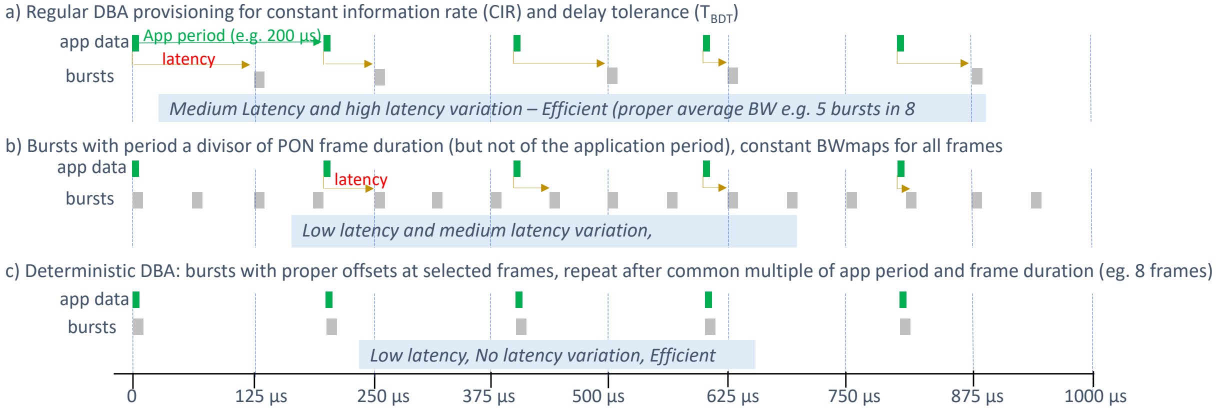

Fig. 1: Options to control the latency of a periodic flow over a TDM-PON: a) TCONT provisioning with constant information rate and delay tolerance descriptors, b) Constant BWmaps with bursts period a divisor of the frame duration, and c) proposed deterministic DBA

detDBA can coexist with a regular DBA (regDBA) to serve both time-critical and best-effort traffic. As a use case, we apply the detDBA to serve industrial flows and demonstrate low latency and jitter, high BW efficiency, while simultaneously serving best-effort traffic.

2. Deterministic DBA (detDBA) process

We assume that we can identify in advance the TCONTs that require deterministic performance and obtain relevant parameters about their patterns. For example, for a periodic flow we need to know its packet size, period and starting time. Such parameters can be passed through the regular PON traffic descriptors, or require extensions/interfaces, which are left outside of this work. An example control plane interface for mobile fronthaul is the co-operative transport interface (CTI) [9] specified by O-RAN. Such input can be processed by a cooperative (Co) DBA [10]. So, in this paper we assume that we are given a set of periodic flows and their relevant parameters as deterministic TCONTs.

The proposed detDBA consists of a scheduler and an enhanced BWmapper that implement the Co-DBA concept. For the set of deterministic TCONTs the scheduler finds the bursts' starting times (and durations - assuming one burst per packet) by solving the non-pre-emptive idling scheduling problem of concrete periodic tasks [11]. It calculates the bursts schedule for the duration T_{detPON}, which is the least common multiple (LCM) of the PON frame duration (125 μs) and the periods of the deterministic traffic TCONTs and repeats the schedule every T_{detPON}. The schedule defines for each flow a burst offset. For a flow with multiple periods in T_{detPON} the offset is applied at each period, defining thus the bursts starting times over T_{detPON}. The offsets are calculated to avoid collisions among flows within the T_{detPON}. The offset is equivalent to the scheduling (queueing) latency of the flow packet at ONU buffer until served by its allocated burst. Since a single offset is calculated per flow and repeated at flow period, all flow's packets have fixed scheduling latency (zero scheduling jitter). This holds for a flow period that is multiple of the PON word (12.86ns in XGS-PON = PON scheduling granularity); else the scheduling jitter will be less than one PON word. Note that scheduling can be complex (NP-hard) [11], but for a low number of requests, a shortest period first heuristic achieves good performance. Moreover, the scheduler runs at provisioning time, so its execution time is not very crucial.

To yield deterministic performance the PON needs to be synchronized and co-scheduled with the interfaced network and/or the applications served. For example, assuming a TSN IEEE 802.1 Qbv network serving industrial (time-critical) flows, the TSN switches would be synchronized by using IEEE 802.1AS, a precision time protocol (PTP) profile, to form a boundary clock distribution tree [12]. Commercial PON systems support PTP, so the PON can be part of the clock distribution tree. The proposed scheduled PON is equivalent to a TSN switch with ports at the Optical Network Units (ONUs) and at the Optical Line Termination (OLT), thus can be part of the TSN network scheduling.

In the proposed solution, the detDBA and regDBA coexist at the OLT (Fig 2). The scheduler calculates the schedule (bursts starting times and durations) for the deterministic TCONTs at provisioning time. The regDBA works in cycles of duration C. The schedule duration T_{detPON} can be larger or smaller than C. So at each start of a regDBA cycle the scheduler informs the regDBA of the BW it has allocated in that cycle, so that the regDBA can allocate the remaining BW to the non-deterministic TCONTs. The scheduler transfers the schedule (in the form of deterministic BWmaps) to the enhanced BWmapper which runs per PON frame. The enhanced BWmapper plays out the deterministic BWmap part for the frame and fills the unused timeslots with bursts for the non-deterministic TCONTs according to regDBA.

3. detDBA PON use case: industrial networking

To demonstrate the capability of detDBA PON to provide deterministic performance to selected TCONTs we implemented a converged TSN-based industrial network testbed (Fig. 3). In such a setting, we envision to integrate or partially replace the TSN switches with the detDBA PON. Serving time-critical applications and interworking with TSN requires to support PTP within the needed accuracy of 1 μs [12] and to co-schedule with the TSN. The testbed consisted of an XGS-PON extended to support the detDBA (implemented through the loading and execution of custom BWmaps), two TSN switches (Hirschmann RSPE35) and an Ethernet traffic generator and analyzer (IXIA Novus). We synthesized two industrial flows: flow #1 was 1250 Bytes at 200 μs period, and flow #2 was 625 Bytes at 250 μs

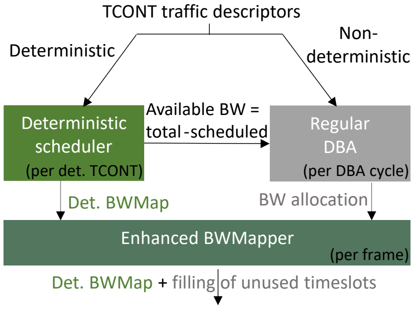

Fig. 2: Flow diagram of the Deterministic DBA (detDBA) and regular DBA

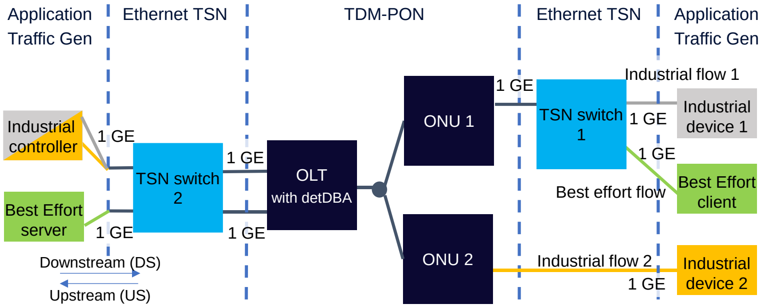

Fig. 3: Diagram of the TSN and detDBA PON converged industrial network setup (PON participated in PTP boundary clock distribution tree with root the controller)

Disclaimer: Preliminary paper, subject to publisher revision

period. We also created best effort (BE) traffic with random size packets (64-1500 Bytes) and 300 Mbps (average) rate. We focused our evaluation on the upstream direction (US), since it is more challenging in TDM-PON.

Initially we measured the latency and peak-to-peak latency variation (jitter) for the two TSN switches alone, scheduled to isolate and follow the industrial flows cycles. We measured upstream average latency of 26.09 µs for industrial flow 1 and of 6.61µs for flow 2. The different latency was because flow #2 crosses only one TSN switch and has a smaller packet size. The measured jitter was below 1 µs for both industrial flows and 73 µs for BE. These measurements are shown in Fig. 4 in orange.

We also measured the performance for the standalone detDBA PON, synchronized to the traffic generator. Note that industrial flow 1 period (200 µs) is not a multiple or divisor of the PON frame duration (125 µs), and thus it is non-trivial to be served with a regular DBA. We calculated the schedule and loaded the BWmaps of T_{detPON}=1 ms duration (8 PON frames = LCM of industrial flow periods and of PON frame) to the OLT and repeated every 1ms. The bursts for industrial flow 1 were of 1280 Bytes every 200 µs and for flow 2 were of 680 Bytes every 250 µs (set in equivalent PON words [8]). BE traffic was served with a 1 Gb/s TCONT and delay tolerance of 250 µs (2 frames). After the ONUs were ranged we disabled ranging. We measured average latency of 36.12 µs for industrial flow 1 and of 28.38 µs for flow 2 and jitter of 230 ns. For comparison, BE flow jitter was 223 µs (according to configured delay tolerance). These measurements are shown in Fig. 4 in blue.

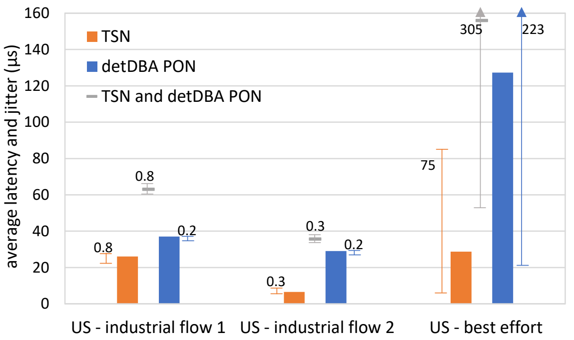

Fig. 4: Measured average latency and peak-to-peak variation (jitter) for i) TSN, ii) detDBA PON, and iii) TSN and detDBA

Then we connected the TSN switches and the detDBA PON (Fig. 3). We made the PON part of the PTP boundary clock distribution tree. We co-scheduled the PON and the TSN switches to minimize the US latency of the industrial flows so that packets do not wait in any queue (ONUs or TSN switches). The measurements are shown in Fig. 4 in grey. We measured jitter below 1 µs for the industrial flows, which was our target, and a much higher jitter for BE. The average latencies of the industrial flows equal to the sum of the latencies of the detDBA PON and of the TSN network (within 1 µs margin) measured individually. This indicates that packets were not additionally queued (as targeted in co-scheduling calculation). In downstream, we relied on TSN scheduling and also achieved low latency and jitter.

In [3] we developed a solution for a converged industrial network using Jitter Compensators (JC) [13] that created a tunnel from an ONU UNI port to an OLT SNI port and yielded constant latency (accounted for as a link in TSN scheduling). The solution in [3] did not require synchronization of the PON with the TSN and allowed the PON to have its own burst period. However, it overprovisioned for the full link rate (1Gb/s) and exhibited latency higher than the minimum (increase equal to the PON burst cycle, set to 31.25 µs in [3]). Compared to [3], the proposed detDBA yielded lower latencies for the industrial flows (the lowest possible in this demonstrator), jitter in the same order, and substantially improved the PON BW efficiency (reserved only the required BW). The disadvantage is that detDBA requires synchronization and co-scheduling.

4. Conclusions

We proposed a deterministic DBA that controls the placement of bursts within PON frames for selected TCONTs to provide deterministic performance. For a TSN network serving industrial periodic flows, we made the PON part of the PTP boundary clock distribution tree and co-scheduled it with the TSN switches. We demonstrated the lowest possible scheduling latency and low jitter for industrial flows, while simultaneously serving best effort traffic.

Acknowledgements: The authors would like to thank R. Sundaralal, J. Poffyn, and W. V. Hoof for their support in the measurements. This work was partly supported by the German BMBF with the EU CELTIC project AI-NET ANTILLAS.

5. References

[1] S. Bidkar, et al., "First Demonstration of an Ultra-Low-Latency Fronthaul Transport Over a Commercial TDM-PON Platform", OFC 2018.
[2] Th. Pfeiffer et al., "PON going beyond FTTH", IEEE J. of Opt. Comm. and Netw., Jan 2022, p.A31-A40.
[3] K. Christodoulopoulos et. al. "Demonstration of Industrial-grade Passive Optical Network", OFC 2022, paper Tu2G.6.
[4] S. Dietrich, et. al., "Performance indicators and use case analysis for wireless networks in factory automation", ETFA, Limassol, 2017.
[5] Industry 4.0 OPC Foundation Webinar, "OPC UA: Industrial interoperability standard from sensor to cloud (and back)", 01. Dec. 2020.
[6] Time-Sensitive Networking (TSN) Task Group, https://1.ieee802.org/tsn/
[7] S.S. Craciunas, et al., "Scheduling Real-Time Communication in IEEE 802.1Qbv Time Sensitive Networks", RTNS, ACM, 2016.
[8] ITU-T G.989.3 (2015) Amd. 3 (03/2020)
[9] O-RAN Alliance O--RAN.WG4.CTI--TCP.0--v01.00 (April 2020), "Cooperative Transport Interface Transport Control Plane Specification"
[10] ITU-T G Suppl. 71 (04/2021), "Optical line termination capabilities for supporting cooperative dynamic bandwidth assignment"
[11] K. Jeffay, D.F. Stanat, C.U. Martel, "On non-preemptive scheduling of period and sporadic tasks", Real-Time Systems Symposium, 1991.
[12] IEEE Standard for Local and metropolitan area networks, IEEE Std 802.1AS-2011.
[13] W. Lautenschlaeger, et al., "A scalable factory backbone for multiple independent time-sensitive networks", J. of Systems Architecture, 2021.

Disclaimer: Preliminary paper, subject to publisher revision