Kent Academic Repository

A Converged Evolved Ethernet Fronthaul for the 5G Era

Philippos Assimakopoulos[ID], *Member, IEEE*, Jim Zou, *Member, IEEE*, Kai Habel,
Jörg-Peter Elbers[ID], *Member, IEEE*, Volker Jungnickel[ID], *Member, IEEE*,
and Nathan J. Gomes[ID], *Senior Member, IEEE*

Abstract— We assess the performance of two distinct functional splits based on latency/latency variation and mapping efficiency, both individually and in unison. By considering hardware-offloading possibilities for a low-layer split (especially a pre-resource mapper split) using an option-6 software-based local thermal equilibrium split as an example, we show how data rate, Ethernet frame size and, in general, traffic generation characteristics will be very important aspects in the design of the future Ethernet mapping function. Then, an integrated Ethernet fronthaul with legacy and new/evolved split functionality operating at 100 Gb/s link rate is presented with the state-of-the art sub-100-ns latency variation for a timing-protocol flow. This is achieved through the application of a gap-filling aggregator, used for the first time in such a mobile fronthaul application.

Index Terms— Mobile fronthaul, Ethernet fronthaul, cloud-radio access network (C-RAN), time-sensitive networking.

I. INTRODUCTION

THE adoption of Ethernet in the fronthaul of the next-generation Radio Access Network (RAN) will extend packet-based networking to the network edge. As a ubiquitous technology with well-established Operations, Administration and Maintenance (OAM) functionality [1]–[3] Ethernet can enable fixed-mobile convergence, multi-vendor operability, reducing costs through newly established economies-of-scale [4], and unified transport over what is now termed the x-haul (x signifying front, mid and back). However, the effects of Ethernet networking on the performance of current 4th Generation (4G) and future 5th Generation (5G) systems and the implications of combining Ethernet with functional splitting must be taken into account. The main constraints will include (the lack of) frequency and time synchronization [5],

both of which have stringent requirements for 4G and 5G processing, especially for Co-ordinated Multi-Point (CoMP), Multiple-Input Multiple-Output (MIMO) antenna systems, and transmit diversity techniques. Synchronization features can be enabled using adapted forms of existing over-Ethernet technologies, namely Synchronous Ethernet (SyncE) (see [6] and accompanying standards) and Precision-Time Protocol (PTP) (see [7] and accompanying standards). In addition, latency caused by port contention/queuing, fronthaul-to-Ethernet mapping functions, subframe timings, Hybrid-Automatic Repeat Request (HARQ) protocol constraints and Ethernet line rates have to be taken into account [8]. Port contention in Ethernet switches/aggregators, in particular, will cause latency variations and will thus require the adoption of time-sensitive networking (TSN) approaches to constrain the variation within acceptable limits.

Current fronthaul, based predominantly on the semi-proprietary Common Public Radio Interface (CPRI) [9] involves In-phase and Quadrature (IQ) time-domain symbol transport and is not able to scale to the increasing data rate demands of 5G use cases (and some of the advanced 4G use cases). The implementation of different functional subdivisions (or "split processing") as a means of reducing the data rate is under consideration by 3GPP [10]. A higher layer split (HLS) has already been agreed upon, but no consensus has been reached (at the time of this paper's writing) for a lower-layer split (LLS). In this paper, 3GPP split terminology is employed. However, other groups have defined their own naming conventions, and a new framework (note, not specification) from the CPRI group has been released, accommodating additional functional splits and transport technologies including Ethernet transport [11].

Additional key requirements for a functional split implementation in a 5G RAN will include the ability to jointly process signals, modularization of functionality, operation within a network slicing and network function virtualisation (NFV) framework, which may be combined with variable split point selection (including the fully centralized split (I/Q based) for reasons of prior investment and backward compatibility).

An example of an *evolved* fronthaul is shown in Fig. 1. The Ethernet-based fronthaul transports a number of LLS-generated flows from Central Units (CUs) to Distributed Units (DUs) which can be fully or partially (through hardware-offloading techniques) virtualised.

Fig. 1. Example of a fronthaul implementation. GPP, Generic Processing Platform; UE, User Equipment; EPC, Evolved Packet Core.

The different flows are identified using Virtual Local Area Network identifiers (VLAN IDs) with priority-based scheduling taking effect in Time-Sensitive Networking (TSN) switching nodes. Each flow includes a number of packet-types (as defined by the control- and data-plane processes) which can themselves be further associated with a Class-of-Service (CoS), or potentially bundled under a single CoS (as shown in Fig. 1). Different flows may be processed according to network slice, whereby slice isolation (hard, soft or softer) is obtained through the allocation of scheduler resources in the switching nodes.

The paper focuses first on the study of two LLS implementations. The first is a "5G-ready" hardware (HW)-based upper-PHY split with custom PHY processing. The second is a 4G software (SW)-based Option-6 split which despite having low-4G data rates does include the whole LTE protocol stack. These splits are then considered in unison (in a HW offloading framework) and assessed in terms of mapping-to-Ethernet methods, and latency/latency variation. Then, a promising TSN aggregation approach operating at very high networking rates is assessed in terms of latency variation performance when transporting an in-band timing protocol, while aggregating the LLS-generated data and additional background traffic flows in an evolved and converged Ethernet fronthaul. Its performance is compared with theoretical predictions of alternative scheduling regimes.

Section II of the paper presents the two LLS split implementations while Section III focuses on TSN methods and presents the gap-filling aggregator for the support of 5G services and measurement results for the aggregated fronthaul. Finally, the paper is concluded in Section IV.

II. THE EVOLVED ETHERNET FRONTHAUL

A number of split points, both static and variable [12], have been suggested leading to an envisioned *flexible RAN* architecture that can accommodate variable split points dependent on application scenario. The choice of split point is not a straightforward one, as each possesses its own advantages and disadvantages [10]–[13]. With functional splitting, the advantages of Ethernet become clear. Certain split points offer data rates that scale with cell load, for which Ethernet aggregators/switches can offer *statistical multiplexing* gains.

Such functional splits will be key enablers for the next-generation 5G RAN. There is currently limited available experimental work, especially within an Ethernet fronthaul

context. In [13] and [14], a software-emulated Option-6 split was presented, concentrating on latency performance [14], and Ethernet data rate contributions from different transport channels [13] while Makris *et al.* [15] concentrated on the real-time operation capabilities of software platforms with functional splitting. However, there is currently a lack of detail of the implementation, especially with regards to mapping functions and their efficiency, and the exposition of multiple flows in an evolved Ethernet fronthaul. With regards, to HW-based splitting, in [16], a 4G-based upper-PHY split was evaluated specifically for CoMP and at 4G data rates.

In the next subsections, we concentrate on the performance of the two LLSs. Important aspects of the implementations presented here include data rate and latency performance, and mapping approaches when migrating to Ethernet transport. For the latter, the IEEE1914.3 Radio-over-Ethernet (RoE) task group has standardized a number of radio-over-Ethernet mappers, specified for CPRI traffic (with some provisions for future functional splits) [17], which can offer a starting point. However, such mappers may offer insufficient information for LLSs (as will be described) especially pre-resource mapper splits that expose a number of traffic flows.

A. The "5G Ready" HW-Based Upper-PHY Split

The "5G-ready" evolved fronthaul, employing a solely HW-based upper-PHY split is shown in Fig. 2(a). No higher (5G RAN) protocol layers are included. Following the Forward Error Correction (FEC) encoder, the mapping function pack-etizes the "backhaul-like" data directly into Ethernet frames resulting in a single data flow. The resulting, encoded FEC blocks are 2040 Octets long necessitating the use of jumbo frames. The block size is a result of a signal processing parallelization approach, required so that the target data rate is achieved. Still, the resulting frame size is at the low end of jumbo frame sizes. The encapsulation format shown in Fig. 2(b) is based on the standard layer-2 (Ethernet) headers with additional custom headers based on the IEEE 1914.3 definitions [17]. The real-time CU processing is implemented in a Xilinx Virtex-7 ultrascale FPGA. The split interface employs 10GbE (Gigabit-Ethernet), while the supported IQ analog bandwidth is 2 GHz.

Fig. 3 shows benchmarking results for fronthaul data rate versus backhaul throughput. Two mapping modes have been tested: latency-optimized and overhead-optimized. For the former, idle data are used to complete the FEC blocks when no additional backhaul data is available. This method of operation is interesting as it reduces latency and provides a more deterministic input traffic characteristic, that is, the traffic flow characteristics become decoupled from the input packet arrival time. However, it is not efficient for intermittent "bursty" traffic with lower traffic loads and small packets. It can be seen that the available fronthaul capacity can be quickly utilized due to low overhead efficiency. However, a clear trend is shown that as the frame size increases, the efficiency approaches that of the overhead-optimized mode. As this mapping mode splits input frames among multiple FEC blocks, the traces do not
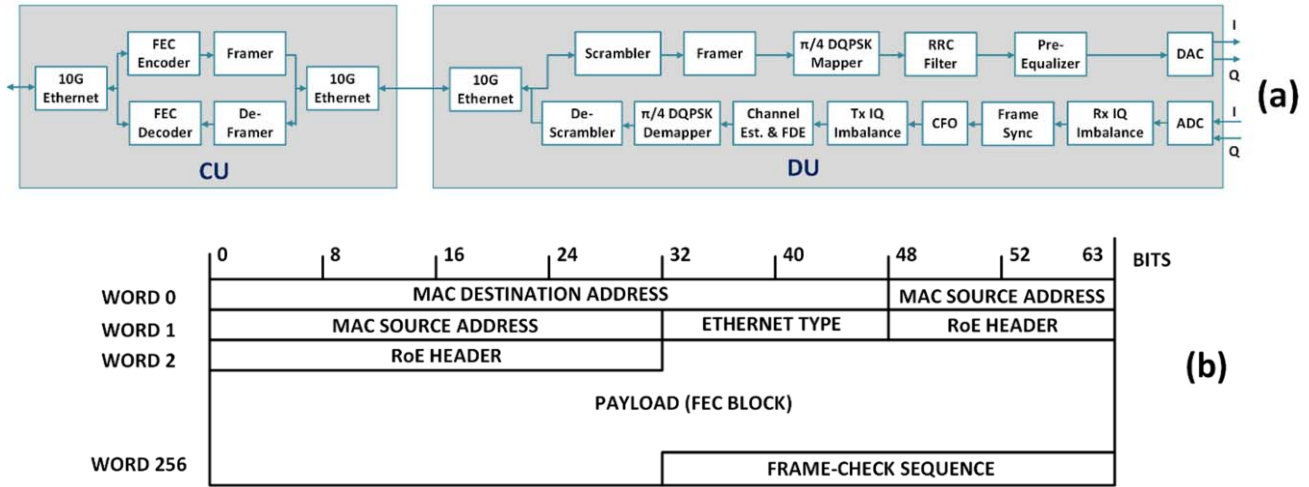
Fig. 2. (a) Real-time transceiver building blocks for the upper-PHY split. DQPSK, Differential Quadrature Phase Shift Keying; FDE, Frequency Domain Equalization; CFO, Carrier Frequency Offset correction; DAC, Digital-to-Analog; ADC, Analog-to-Digital. (b) The Ethernet-based header format for the upper-PHY split with the addition of IEEE 1914 (RoE)-compatible header.


Fig. 3. Benchmarking results for fronthaul data rate versus backhaul throughput for different input frame lengths, for the HW-based upper-PHY split.

necessarily show a linear dependence on load (as can be seen for the "random" trace for example).

The overhead-optimized mode may lead to larger latency but can provide higher statistical multiplexing gains. For this mode a linear dependence on load is obtained irrespective of frame size and data rate. In general, for both modes, the custom header's fixed size leads to an overhead that decreases as the frame size becomes larger but the overhead-optimized mode results in significantly better overall efficiency (as no idle data are used), with minimal improvement gains for frame sizes larger than 512 octets.

Fig. 4 shows the processing latency results for the overhead-optimized mapping scheme. The peaks in latency occur at very high data rates; approximately at 9 Gbps input traffic, when the available capacity is saturated. Two main trends are observed: the latency reduces with frame size due to a reduction in the processing per frame, and the latency reduces with data rate as the time required to fill an FEC block reduces. At approximately 2 Gbps input data rate, the latency approaches its lowest values.

With a potential latency of 10 to 20 μs at the CU, the system allows enough margin for the optical transmission (approximately 100 μs for 20 km) and the processing at the DU and networking equipment. However for such a system to be used efficiently in a HW offloading application, where the input data will be based on higher 5G layers (for example MAC PDUs), the input traffic characteristics have to be taken into account. The overhead-optimized approach is promising provided large 5G-type data rates are maintained (>1 Gbps), as it does not have a strong dependence on input frame size. This is important as a functional split that includes the entire LTE/5G protocol stack will generate/expose a number of traffic flows (including control primitives), as will be shown in subsection IIB, which will have significantly varying frame size. But, if such data rates are not present the latency may be significant (see the low-end of the data rate in Fig. 4) and has to be carefully considered. In such a case, the latency-optimized approach could be used. Balancing between these two approaches will become significant for pre-resource mapper splits where multiple flows are exposed over the fronthaul and in cases where large variations between user-plane throughputs are present. Both of these aspects will become clearer when considering a full protocol implementation as presented in the next subsection.

B. The SW-Based LLS Functional Split

The MAC/PHY functional split (Option 6) implementation includes the full LTE protocol layer stack in a software emulation environment. The mapping from LTE MAC Protocol Data Units (PDUs) to Ethernet frames is carried out by a *Fronthaul Interface Library (FIL)*, which is inserted between OpenAirInterface (OAI) [18] software modules that perform the LTE processing functions. Packetization is carried out using a raw Linux socket abstraction. The implementation allows the separation and different treatment of flows that are generated at the MAC/PHY boundary. These flows are in turn associated with a packet type that includes Downlink Control

Fig. 4. Benchmarking results for average processing latency versus backhaul throughput for the overhead-optimized mapping to FEC blocks.

Information (*packet_type_DCI*), Downlink Shared Channel (*packet_type_DLSCH*), System Information (*packet_type_SI*) and Random Access Response (*packet_type_RAR*). The preferred choice here was to employ a common encapsulation format for all packet types using standard layer-2 (Ethernet) headers (as in Subsection IIA). However, due to the different packet types, additional header fields are required. Unlike the upper-PHY split where a single flow was present, the MAC/PHY split generates multiple flows and exposes a number of inter-layer interfaces that need to exchange information (primitives).

The Ethernet payload section contains a number of header fields that are specific to the LLS split; these include LTE subframe and system frame numbers, packet-type IDs, length and FIL payload fields. The FIL receiver function is able to parse these header fields accordingly. The FIL payload section further includes packet-type specific fields, with examples for the DCI packet type shown in Table I.

The distinction between different packet types brings possibilities, which include using different CoS for different packet types and extending buffer time-outs for packet types that are assigned a lower CoS. Furthermore, it offers efficiency in the protocol primitive exchange over the fronthaul by using the DCI packet type in effect as a MAC/PHY *primitive carrier*. The DCI information also includes the user allocations for the next LTE subframe allowing the DU to setup its receive buffers. For CoS assignments, VLAN IDs are applied which can be used to either offer CoS differentiation for the different packet types or alternatively apply CoS bundling. The latter is used to assign a common CoS for all packet-types within the same subframe and is a sensible implementation as DU processing takes place once all packet types (including multiple DLSCH packet types for multiple users) have been received. Therefore, defining time-outs is an important split-specific design aspect as delayed (or dropped) packets may in turn cause large delays in the DU. With buffer time-outs, subframe timings can be maintained over the air. By using Sequence Numbers (SNs) for user allocations, the DU can keep track of which DLSCH packets have been received and

TABLE I
THE PAYLOAD SECTION FOR THE DCI PACKET_TYPE

PKT_DCI (Payload section)	
unsigned 8-bit integer	# of UE-specific DCI allocations = A
unsigned 8-bit integer	# of common DCI allocations = B
unsigned 8-bit integer	# of PDCCH symbols
15 Octets × (A+B)	DCI allocations [A+B]

Fig. 5. The Ethernet fronthaul testbed set-up for the Option-6 LLS. VLAN, Virtual Local Area Network; IP, Internet Protocol; GbE, Gigabit-Ethernet.

TABLE II
PACKET AND LTE SUBFRAME ONE-WAY END-TO-END LATENCIES AND LATENCY STANDARD DEVIATION

KPI	1 GbE	10 GbE
Average packet latency/ μs	45	43
Packet latency STD/ μs	9	6.9
Average subframe latency/ μs	54	52
Subframe latency STD/ μs	6.3	4.1

which ones must be ignored due to delays in the fronthaul. Then the FIL can insert nulls in the place of the missing user allocations and defer them to the corresponding HARQ retransmission process.

An experimental analysis of HARQ retransmissions due to fronthaul-induced latency variation and the transmission of nulls in place of missing user data has been carried out in [8], albeit for an Option-8 LLS.

A testbed set-up for benchmarking characterization of the LLS is shown in Fig. 5 with measurement results of latency shown in Table II. The LTE bandwidth is 5 MHz with three emulated UEs. The average packet size is 300 octets while the maximum (for packet_type_DLSCH) is 1000 octets, thus limiting the TB sizes to within the 1492-octet Ethernet standard Maximum Transmission Unit (MTU). The choice was made here to avoid the use of jumbo frames, which although offering higher overhead efficiency can lead to large delay variations for other fronthaul flows. The encapsulation overhead resulting from the FIL for the DLSCH packet type is in the order of 3% but is significantly higher for the other packet types.

The one-way packet latency (average of all the packet-types) for 1GbE is approximately 45 μs for a back-to-back (b2b) connection while the latency standard deviation (STD)
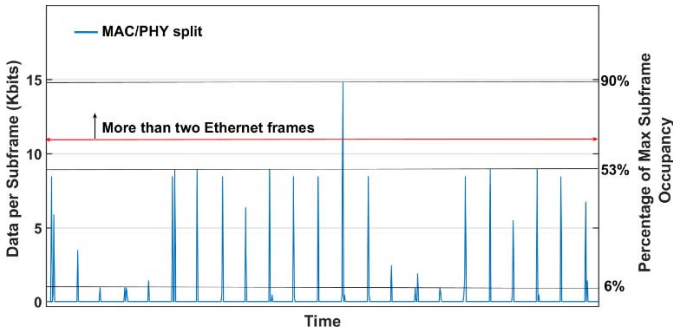
Fig. 6.　Data per subframe measurement results for the Option 6 LLS, for a bandwidth of 5 MHz, showing fronthaul occupancy varying with cell load, for a resolution of 1 ms.

is approximately 9 μs. The incorporation of 10GbE improves the latency by approximately 2 μs.

The subframe latency is not a KPI usually quoted in the literature but is important for fronthaul as it determines the buffer size at the DU (assuming that DU processing for a given subframe takes place only once all packets destined for that subframe have been received). The results in Table II show that the subframe latency is larger than the packet latency due to the fact that a subframe consists of a number of packet types. However, the STD is smaller and it is this that has to be taken into account when designing the buffer algorithm and potential buffer time-out implementation in the DU.

Fig. 6 shows a data per subframe measurement result for the LLS with a channel bandwidth of 5 MHz. The fronthaul occupancy scales mainly due to the DLSCH packet types (corresponding to the cell load), resulting in periods of available link capacity that can be used by other flows. Contrasting this with IQ transport and its resulting constant data rate, irrespective of cell load, the potential benefits in capacity provisioning through statistical multiplexing gains obtained with the LLS are clear. The right y-axis of Fig. 6 indicates the percentage occupancy of the LTE subframe (for a 5MHz bandwidth the maximum data per subframe is approximately 17 kbits). The aim here is to offer a view of subframe occupancy normalized to the maximum subframe capacity.

This variability in generated traffic has to be taken into account when employing HW-offloading such as in the upper-PHY split presented in subsection IIA. As certain subframes will require multiple Ethernet frames to be transported, with largely varying sizes, the overhead-optimized mapping would be the most efficient method of providing HW offloading, provided that high data rates (>1Gbps) are maintained. This is attested by the results of Figures 3 and 4 for the "random" trace corresponding to a random frame size for the input traffic. For smaller data rates including 4G-type ones (e.g. low-data rate DLSCH flows) and for control and primitive flows (DCI, SI etc.) the latency-optimized mapping would be beneficial with the highest efficiency obtained for packet types that are as close in size to the FEC blocks as possible. The smaller size packet-types (control and primitives) would result in low overhead efficiencies but their contribution in the total fronthaul data rate would be small. In addition to different

mapping regimes for control and user flows, users with bad channel conditions also experience low data rates for their DLSCH flow. In such cases, the latency-optimized approach will again be beneficial.

III. SYNCHRONIZATION AND TIME-SENSITIVE NETWORKING

While statistical multiplexing gains can help reduce over-provisioning of fronthaul links (and thus reduce operator costs), they also mean that latency variation (also termed Frame-Delay Variation, *FDV*, or packet-delay variation,*PDV*) from the aggregation/switching process will occur and has to be constrained to acceptable levels. PTP traffic will have the tightest constraint, so as not to produce significant timestamping errors. On the other hand, traffic arising from new split points will generally have reduced latency variation requirements but will require proper buffering to absorb such variation, especially with contention.

The IEEE P802.1CM (Time-Sensitive Networking for Fronthaul) standardization effort is in the process of selecting/adapting time-sensitive networking profiles for use in a bridged fronthaul network [19]. The standard aims to define the requirements for different base station functional decompositions (these are separated into *Classes*) and specifies two profiles for meeting class requirements, one employing strict priority (SP) scheduling, and the other frame pre-emption based on P802.1Qbu [20]. However, the standard does not take into account transport of a synchronization flow, for example, PTP and how this should be treated by a switch/aggregator, and instead assumes that network-wide synchronization is present through unspecified means (potentially through Global Navigation Satellite System (GNSS) and/or transparent PTP clocks).

The aim of this section is to compare experimental results using a novel aggregation approach, termed the *gap filling* aggregator, with theoretical results of other TSN schedulers, specifically when PTP has to be transported by the evolved fronthaul through non-PTP-aware switches/aggregators (i.e. there is a lack of transparent clocks). The experimental results are obtained from a converged fronthaul transporting a number of traffic flows with different traffic generation characteristics.

A. Switching/Bridging Scheduling Techniques

Assuming an arbitrary distribution of frames between High Priority (HP) and Low Priority (LP) flows, the peak node (switch or aggregator) latency, $T_{node,peak}$, is obtained from the maximum LP packet size and is given as

$$T_{node,peak} = T_f + T_s + T_{q,peak}, \qquad (1)$$

where T_f is the switch/aggregator processing (fabric) delay, T_s, is the port serialization delay for the HP frame and $T_{q,peak}$ is the peak queuing delay (based on the maximum frame size of the LP traffic flow) and is dependent on scheduler implementation. The corresponding minimum node latency is given as

$$T_{node,min} = T_f + T_s. \qquad (2)$$

Fig. 7. Different scheduling regimes and corresponding peak FDVs. T_s here corresponds to the worst case (i.e. maximum frame size for the background traffic). TAS: Time-Aware Scheduler; GP, Guard Period; TW, Time Window.


Fig. 8. Cumulative distribution functions of measurement results for IQ transport with contending background traffic of different frame sizes. NP, No Priority (single-queue); WRR, Weighted Round Robin; SP, Strict Priority.

contending flows (the single queue scheduler will attempt to balance the number of bytes selected for transmission from the input queues). Both mean latencies and STDs increase with background traffic frame size. By applying SP or WRR, the mean latency decreases and does so more significantly for the SP case.

As a result, a smaller per-subframe latency would have to be absorbed by the receive buffer of the CU/DU resulting in a smaller increase in end-to-end latency. However, the peak latency component remains bounded and given by the inset equation in Fig. 7, necessitating the use of pre-emption. It is evident from these results that while the use of larger frames (for example jumbo frames) is beneficial due to larger overhead efficiency, for multiple transported flows, in terms of overall end-to-end latency, smaller frame sizes lead to smaller per hop delays allowing for an increase in the network reach.

B. Gap-Filling: TSN Aggregation Approach

The TSN aggregation presented here employs gap-filling between HP frames [25]. The aggregator multiplexes HP traffic streams with LP streams by taking advantage of the inter-packet gaps between HP frames and using them to transmit LP frames. For this method to work, a deterministic delay is added to an outgoing HP stream, which is equal to the maximum transmission time of a LP frame. A gap detector is used to obtain the inter-packet gaps and a scheduler then chooses an LP frame that can fit within each inter-packet gap. Note that this approach does not require any additional (in-band or out-of-band) form of synchronization (as is the case with TAS for example).

The implementation of the 100GbE aggregator and theoretical delay budget are shown in Fig. 9. Table III summarizes benchmarking latency and latency variation measurement results for a generic traffic scenario, where a MTU of 16000 octets is used for the HP flow and a packet size of 9622 octets for two LP flows (Table III only shown the results for one of the LP flows). The HP results have a minimum latency of 14180 ns, and a maximum latency of 14316 ns resulting in a maximum latency variation of 136 ns.

The LP results have a minimum latency of 8769 ns and a maximum latency of 11177 ns resulting in a maximum latency variation of 2408 ns. It is important to note that the latency scales with MTU length proportionally.

Subtracting (2) from (1) results in the peak FDV, FDV_{peak}, for a given scheduler.

Fig. 7 shows a number of TSN scheduling techniques and the corresponding expressions for peak FDV. The simplest scheduler is based on a single output queue. Note that this is the case when flows are not differentiated based on priority setting or when multiple flows have the same priority setting and therefore share the same output queue. A comparison based on experimental results of a number of traditional schedulers for IQ traffic in an Ethernet fronthaul has been carried out in [21], while the Time-Aware Shaper (TAS) is based on [22] and simulation results have shown that in principle such a scheduler can eliminate contention-induced FDV [23], [24].

Fig. 8 shows cumulative distribution functions of measured Ethernet frame inter-arrival delays in the fronthaul testbed of Fig.5, where contention occurs between two streams comprising LTE (IQ based)-carrying Ethernet frames and background Ethernet traffic, for different background traffic frame sizes. The baseline value (no contention) of the distributions is approximately 32 µs. The NP (No Priority) trace corresponds to the single-queue scheduler while the remaining traces are used to compare SP and Weighted-Round Robin (WRR) scheduling. The peak deviation for the single-queue case can vary based on the distribution of relative frame size of the
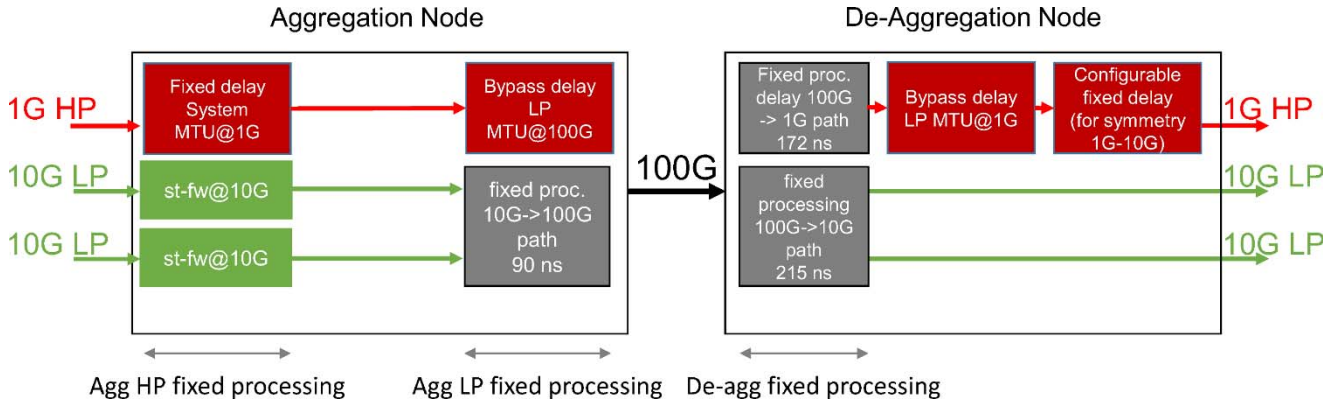
Fig. 9. Schematic of gap-filling 100GbE aggregator and theoretical latency budgets.

TABLE III

BENCHMARK MEASUREMENT RESULTS OF LATENCY AND LATENCY
VARIATION FOR THE GAP-FILLING AGGREGATOR, USING AN
MTU SIZE OF 16000 OCTETS FOR HP AND A PACKET
SIZE OF 9622 OCTETS FOR LP TRAFFIC

KPI	HP	LP
Minimum Latency (ns)	14180	8769
Average Latency (ns)	14246	8991
Peak Latency (ns)	14316	11177
Peak Latency Variation (ns)	136	2408

Fig. 10. The converged Ethernet evolved fronthaul with gap-filling TSN aggregation and PTP transport. GM, Grand-Master; GNSS, Global Navigation Satellite System; SMF, Single-Mode Fiber; Map., Mapper; RRH, Remote Radio Head.

C. Overall Aggregation: The Evolved Ethernet Fronthaul

Fig. 10 shows the Ethernet fronthaul testbed that aggregates a number of traffic flows. In addition to the two LLSs presented in Section II, the testbed includes a number of additional flows that are treated as background traffic. These include Option-8 CPRI-over-Ethernet (CPRIoE) with constant-packet rate characteristics, generic Ethernet background traffic with bursty characteristic, EPC backhaul traffic to an LTE femtocell (bursty traffic characteristics) and Option-8 generic IQ over Ethernet traffic (bursty traffic characteristic). Access to the gap-filling aggregator is provided through 10GbE ingress/access ports while aggregation takes place on a 100GbE trunk. The utilization of the trunk is approximately 22%.

We note that the effects of aggregation on the performance of the functional split and background flows was negligible


Fig. 11. Packet-delay variation of PTP messages using the set-up of Fig. 10.

(no, HARQ retransmissions for the Option-6 split, no loss of synchronization for the upper-PHY and CPRIoE splits) due to the large link rates. We defer a more detailed investigation of the performance effects on these flows to future work, especially at higher trunk utilisations.

Finally, a PTP flow is provided through a PTP grandmaster (GM) and a GNSS-disciplined OSCILLOQUARTZ OSA 5410 time source, while a similar unit is employed as the PTP slave clock at the receiver side.

To precisely measure the PDV of the PTP sync and delay request messages, a VIAVI MTS-5800 was used as a tester with a separate GNSS input to analyze the PTP stream broadcast to the 10GbE egress port at the DU, as shown in Fig. 10.

Measured PDV results for a PTPv2 flow are shown in Fig. 11. Within a 19.5 hour test duration, the average PDV for both sync and delay request messages was less than 100 ns. From this total value, approximately 50 ns and 27 ns are attributed to the 10GbE MAC/PHY and 100GbE MAC/PHY, respectively.

In fact, the additional PDV introduced by the aggregator firmware is only approximately 19 ns.

In addition to timing synchronization, frequency syntonization is equally important. A mechanism compatible to SyncE has been implemented to forward the clock from the

TABLE IV

SUMMARY OF MEASURED AND THEORETICAL KEY-PERFORMANCE INDICATORS FOR THE FUNCTIONAL SPLITS AND TSN REGIMES

KPI	Requirement	4G SW-based Option-6 split	5G Upper-PHY split	TSN Schedulers				
Frequency accuracy (ppb)	2[1]	N/A	0.2	**WRR**	**SP**	**Frame preemption**	**TAS**	**Gap-filling aggregator**
One-way latency (μs)	220[1] 75 (CoMP)	43 (53)[8]	10-20[4,9] (20-30)[8,9]					
Peak packet jitter (for HP flow) (ns)[3]	64[2]	N/A	N/A	>SP	1300	224	≥100[7]	<100 ns
Experimental data available (@ 100GbE aggregation)				NO	NO	NO	NO	YES
In-band timing protocol requirement[5]				NO	NO	NO	YES	NO
Requirement for global scheduling[6]				NO	NO	NO	YES	NO
Slice isolation				Softer	Softer/soft	Soft	Soft/Hard[7]	Hard

[1]From [4]; [2]From [23]; [3]Target application is PTP transport at 10GbE, assumes 50 ns node processing-induced jitter and 1500 Octet LP frames; [4]For CU processing assuming efficient per-flow mapping; [5]Required for the operation of the scheduler; [6]Note that some complex network slicing scenarios (e.g. dynamic CoS reconfiguration) would potentially require network wide schedulers; [7]Depends on guard period size used [24]; [8]With gap-filling aggregator; [9]Assuming efficient per-flow mapping.


Fig. 12. Histogram of frequency deviation in ppb for the HW-based upper PHY-split.

upper-PHY split CU to its respective DU. A histogram showing the deviation of the radio unit clock compared with the data source is depicted in Fig 12. The maximum deviation observed is $+/-0.2$ppb, which is well within the limits of frequency accuracy expected for 5G [4].

Table IV summarizes the measured and theoretical KPIs for the LLS implementations and scheduling techniques. The one-way latency for the Option-6 split through the aggregator is approximately 53 μs, which is within both quoted latency requirement figures, including the more stringent 75 μs one, for CoMP, but leaving a very small margin for fiber delays. Use of HW-offloading using for example, the upper-PHY split platform presented here, is expected to reduce this latency figure considerably while supporting much higher data rates.

The gap- filling aggregator approach presented here is the only technique that has available experimental data at 100GbE aggregation showing performance that is promising for meeting the stringent latency variation requirements when PTP is transported. By comparison, the other scheduling techniques considered only provide improvements in average

latency and latency variation with varying levels of success. Thus, the simpler techniques such as SP and WRR may offer improvements when only softer/soft isolation between slices is necessary. TAS may in theory eliminate contention-induced delay variation but would require the implementation of a global scheduler (when the number of aggregation/switching nodes increases), and it is still an open question whether such a scheduler can operate at high networking rates (e.g. 40GbE and 100GbE).

IV. CONCLUSION

Two functional splits are assessed both individually and in unison, in terms of latency, latency variation and Ethernet mapping efficiency. A hardware-based upper-PHY split able to provide 5G-type data rates is assessed as a potential hardware-offloading platform. Then a software-based 4G Option-6 split that includes the full LTE protocol stack is used to obtain a number of packet types corresponding to exposed user- and control-plane primitive flows.

It is shown that in order to obtain high statistical multiplexing gains and significant overhead efficiency, the future fronthaul needs to treat the newly exposed traffic flows separately, applying mapping on a per-flow basis based on their respective traffic generation characteristics (e.g. large data rate or channel condition differences between users). Different mapping is also required between user- and control-plane flows. Such an approach is efficient but will lead to design complications in multi-operator/multi-user slicing scenarios so must be taken into account in future research. Potential alternative methods such as efficient aggregation of flows exist, but these may complicate slice isolation and class-of-service differentiation.

As the latency/latency variation performance of the evolved fronthaul needs to be treated holistically, in addition to functional split characterization, time-sensitive networking was considered. Experimental results for a gap-filling aggregator over a converged Ethernet fronthaul operating at 100 GbE rates and transporting functional split data in addition to background

traffic flows, show state-of-the art sub-100 ns latency variation for a PTP flow. Through a comparison with theoretical predictions for other scheduling techniques, the aggregator approach is shown to be the most promising in meeting stringent future mobile network requirements, and could additionally be employed in hard slice isolation scenarios.

ACKNOWLEDGMENT

REFERENCES

[1] *Part 3: Carrier Sense Multiple Access With Collision Detection (CSMA/CD) Access Method and Physical Layer Specifications Amendment: Media Access Control Parameters, Physical Layers, and Management Parameters for Subscriber Access Networks*, IEEE Standard 802.3ah, Sep. 2004. [Online]. Available: http://www.ieee802.org/21/doctree/2006_Meeting_Docs/2006-11_meeting_docs/802.3ah-2004.pdf

[2] *Connectivity Fault Management*, IEEE Standard 802.3ag, 2007. [Online]. Available: http://www.ieee802.org/3/ag/index.html

[3] Metro Ethernet Forum. (Jan. 2006). *MEF 16: Ethernet Local Management Interface (E-LMI)*. [Online]. Available: http://www.mef.net/Assets/Technical_Specifications/PDF/

[4] iCIRRUS. (Jul. 2016). *D3.2: Preliminary Fronthaul Architecture Proposal*. [Online]. Available: http://www.icirrus-5gnet.eu/category/deliverables/

[5] N. J. Gomes, P. Chanclou, P. Turnbull, A. Magee, and V. Jungnickel, "Fronthaul evolution: From CPRI to Ethernet," *Opt. Fiber Technol.*, vol. 26, pp. 50–58, Dec. 2015.

[6] International Telecommunication Union, document ITU G.8264, *Distribution of Timing Information Through Packet Networks*, May 2014. [Online]. Available: http://www.itu.int/rec/

[7] International Telecommunication Union, document ITU G.8275.1, *Precision Time Protocol Telecom Profile for Phase/Time Synchronization With Full Timing Support From the Network*, Jul. 2014. [Online]. Available: http://www.itu.int/rec/

[8] P. Assimakopoulos, M. K. Al-Hares, and N. J. Gomes, "Switched Ethernet fronthaul architecture for cloud-radio access networks," *OSA/IEEE J. Opt. Commun. Netw.*, vol. 8, no. 12, pp. B135–B146, Dec. 2016.

[9] CPRI. (Oct. 2015). *CPRI Specification V7.0, Interface Specification*. [Online]. Available: http://www.cpri.info/spec.html

[10] *Study on New Radio Access Technology; Radio Access Architecture and Interfaces (Release 14)*, document 3GPP TR 38.801 V0.4.0, Aug. 2016. [Online]. Available: http://www.3gpp.org/DynaReport/38-series.htm.

[11] CPRI. (Aug. 2017). *eCPRI Specification V1.0, Interface Specification*. [Online]. Available: http://www.cpri.info/spec.html

[12] iCIRRUS. (Jul. 2015). *D2.1: iCIRRUS Intelligent C-RAN Architecture*. [Online]. Available: http://www.icirrus-5gnet.eu/category/deliverables

[13] G. S. Birring, P. Assimakopoulos, and N. J. Gomes, "An Ethernet-based fronthaul implementation with MAC/PHY split LTE processing," in *Proc. Global Commun. Conf.*, Singapore, Dec. 2017, pp. 1–6.

[14] G. Mountaser, M. L. Rosas, T. Mahmoodi, and M. Dohler, "On the feasibility of MAC and PHY split in cloud RAN," in *Proc. IEEE Wireless Commun. Netw. Conf. (WCNC)*, San Francisco, CA, USA, Mar. 2017, pp. 1–6.

[15] N. Makris, P. Basaras, T. Korakis, N. Nikaein, and L. Tassiulas, "Experimental evaluation of functional splits for 5G cloud-RANs," in *Proc. Global Commun. Conf.*, Paris, France, May 2017, pp. 1–6.

[16] K. Miyamoto, S. Kuwano, T. Shimizu, J. Terada, and A. Otaka, "Performance evaluation of Ethernet-based mobile fronthaul and wireless CoMP in split-PHY processing," *J. Opt. Commun. Netw.*, vol. 9, no. 1, pp. A46–A54, 2017.

[17] *Standard for Radio Over Ethernet Encapsulations and Mappings*, IEEE Standard P1904.3, 2016. [Online]. Available: https://standards.ieee.org/develop/project/1904.3.html

[18] OpenAirInterface. *OpenAirInterface Software Alliance*. Accessed: Oct. 10, 2018. [Online]. Available: http://www.openairinterface.org/

[19] *Time-Sensitive Networking for Fronthaul*, IEEE Standard P802.1CM, 2018. [Online]. Available: http://www.ieee802.org/1/pages/802.1cm.html

[20] *Frame Preemption*, IEEE Standard 802.1Qbu, 2015. [Online]. Available: http://www.ieee802.org/1/pages/802.1bu.html

[21] M. K. Al-Hares, P. Assimakopoulos, S. Hill, and N. J. Gomes, "The effect of different queuing regimes on a switched Ethernet fronthaul," in *Proc. IEEE Int. Conf. Transp. Opt. Netw. (ICTON)*, Trento, Italy, Jul. 2016, pp. 1–4.

[22] *Enhancements for Scheduled Traffic*, IEEE Standard 802.1Qbv, 2016. [Online]. Available: http://www.ieee802.org/1/pages/802.1bv.html

[23] T. Wan and P. Ashwood-Smith, "A performance study of CPRI over Ethernet with IEEE 802.1Qbu and 802.1Qbv enhancements," in *Proc. Global Commun. Conf.*, San Diego, CA, USA, Dec. 2015, pp. 1–6.

[24] M. K. Al-Hares, P. Assimakopoulos, D. Muench, and N. J. Gomes, "Modeling time aware shaping in an Ethernet fronthaul," in *Proc. Global Commun. Conf.*, Singapore, Dec. 2017, pp. 1–6.

[25] R. Veisllari, S. Bjornstad, J. P. Braute, K. Bozorgebrahimi, and C. Raffaelli, "Field-trial demonstration of cost efficient sub-wavelength service through integrated packet/circuit hybrid network," *J. Opt. Commun. Netw.*, vol. 7, no. 3, pp. A379–A387, 2015.

Philippos Assimakopoulos (S'09–M'16) received the B.Eng. degree in electronic engineering from the University of Bath, Bath, U.K., in 2003, and the M.Sc. degree in broadband and mobile communication networks and the Ph.D. degree in electronic engineering from the University of Kent, Canterbury, U.K., in 2007 and 2012, respectively.

He is currently with the Communications Research Group, University of Kent. He has participated in various EU FP7, Horizon 2020, and U.K. EPSRC research projects. His research interests include distributed antenna systems, low-cost microwave radio-over-fiber networks for indoor and outdoor applications, and the design of cloud-radio access network for 4G and 5G applications.

Jim Zou (S'11–M'16) received the B.Eng. degree in communication and information engineering and the M.Sc. degree in electrical circuits and systems from Shanghai University, China, in 2008 and 2011, respectively, and the Ph.D. degree from the Eindhoven University of Technology, The Netherlands, in 2015. He conducted research work in the area of broadband indoor fiber-wireless networks with the Electro-Optical Communication Group, COBRA Research Institute, Eindhoven University of Technology.

He is currently a Senior Engineer with the Advanced Technology Department, ADVA Optical Networking SE, where he is involved in the internal prototype development related to the next-generation optical access. He has been participating in various EU FP7 and Horizon 2020 research projects.

Kai Habel received the Diploma degree in electrical engineering from Technical University Berlin in 2001.

He was a Research Assistant with the Fraunhofer HHI, Berlin, Germany, where he focuses on optical metro, access and in-house networks, and optical wireless systems. His expertise ranges from physical layer transmission technology to media access and higher layer protocols.

Jörg-Peter Elbers (M'01) received the Diploma and Dr.-Ing. degrees in electrical engineering from Technical University Dortmund, Germany, in 1996 and 2000, respectively.

From 1999 to 2001, he was with Siemens AG, where he was the Director of network architecture with the Optical Networks Business Unit. Afterwards, he was the Director of technology with the Optical Product Unit, Ericsson. In 2007, he joined ADVA Optical Networking, where he is currently a Senior Vice President of advanced technology, standards, and IPR. He has authored or co-authored over 100 technical presentations and three book chapters, and holds 15 patents. He is the Head of the German VDE ITG Expert Committee on communication technologies. He serves on the board for the European Technology Platforms on Photonics and Networks.

Nathan J. Gomes (M'92–SM'06) received the B.Sc. degree in electronic engineering from the University of Sussex, Sussex, U.K., in 1984, and the Ph.D. degree in electronic engineering from University College London, London, U.K., in 1988.

From 1988 to 1989, he was a Royal Society European Exchange Fellow with ENST, Paris, France. Since 1989, he has been with the University of Kent, Canterbury, U.K., where he is currently a Professor of optical fiber communications. His current research interests include fiber-wireless access systems and networks, fronthaul, and radio-over-fiber technology. He was the TPC Chair of the IEEE International Conference on Communications (ICC 2015), London.

Volker Jungnickel (M'99) received the Doctorate degree in physics from Humboldt University Berlin in 1995 and the Habilitation degree in communications engineering from Technical University Berlin in 2015.

In 1997, he joined the Fraunhofer Heinrich Hertz Institute, Berlin, where he contributed to high-speed optical wireless communications, the first 1 Gb/s mobile radio link, the first real-time trials of local thermal equilibrium (LTE), and first coordinated multipoint trials for LTE. Since 2002, he has been teaching courses on multiple-input multiple-output and adaptive transmission with Technical University Berlin. He currently leads the Metro-, Access and In-house Systems Group, where he focuses on short-range optical technologies. He also serves as the Chair for the IEEE 802.15.13 Task Group on Multi-Gbit/s Optical Wireless Communications and a Technical Editor for the IEEE 802.11bb Project on Light Communication.