

Effects of Contention and Delay in a Switched Ethernet Evolved Fronthaul for Future Cloud-RAN Applications

Philippos Assimakopoulos, Gurtej S. Biring, Nathan J. Gomes

Communications Research Group, University of Kent, Canterbury, UK
(P. Assimakopoulos@kent.ac.uk)

Abstract: *A Switched-Ethernet fronthaul transporting data generated by a Long-Term Evolution software base station with a MAC/PHY functional split is presented. Contention effects arising from the Ethernet fronthaul and the effects of priority-based scheduling are characterised.*

Introduction

The implementation of different Long-Term Evolution (LTE) functional subdivisions, instead of the conventional centralised one, has been proposed as a means of reducing data rates over the fronthaul section of Centralised (or Cloud) Radio Access Networks (C-RANs)¹. Conventionally, centralised processing involves transportation of the digitised In-phase and Quadrature (IQ) radio samples, predominantly using the specifications of the Common Public Radio Interface (CPRI)². However, with the advent of multiple antenna techniques (e.g. massive multiple-input multiple-output, *mMIMO*), carrier aggregation and new 5th generation (5G) signal bandwidths, the transportation of these IQ signals using current technologies becomes challenging³, hence the interest in different functional subdivisions (or “splits”) which can reduce the resulting data rate requirements, considerably. A number of standardisation bodies are in the process of developing consensus over the choice of split points. The choice is by no means straightforward, as each split point comes with advantages and disadvantages⁴. Deciding factors will relate mainly to latency and latency variation constraints, ease of implementation, data rate reduction and support for advanced joint processing techniques (e.g. coordinated multi-point, CoMP). A promising split that offers a good overall balance is one at or close to the LTE medium access control (MAC) and Physical (PHY) layer interface³. The use of Ethernet has recently been proposed for the future fronthaul³. Due to its ubiquitous nature, Ethernet can enable structural and operation convergence for operators leading to a converged infrastructure termed the x-haul (x standing for back, front and mid). However moving to a packet based networking approach also comes with significant challenges. Latency and latency variation will need to be controlled, while synchronisation features (not part of native Ethernet) will need to be provided. Promising technologies for the latter are the Precision-Time Protocol (PTP) and

Synchronous-Ethernet (SyncE) with a combination of both likely to be necessary for 5G application requirements. With regards to latency, a number of priority-based scheduling schemes for the fronthaul are under standardisation⁵. However, these are currently focusing on the centralised approach with requirements based on CPRI. Equivalent requirements will be required for fronthauls transporting functional split data.

Limited work has been carried out with physical implementations of functional splits and this only including limited parts of the LTE PHY layer functionality⁶.

In this paper, we present a testbed for an Ethernet fronthaul transporting signals arising from an LTE split at the MAC/PHY interface. The LTE functionality runs in a software emulation environment and includes *all* eNodeB protocol layer functionality. Contention and the effect of applying priority-based scheduling are measured.

Test-Bed

The testbed presented in this paper is used to demonstrate certain principles of advanced 4G and 5G fronthaul networks shown in Fig. 1(a). The future fronthaul will transport traffic from multiple baseband processing unit (BBU)/digital unit (DU) pools. The BBU/remote unit (RU) pairs will form part of the legacy/centralised splits while DU/RU pairs will transport evolved fronthaul data. The BBU/DU pools will employ hardware abstraction techniques with flexible processing running on general-purpose processing modules. Multi-operator and multi-tenancy scenarios will be possible by assigning a subsection of the RAN and networking processing/resources to tenants with different quality-of-service specifications. In Fig. 1(a) for example, RU1 is serviced by “bundling” its packet-types (or data flows) into a single traffic class of service (CoS). RU 1s flows will contend for access in trunk links within the fronthaul with multiple other flows arising from both legacy and evolved fronthauls and belonging to different users/tenants. Thus, the

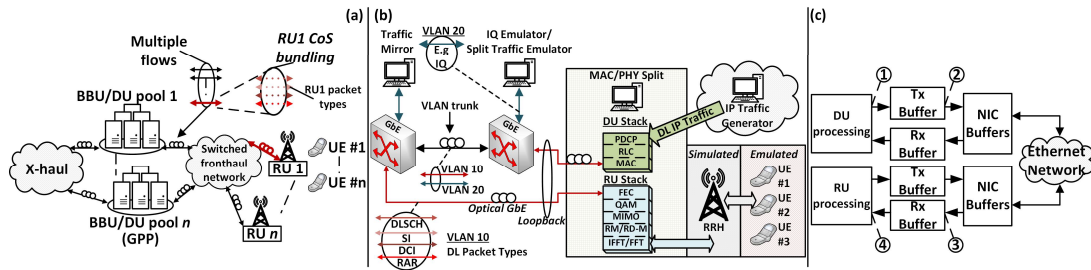


Fig. 1: (a) Advanced 4G and 5G fronthauling principles, (b) The Ethernet fronthaul testbed and (c) High-level view of the buffering stages and measurement interface points. CoS, Class-of-service; GPP, General-purpose processing; RRH, Remote radio head; NIC, Network interface card; IP, Internet protocol

contribution of RU1's data rate into the total utilisation of trunk links within the fronthaul can be very low. The testbed is shown in Fig. 1(b). The RAN processing modules are based on the open source OpenAirInterface software environment⁷. The functional split is implemented at the MAC/PHY interface (before/after the forward error correction in the downlink/uplink direction). The DU implements all LTE layers down to and including the MAC layer (NAS, RRC, PDCP, RLC, MAC) while the RU implements the PHY layer. Internet protocol (IP) traffic is directly fed into the PDCP layer. The mapping between the MAC and PHY layers is carried out through a *fronthaul interface library (FIL)*. The FIL performs all the necessary mapping/abstraction functions between the entities that are being split. It encapsulates the different flows using appropriate encapsulation formats. The produced flows are identified by *packet-type* and include system information (SI), downlink shared channel (DL-SCH), random access response (RAR) and downlink control information (DCI). The latter in addition to carrying the cell DCI, that is transmitted over the air at the start of each subframe (1 ms transmission time interval, *TTI*), also acts as a MAC/PHY interface primitive carrier. The RU uses the information in the packet to set up its own state variables, based on the number of allocations and the number of packets expected for the current radio subframe. The receive (Rx) buffers are then configured appropriately to accommodate the different packets. In the uplink, for simplicity, a single packet-type is used to aggregate all data flows. A traffic generator is used to produce background traffic configured as one that would be produced by centralised splitting (albeit through generic IQ transport and not CPRI based) or by additional splits close to the MAC/PHY interface (bursty traffic). The receiver is a traffic mirror that reflects incoming traffic back to the traffic generator. The MAC/PHY split and background traffic flows are logically separated through virtual local area network identifiers (VLAN IDs) and are made to contend in the optical trunk link formed between

two standard GbE (gigabit Ethernet) switches. All links are optical 1 GbE Ethernet.

Measurement Results

For the measurement results presented in this section, three UEs are attached to the network with a set application data rate per UE of 1.7 Mbps. The total DL-SCH allocation per subframe is divided to approximately 1000 Octets and is divided amongst the different UEs (note that DL-SCH is the carrier of user-plane data). Fig. 2 shows the per-packet (including all packet-types) and per-subframe mean and standard deviation (STD) delays with bursty background traffic and two different packet sizes, here set at a constant 200 Mbps rate, so increasing burst size results in fewer (but larger) bursts in a given time interval. Fig. 3 shows the same results but with constant packet-rate background traffic, at a fixed data rate of 400 Mbps. The per-packet delay is measured between interface points 1 and 3 in Fig. 1(c). The per-subframe delay takes into account the processing time of each subframe and is measured at the same interface points but taking into account the first packet in each subframe (a DCI packet-type) at interface point 1 and the last packet for that same subframe (DL-SCH, SI or RAR packet types) at interface point 3. The subframe delays are obviously larger than per-packet delays, with both traces following the same trend as burst size is increased. However, the opposite behaviour is seen in the STD results. The subframe delays have smaller STD than the per-packet results. This is the case even for the baselines (i.e. without background traffic) and is a result of the packet-types having different packet lengths (and therefore different serialisation delays). In addition, the traces of per-packet and per-subframe STDs converge with larger burst size and converge faster for the larger packet sizes. Similar trends are observed for the CPR results of Fig. 3. The reason behind this behaviour is that contention in the network will tend to absorb delays in between the packets that make up the subframe. Fig. 4 shows delay results of the individual packet types and the subframes with and without priority-based

scheduling using the strict priority (SP) algorithm for a background traffic packet size of 2000 Octets. For the subframe results, the measurement methodology is changed, now taking into account the first packet in each subframe (a DCI packet-type) at interface point 3 and the last packet for that same subframe at interface point 3. The baseline for this measurement is 14 μs while for the DLSCH and DCI packet-types it is 65 μs and 51 μs respectively. Note that SP has recently attracted attention as it is being considered in fronthaul standardization⁵. The results show that in order for SP to have a significant effect, the utilisation of the trunk and the burst size have to be increased considerably. It can be seen for example that for the smaller burst size (50), SP results in slight improvement and that only at 90% utilization. A similar trend is observed for the 4000 Octet case in Fig. 5. Note also that in general, a reduction in packet-type delays results in subframe delay reduction, but the levels of reduction may not be equivalent.




Fig. 2: Mean and STD of the per-packet and per-subframe delays versus burst size, with bursty background traffic at 20% utilization, for two different packet sizes




Fig. 3: Mean and STD of the per-packet and per-subframe delays versus packet size with constant packet-rate background traffic at 40% utilization




Fig. 4: Comparison of individual packet-type and subframe delays versus trunk utilization with and without strict priority for two different burst sizes, for 2000 Octet packets

Conclusions

A characterisation of contention and priority-based scheduling (employing strict priority, SP) effects in an evolved Ethernet fronthaul is




Fig. 5: 4000 Octet results. Comparison of individual packet-type and subframe delays versus burst size.

presented for the first time. It is shown that contention and implementing SP affect the individual packet-type and LTE subframe delays to different extents. This is a result of the injected delays absorbing intra-subframe delays. Characterising subframe delays is important for future fronthaul implementations where the buffer management in the RUs will have to handle late intra-frame packets, a crucial requirement for maintaining over-the-air subframe timings. It is shown that when the functional split traffic contribution to the trunk utilisation is very low, SP has a positive (albeit limited) effect only when the background traffic has very high data rate and very large burst size.

Acknowledgements

This work was carried out within the framework of the European Union's Horizon 2020 research and innovation programme under grant agreement No 644526 (iCIRRUS project) and EPSRCs "Towards an Intelligent Information Infrastructure (T13)" programme (NIRVANA project). Philippos Assimakopoulos acknowledges the funding by the NIRVANA project. Data used in this work is stored in Kent Academic Repository (<https://kar.kent.ac.uk/>).

References

- [1] U. Dötsch et al, "Quantitative analysis of split base station processing and determination of advantageous architectures for LTE," in *Bell Labs Tech. J.*, vol. 18, no. 1, pp. 105-128, June 2013
- [2] CPRI (Oct. 2015), "CPRI Specification V7.0, Interface Specification" [Online]. Available: <http://www.cpri.info/spec.html>.
- [3] iCIRRUS (July 2015), "D2.1: iCIRRUS intelligent C-RAN architecture," [Online]. Available: <http://www.icirrus-5gnet.eu/category/deliverables/>
- [4] 3GPP, "Study on New Radio Access Technology; Radio Access Architecture and Interfaces (Release 14)," 3GPP TR 38.801 V0.4.0, Aug. 2016 [Online]. Available: <http://www.3gpp.org/DynaReport/38-series.htm>
- [5] "Time-Sensitive Networking for Fronthaul," IEEE Standard P802.1CM [Online]. Available: <http://www.ieee802.org/1/pages/802.1cm.html>
- [6] K. Miyamoto, S. Kuwano, T. Shimizu, J. Terada, A. Otaka, "Performance Evaluation of Ethernet-based Mobile Fronthaul and Wireless CoMP in Split-PHY Processing," in *OSA/IEEE J. Optical Commun. and Netw.*, vol: 9, no. 1, pp. A46-A54, Jan. 2017
- [7] "OpenAirInterface (OAI)", OpenAirInterface Software Alliance [Online]. Available: <http://www.openairinterface.org/>