

Experimental Demonstration of SDN-controlled Variable-rate Fronthaul for Converged LTE-over-PON

Pedro Alvarez, Frank Slyne, Christian Bluemm,
Johann M. Marquez-Barja*, Luiz A. DaSilva, Marco Ruffini

CONNECT Centre, Trinity College Dublin, Ireland. *University of Antwerp - imec
{pinheirp,fslyne,blummc,dasilval,marco.ruffini}@tcd.ie, johann.marquez-barja@uantwerpen.be

Abstract: We introduce the concept of variable-rate fronthaul and provide experimental validation over PONs. Our SDN controller dynamically modifies the wireless cell bandwidth depending on load, thus varying the fronthaul rate with sub-second end-to-end reconfiguration times.

OCIS codes: (060.4251) Networks, assignment and routing algorithms, (060.2330) Fiber optics communications

1. Introduction

Cloud Radio Access Network (C-RAN) is a key technology for massive cell densification in future mobile networks. A main advantage is the simplification of the Remote Radio Head (RRH) functionality, which decreases its size, its power consumption, and its overall cost of ownership. Another advantage is the centralization of Base Band Unit (BBU) processing, which enables resource sharing, through BBU pooling, and inter-cell and multi-user coordination. C-RAN can adopt different types of processing fragmentation, i.e. split the signal processing and protocol functionality between the RRH and the BBU [1] in different ways. This has a direct impact on the resource and performance requirements for the link between RRH and BBU. Fronthaul typically refers to a functional split where the RRH performs only RF signal up-/down-conversion, digital-/analogue-conversion, and basic signal conditioning. All the baseband processing is carried out at the BBU. With this solution, the BBU-RRH link exchanges pure or compressed I/Q-data under a conventionally fixed data rate, independent of the actual demand of the mobile subscriber. This approach accounts for the maximum potential demand, which translates into a maximum transmission rate. While Passive Optical Networks (PONs) could in principle provide cost-effective fronthaul to multiple RRHs [2], by multiplexing multiple cells over each wavelength channel, the fixed fronthaul rate highly reduces such advantages, as it eliminates the ability to statistically multiplex across the fronthaul streams.

We propose the concept of *adaptive variable-rate fronthaul*, enabled by Software Defined Network (SDN) and statistical Time Division Multiplexing (TDM). The SDN controller interacts with the BBU to monitor the cell usage, and adapts the cell wireless bandwidth accordingly. A reduction in the wireless bandwidth requirement, for example when the cell usage is low, triggers a reduction of the fronthaul sampling rate, which reduces the required capacity over the PON. The SDN controller coordinates such capacity adaptation between the BBU, the RRH, and the PON, so that any freed-up capacity can be re-used by other lower-priority services (e.g., residential broadband). Our approach, in addition to reducing the fronthaul bandwidth requirements over the PON, enables the controller to coordinate spectrum reuse across multiple cells, with potential benefits to spectral efficiency and the performance observed by cell edge users. In this paper, we demonstrate the performance of an adaptive variable-rate fronthaul using a testbed facility that converges C-RAN, LTE and PON technologies. Our experimental results show dynamic resource reconfiguration times for the converged fixed-mobile network of around 750 ms.

2. Experimental Setup

Our testbed setup is shown in Fig. 1. The PON system shares two downstream links, one serving an LTE PHY downstream fronthaul link, and the other a lower priority traffic, representing residential users and called background link here. A Metro-Access switch (Pronto 3270 Openflow) separates the two links logically via VLAN tags, before aggregating them for transmission. Aggregation works by statistical TDM, according to commands from an SDN controller. The result, a combined downstream, is passed on to the OLT, which translates VLAN tags into the XGS-PON addressing format XGEM. After transmission over fibre, the stream reaches a single FPGA board, which hosts both ONU implementations. They are assigned to different 10G Ethernet ports at the customer side. Depending on the XGEM port identifier, the traffic is either routed to the port connected to the RRH or to the one emulating other PON users




Fig. 1: Experimental setup

(i.e., the background traffic). The PON OLT and two ONUs are implemented on Xilinx FPGA development boards VC709 [3] and comply with the ITU-T XGS-PON standard to all parts of interest [4]. Implementing two ONUs on a single device was a design choice due to missing burst-mode reception capabilities of the VC709 FPGA boards, which is a prerequisite for multipoint-to-point upstream operations. Nevertheless, our setup implements a classical PON, with the only difference that the multiplexing of upstream ONUs data is carried out in the FPGA rather than in the fibre. This does not affect our results, as we do not aim to demonstrate physical layer transmission, but rather the dynamic capacity (re-)allocation mechanism for mobile downstream. Previous work [5] has demonstrated the feasibility of carrying fronthaul-over-PON, by coordinating the DBA and BBU scheduling. Our work aims instead at dynamically optimising the mobile cell bandwidth, and consequently the fronthaul rate, to restore the statistical multiplexing gain, which is not available in legacy fronthaul systems.

The BBU part of the wireless system implements the eNodeB’s PHY layer, using the open source srsLTE software [6], compliant with 3GPP LTE release 8. We have modified the implementation of the eNodeB’s Physical Downstream Shared Channel (PDSCH), to add dynamic bandwidth reconfiguration. This is achieved by reconfiguring the number of Physical Resource Blocks (PRBs) used by the waveform, which modifies the bandwidth, sampling rate, FFT size, and other signal processing blocks. All PRBs are configured for 16 QAM modulation. The RRH part is implemented using an USRP X310 radio device, which directly connects to the ONU through a 10G Ethernet interface. In the downstream direction, the BBU sends I/Q samples over the PON towards the USRP board, which operates digital-to-analogue conversion and upconverts the signal to the 2.5 GHz ISM band. The LTE user equipment (UE) is represented by a second USRP X310 and linked to a server implementing the UE LTE PHY layer (also from srsLTE). In the upstream, the USRP (RRH) sends control messages to inform the BBU periodically about mobile network load. The BBU adds a Cell ID to these statistics before forwarding them to the real-time (Ryu-based) SDN controller. In order to guarantee the requisite fronthaul I/Q sample bandwidth, the SDN controller utilises meter tables to shape the fronthaul traffic according to its expected rates, listed in Table 1. These have been defined experimentally in our setup by correlating wireless channel bandwidth, number of PRBs, fronthaul rate and cell capacity, according to the modulation and coding scheme used. The SDN controller chooses the lowest fronthaul rate capable of serving the demanded cell capacity. Remaining capacity available over the PON can then be assigned to other users.

Table 1: Bandwidth and traffic schemes

Wireless Bandwidth	PRB Number	Fronthaul Rate	Max Cell Capacity
1.4 MHz	6	61 Mbps	1.8 Mbps
3 MHz	15	121 Mbps	4.584 Mbps
5 MHz	25	182 Mbps	7.736 Mbps
10 MHz	50	364 Mbps	15.264 Mbps
15 MHz	75	485 Mbps	22.92 Mbps
20 MHz	100	730 Mbps	30.576 Mbps

In our experiment, we reproduce a scenario where the data usage of a mobile user (Event 1, indicated in Fig. 1), changes over time. This change generates an decrease (or increase) in capacity at the BBU side, which is detected by the SDN controller (Event 2). The controller decreases (or increases) the mobile capacity by sending the BBU instructions to change the number of PRBs (Event 3), and simultaneously reconfigures the meter tables in the Openflow




Fig. 2: Measured Background/LTE switching




Fig. 3: Time diagram of switching events

switch accordingly to guarantee capacity to the fronthaul link as needed (Event 4). After the BBU updates its wireless transmission bandwidth, it broadcasts information about the bandwidth used by the cell through Master Information Block (MIB) messages, as per LTE specifications. The UE listens to these messages and reconfigures its sampling rate, FFT size, etc., also on the fly (Event 5).

3. Testbed results

Our experimental results confirm the practical feasibility of dynamic fronthaul reconfiguration in a converged LTE-PON scenario and provide the end-to-end reconfiguration time achieved by our implementation. The total time is calculated from the moment the SDN controller detects increasing load in the cell to the moment when the LTE capacity is increased towards the UE and is made available at the application level (i.e. delivered to the LTE sink).

Fig. 2 shows measured dynamics of fronthaul vs. background traffic adjustments, depending on the LTE application load. The LTE traffic from the mobile user is dynamically changed step by step to show operation at all of the rates/schemes reported in Table 1. In-between these changes, the cell always reverts to the 1.4 MHz scheme, corresponding to the minimal cell bandwidth and minimal fronthaul traffic. The background traffic is de-prioritized, so that any time the fronthaul traffic increases, its Quality of Service (QoS) remains guaranteed across the PON. Since the maximum fronthaul rate for a 20 MHz cell is 730 Mb/s, we limit the port rate of the Openflow switch to 1 Gb/s, so that we can reproduce the contention between the fronthaul and other background traffic.

Fig. 3 gives measured timings of the process of changing the cell's bandwidth. The events 1-4 show how the controller reconfigures the BBU from an initial state of 3 MHz to 5 MHz. After increasing the LTE application traffic, corresponding to Event 1 in Fig. 3, the controller sets the number of PRBs to 25 (5 MHz) and reconfigures the Committed Information Rate (CIR) of the switch to cope with the extra load. Using traffic traces, we measure a time period of about 150 ms for the control plane to react to the load increase in terms of reporting to the BBU and reconfiguring the Openflow switch. The final step towards Event 5, the change to the sampling rate, takes most of the reconfiguration time. This is due to the fact that the USRP has old samples buffered, belonging to the previous configuration (i.e., 15 PRBs). These samples will be transmitted at the new rate (25 PRBs), which is the wrong rate for them. Consequently, the UE is momentarily desynchronized from the transmission, leading to an overall reconfiguration time of 750 ms.

Acknowledgements

Financial support from Science Foundation Ireland (SFI) grants 14/IA/2527 (O'SHARE) and 13/RC/2077 (CONNECT) and European Union's Horizon 2020 grant 688941 (FUTEBOL) is gratefully acknowledged.

References

1. U. Dötsch, et al., Quantitative analysis of split base station processing and determination of advantageous architectures for LTE. *Bell Labs Technical Journal* **18**, 1 (2013).
2. M. Ruffini, Multidimensional Convergence in Future 5G Networks. *J. of Lightwave Technology* **35**, 3 (2017).
3. S. 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. *J. of Lightwave Technology* **34**, 18 (2016).
4. ITU-T, 10-Gigabit-capable symmetric passive optical network (XGS-PON). G.9807.1, (2016).
5. T. Tashiro, et al., A Novel DBA Scheme for TDM-PON based Mobile Fronthaul, OFC2014, paper Tu3F.3.
6. Software Radio Systems, "srsLTE Library" (2017), *online 2017-10-01*, <https://github.com/srsLTE>.