Future Metro bottle-neck – how will optical networks cost-effectively deliver anticipated 5G services?

A.Lord [(1)], A. Rafel [(1)] and P. Pavon[(2)]

[(1)] BT Applied Research, Orion 5, Adastral Park, Ipswich, IP5 3RE, United Kingdom (andrew.lord@bt.com)
[(2)] Universidad Politécnica de Cartagena and E-lighthouse Network Solutions, Cartagena, Spain

Abstract *Summarising the EU Metro-Haul project, we assess the impact of end-to-end orchestration for 5G, showing that a static control plane is essential to enable instantiation of myriad applications; but for the most cost-effective management of resources, a responsive, dynamic orchestration will be required.*

Introduction

The overall Metro-Haul objective is to architect and design cost-effective, energy-efficient, agile and programmable metro networks that are scalable for 5G access and future requirements, encompassing the design of all-optical metro nodes (including full compute and storage capabilities), which interface effectively with both 5G access and multi-Tbit/s elastic core networks.

The work has progressed in a classical way, by: (i) taking KPIs arising from the 5G paradigm and converting these into KPIs relevant to an optical network, (ii) designing architectures (with multiple components including physical, control plane, monitoring etc) to potentially meet these requirements, (iii) conducting deep domain-specific research into these individual components, (iv) building an integrated techno-economic framework, taking the component inputs, with the aim of validating the original KPIs – specifically capacity, cost, power consumption, (v) completing project demonstrations to show specific use cases and fully-working technology.

This paper will report specifically on the outcomes of step (iv) above, quantifying the benefits to 5G of building optical networks, including compute, storage and monitoring facilities that can be provisioned when required and according to the applications that need these specific resources.

Architecture

As shown in Fig.1, the scope of the work focuses on the metro region, including compute / storage resources at both the access- and core-facing nodes. Individual 5G applications could potentially require a wide set of resources, with respective KPIs, and exhibiting wide variation (e.g. some may be bandwidth hungry whilst others may require a tight latency and jitter guarantee)

There are a wide range of challenges here, but for the purposes of this paper we can differentiate them into two classes

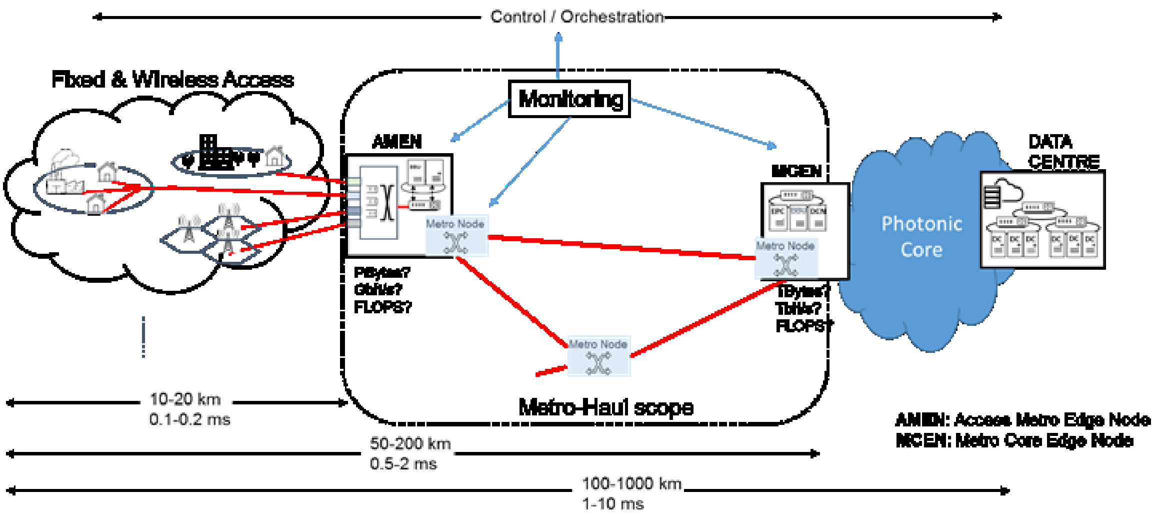

Fig.1. Overall network schematic with focus on metro components, including transport and DC resources

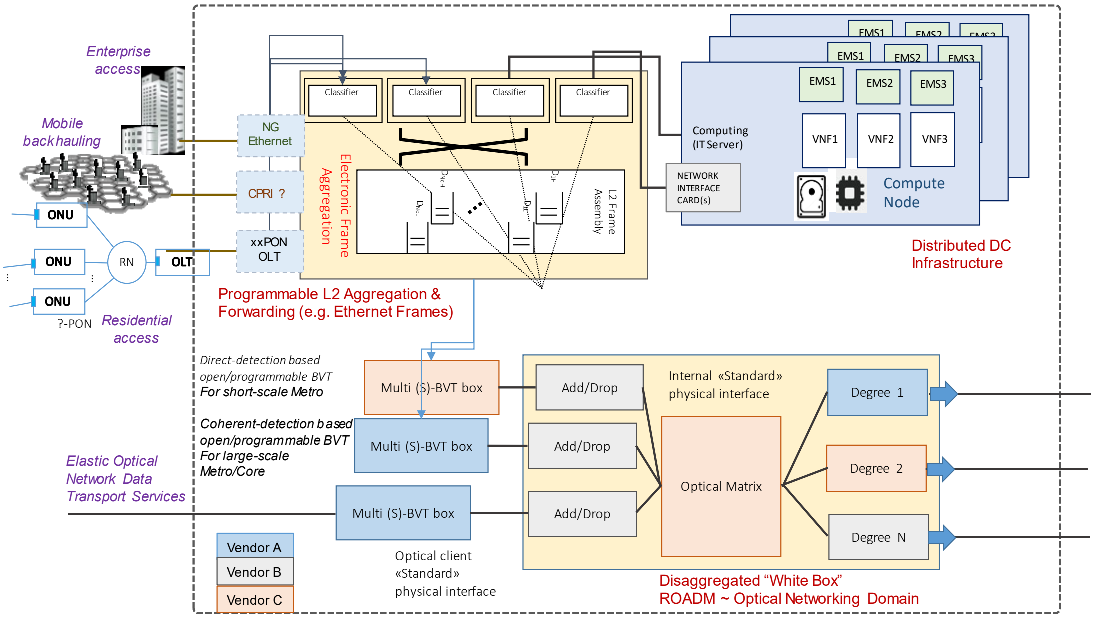

Fig.2. Metro-Haul node architecture including functional blocks

Static. Technology that allows the setting up of resources to facilitate multiple 5G applications with limited or no manual intervention. For a 5G eco-system to work effectively, applications should be instantiated within minutes rather than potentially weeks, and the orchestration of multiple disparate resources to handle this is already a complex challenge. Without the Metro-Haul technology to statically provision 5G applications, it will be very difficult to build an effective ecosystem.

Dynamic. Just having a static capability to automate the highly complex setting-up of resources needed to service multiple 5G applications is one thing, but to fully appreciate the benefits of the technologies behind Metro-Haul, the entire end-to-end system should be operated dynamically. In this mode, resources can be re-apportioned as required, limiting the amount of over-provisioning needed. 5G applications needing low latency can be handled closer to the edge, but those which have a less critical latency KPI can be backhauled through the metro to Core DCs. Although this can be done in the static case, it becomes more beneficial when operated dynamically, especially in terms of overall cost of expensive edge resources and of course, power consumption.

Drilling down into the details of the nodes (AMEN or MCEN) results in fig.2, showing the constituent functional blocks required to achieve this. The optical / physical layer transport is clear, as is the distributed DC infrastructure and facilities for higher layer (ethernet/ PON etc) transport.

Techno-economic approach

Throughout Metro-Haul a range of tools have been used, but a major focus has been on the NET2PLAN framework (www.net2plan.com [2]). As will be appreciated, the enormous task of modelling a complete system, including components from multiple layers, with distributed architectures etc is too large to be able to summarise in a short paper such as this. Indeed the full description occupies over 100 pages in the key Metro-Haul deliverables [3]. The methodology comprised taking the following inputs: detailed architectural designs, a comprehensive cost / price template for each component, traffic growth estimation, electrical power dissipation data. The following figure shows the overall flow for the techno-economic study:

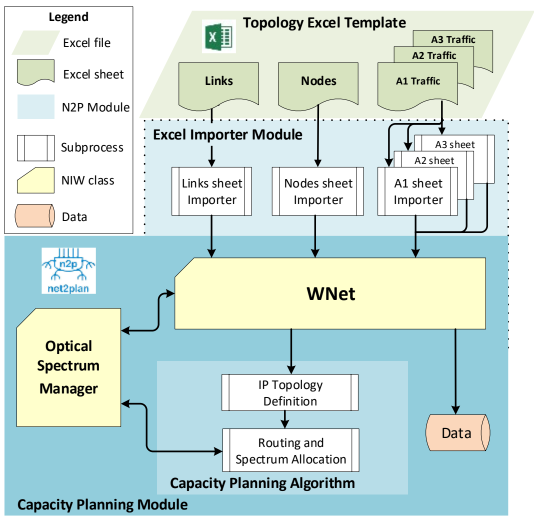

Fig.3. NET2PLAN workflow for evaluation of techno-economic analysis

Results

This paper has summarised the techno-economic evaluations that have taken place as part of the Metro-Haul project. There are of course many results, and always dependent on the input data and specific scenarios modelled. In an attempt to generalise project outcomes:

- The integration of Data Centre (DC) functionality(e.g. compute, storage) in Central Offices with a shared switching fabric provides massive reductions (above 70%) in packet switching hardware compared to traditional siloed approaches.
- The adoption of flexible switching technology at the optical layer, based on low-cost, few-degree ROADMs or filterless solutions can increase optical bandwidth while providing flexibility to redirect traffic flows to the most appropriate DC. Thus, the investment in the optical layer brings about reduced DC requirements, which can lower total cost by 20-30%. Filterless architectures in general showed modest savings of around 5-10%.
- Introducing disaggregation at the optical layer, which has been practically demonstrated in various Metro-Haul testbeds, can further reduce the deployment cost of optical components by 10-40%.
- The combination of a flexible optical layer, with sliceable transceivers and reconfigurable switching, and an agile control-plane with end-to-end orchestration capabilities across packet/optical/DC resources can reduce both the optical (transponders) and IT (servers, storage) resources significantly. Transponder count reduction is estimated to lead to around 20% savings.
- Reduction of 20% in NfV via optimisation of service chaining also emerged as a key result.

Each of these constituent components have different relative contributions to the overall CapEx, but the dominance of the first item guarantees an overall saving of more than 50%.

Conclusions and lessons learnt

This capability enabled by the Metro-Haul control and orchestration architecture has demonstrable CapEx and power consumption savings in various practical use-cases: optimization of 5G RAN CU placement across metro transport nodes through machine learning based traffic prediction, multi-layer service chaining exploiting grooming in the optical layer and lightpath/VNF reconfiguration in response to tidal traffic fluctuations throughout the day.

Overall, the aggregate analysis suggests that expected CapEx reductions would be slightly above the targeted 50%, due mostly to the packet switching reduction afforded by hybrid CO/DC nodes and the optimization of traffic through the optical layer reducing DC overprovisioning.

The power consumption reduction approaches the target of 50%, again due to the ability to flexibly direct flows across optical pipes to the most efficient VNF hosts, although more detailed modeling on the packet/DC power consumption trade-off is required going forward.

The capacity of the infrastructure can be increased more than ten-fold through a combination of more optical capacity and more edge processing capabilities. Increasing it further requires an optical infrastructure that can scale the bandwidth capacity per fiber and also provide transparent connectivity across different network segments without the cost and latency overhead of multiple packet aggregation stages in each traditional network domain.

ACKNOWLEDGEMENT

This paper represents the work of a large number of people, associated with the EU METRO-HAUL project.

References

[1] A. Ferrari et al., "Assessment on the Achievable Throughput of Multi-Band ITU-T G.652.D Fiber Transmission Systems," *J. Lightw. Technol.*, vol. 38

[2] P. Pavon-Marino, J.L. Izquierdo-Zaragoza, “Net2Plan: an open source network planning tool for bridging the gap between academia and industry”, IEEE Network, vol. 29, no. 5, pp. 90-96.

[3] Metro-Haul deliverables available via https://metro-haul.eu/