Politecnico di Milano
Dipartimento di Elettronica, Informazione e Bioingegneria

DOTTORATO DI RICERCA IN INGEGNERIA DELL'INFORMAZIONE

BaseBand Unit Hotelling Architectures for Fixed-Mobile Converged Next-Generation Access and Aggregation Networks

Doctoral Dissertation of:
Nicola Carapellese

Advisor:
Prof. Massimo Tornatore

Tutor:
Prof. Antonio Capone

Supervisor of the Doctoral Program:
Prof. Carlo Fiorini

2015 – Cycle XXVII

Research in International Institutions

Nicola Carapellese spent four months from January 20th to May 16th, 2014 as a visiting PhD student in Orange Labs Networks, Lannion (France), under the supervision of Dr. Anna Pizzinat, Dr. Philippe Chanclou and Dr. Stéphane Gosselin. During this period, he carried on collaborative activities in the framework of FP7 COMBO project.

He also visited UC Davis, CA (USA), from March 11th to March 16th, 2013, for a collaboration with students working in the Computer Networks Lab held by Prof. Biswanath Mukherjee, regarding the deployment of the WOBAN test-bed for the “5 per mille” project.

Abstract

To enable and sustain the “Internet Society” in the next future, characterized by an exponential growth of bandwidth and Quality-of-Service requirements by users, communication networks must be continuously evolved, by resorting to novel technologies and architectural solutions to improve cost and energy efficiency. In the realm of access and aggregation networks, such process causes severe issues to network owners. In fact, from one side such networks are much more expensive to evolve and consume a relevant quota of the energy consumption, with respect to core/backbone; from the other side, they often constitute the bottleneck of the whole network performance.

Among the main trends that are expected to guide the evolution of such networks towards cost and energy efficiency, we focus on two promising principles. The first one is the Fixed/Mobile Convergence (FMC), i.e., the concept of designing and optimizing networks as “a whole” resorting to infrastructure and equipment sharing among fixed and mobile networks. The second one is the BBU Hotelling (also known as C-RAN, Centralized/Cloud Radio Access Network), i.e., the new mobile access paradigm in which base stations are splitted among BaseBand Units (BBU) and Remote Radio Heads (RRH) and BBUs are centralized into hotels, with the consequent introduction of the new “fronthaul” traffic.

In this thesis, we investigate some FMC architectures incorporating the concept of BBU hotelling, for next-generation access and aggregation networks. After a survey of the relevant state-of-the-art BBU hotelling technological solutions, we make an energy-consumption comparison of some mobile network architectures that enable such schemes. Then, we identify optical WDM aggregation networks as the ideal substrate to perform hotelling, therefore we devise some alternative architectures. For these, a novel “BBU placement” optimization problem can be identified. An energy-minimization version of the model is formalized by Integer Linear Programming (ILP) and solved for multistage trees topologies. To deal with larger instances, a greedy-based heuristic algorithm is formulated, with a generic cost function. Finally, the model is further extended to consider the joint placement of hotels, the installation of aggregation electronic switches, and different options for fronthaul transport.

The results show the cost and energy advantages of BBU hotelling architectures with respect to classical RANs, due to “consolidation” of BBUs into a few sites, and give insights on the interaction of BBU placement with other degrees of freedom, like electronic aggregation switching, wavelength routing and number of available wavelengths and fibers.

Sommario

Per rendere possibile e sostenere la società basata su Internet nel prossimo futuro, caratterizzata da una crescita esponenziale della banda e dei requisiti di Qualità del Servizio da parte degli utenti, le reti di comunicazione devono essere costantemente evolute, affidandosi a tecnologie e architetture innovative per migliorare l'efficienza di costo ed energetica. Nell'ambito delle reti di accesso e aggregazione, questo processo è causa di questioni critiche per i proprietari della rete. Infatti, da un lato tali reti sono ben più costose da aggiornare e consumano una parte rilevante del fabbisogno energetico; dall'altro esse rappresentano il collo di bottiglia per la prestazione dell'intera rete.

Tra le varie tendenze che ci si aspetta guideranno l'evoluzione di tali reti verso l'efficienza di costo e energetica, ci si sofferma su due principi promettenti. Il primo è la Convergenza Fisso/Mobile (FMC), cioè il concetto di progettare e ottimizzare le reti "come un tutt'uno" grazie alla condivisione delle infrastrutture e degli apparati tra differenti reti fisse e mobili. Il secondo è il BBU "Hotelling" (anche noto come C-RAN, Centralized/Cloud Radio Access Network), cioè il nuovo paradigma di accesso mobile secondo il quale le stazioni radio base sono separate in BaseBand Unit (BBU) e Remote Radio heads (RRH) e le BBU vengono centralizzate in siti denominati "hotel", con conseguente introduzione del nuovo traffico "fronthaul".

In questa tesi, vengono investigate alcune architetture FMC che incorporano il concetto di BBU hotelling, per reti di accesso e di aggregazione di prossima generazione. Dopo una panoramica dello stato dell'arte delle soluzioni tecnologiche rilevanti per il BBU hotelling, viene presentato un confronto di alcune architetture di rete mobile che abilitano tali soluzioni. Successivamente, vengono indentificate le reti di aggregazione ottiche basate su WDM come il substrato ideale per l'hotelling, per le quali vengono definite diverse architetture. Per queste, un nuovo problema di ottimizzazione denominato "Posizionamento di BBU" può essere identificato. Una versione di tale problema, volta alla minimizzazione del consumo energetico, è formalizzata tramite Programmazione Lineare Intera (ILP) e risolta per topologie ad albero multi-stadio. Per gestire istanze di dimensioni maggiori, un algoritmo euristico di tipo "greedy"

è formulato, adoperante una funzione di costo generalizzata. Infine, il modello è esteso per considerare il posizionamento congiunto di hotel, l'installazione di apparati elettronici di aggregazione nei nodi intermedi, e diverse opzioni per il trasporto del traffico fronthaul.

I risultati mostrano i vantaggi in termini di costo e di energia del BBU hotelling, rispetto alla RAN tradizionale, dovuti alla "consolidazione" delle BBU in un numero inferiore di siti. Inoltre, essi gettano luce sull'interazione tra il posizionamento delle BBU e altri gradi di libertà, come la modalità di aggregazione elettronica, il tipo di instradamento delle lunghezze d'onda e il numero di canali WDM e fibre disponibili.

List of Publications

P1 Nicola Carapellese, Massimo Tornatore, and Achille Pattavina, “Placement of Base-Band Units (BBUs) over Fixed/-Mobile Converged Multi-Stage WDM-PONs,” in *2013 17th International Conference on Optical Network Design and Modeling (ONDM)*, April 2013, pp. 246–251.

The material of this publication can be found in Chapter 4.

P2 Nicola Carapellese, Massimo Tornatore, and Achille Pattavina, “Energy-Efficient Baseband Unit Placement in a Fixed/Mobile Converged WDM Aggregation Network,” *IEEE Journal on Selected Areas in Communications*, vol. 32, no. 8, pp. 1542–1551, Aug 2014.

The material of this publication can be found in Chapter 4.

P3 Nicola Carapellese, Anna Pizzinat, Massimo Tornatore, Philippe Chanclou, and Stéphane Gosselin, “An Energy Consumption Comparison of Different Mobile Backhaul and Fronthaul Optical Access Architectures,” in *2014 European Conference on Optical Communication (ECOC)*, Sept 2014, pp. 1–3.

The material of this publication can be found in Chapter 3.

P4 Nicola Carapellese, M. Shamsabardeh, Massimo Tornatore, and Achille Pattavina, “BBU hotelling in Centralized Radio Access Networks,” Chapter 3.3 of book *“Fiber Wireless Convergence” (Springer Optical Networks Series)*, ACCEPTED FOR PUBLICATION (2015).

The material of this publication can be found in Chapter 5.

P5 Nicola Carapellese, Massimo Tornatore, Achille Pattavina, and Stéphane Gosselin, “BBU Placement over a WDM Aggregation Network Considering OTN and Overlay Fronthaul Transport,” SUBMITTED TO: *2015 European Conference on Optical Communication (ECOC)*, 2015.

The material of this publication can be found in Chapter 6.

P6 Nicola Carapellese, Massimo Tornatore, and Achille Pattavina, “Joint BBU and Electronic Switch Placement in a FMC WDM Multifiber Aggregation Network,” SUBMITTED TO: *Journal of Optical Communications and Networking*, 2015.

The material of this publication can be found in Chapter 6.

Contents

List of Figures

List of Tables

Introduction 1

1.1 Motivation

The historical moment we are living now can be denoted as “information society” or “information age”, i.e., a period in which the creation, use and exchange of information play a crucial role in almost all human activities, being them scientific, economical, political or cultural. The rise and expansion of our information society has been possible thanks to the constant development of new technologies to store, compute, process and communicate information, which are known as Information and Communication Technologies (ICT), and their progressive penetration to larger and larger portions of the population. For instance, the progress of computer science and the advance in electronic miniaturization has led to a large-scale production of digital processors featuring more computational power, smaller size and lower cost, which in turn contributed to the widespread adoption of personal computers in workplaces and households. In addition to this, the progress in digital transmission and networking technologies opened the doors to some new telecommunication services that are the evolution of the traditional wireline telephony. The most successful of them are the mobile (or cellular) telephony and the access to the global system of interconnected telecommunication networks, best known as “the Internet”.

Within the wide and rapidly evolving ICT panorama, the Internet has eventually gained a central position, thanks to its unprecedented features, e.g., unified protocol structure based on Internet Protocol (IP) and Transmission Control Protocol (TCP), scalable and robust architecture, public access, inherent neutrality, transparency to all kinds of multimedia information, and fully decentralized structure. Indeed, the Internet “boom” started during the second half of the 90s, thanks to the increasing popularity of services like: World Wide Web (WWW), email, file transfer, Voice-over-IP (VoIP), peer-to-peer file sharing, video streaming, IP Television (IPTV) online games, social networks, location-based services, cloud storage, and many others. In few years, the Internet has become the symbol itself of our information society, up to evolving from a mere technological medium designed to satisfy the needs of the society to an entity that actively molds the structure of the society. Such process is so radical that it is thought to be marking another transition, from an information society to a network society.

Apart from sociological and cultural considerations (that also make a hot subject of interest nowadays) from a technical standpoint the most critical issue that ICT engineers are experiencing is the exponential growth of data traffic. It is a consequence of the continuous increase of the number of people using Internet, combined with the large variety of offered services, which require more and more capacity, coupled with stricter requirements in terms of Quality of Service (QoS). In addition, the average user is connected to the Internet through multiple devices, including mobile hand-held ones like netbooks, tablets and smartphones. Thanks to the portability of such devices, the user is becoming “always online”, i.e., he/she uses some services in the same way independently from his/her location and the time. As a consequence, the traffic originated by mobile devices is increasing at a greater pace than the conventional fixed traffic. Another factor which is forecasted to have a relevant impact for traffic growth in the next years is the “Internet of Things” (IoT), a further evolution in which potentially all objects become Internet users and exchange traffic independently from human interaction.

1.2 Evolving The Internet: Cost and Energy Issues

To cope with traffic growth and increasing QoS requirements, the whole Internet infrastructure needs to be continuously updated, by resorting to novel technologies, network architectures and communication protocols developed from industrial and academic research. To sustain this evolution, a rather complex economical ecosystem has developed around the Internet, in which the main actors are users, telecommunications operators, service providers, manufacturers, regulation and standardization entities. A detailed analysis of how such ecosystem works certainly is out of our scope, but it is worth identifying some basic issues that contribute to motivate the work of this thesis.

First of all, this process is a major source of costs by network owners (both operators and service providers), mainly consisting in Capital Expenditures (CapEx), i.e., for buying, building and installing infrastructure and equipment, but also Operational Expenditures (OpEx), i.e., for operation and management. Therefore, careful design choices must be made to achieve an acceptable tradeoff between capacity/QoS and costs, or in other words, “cost efficiency”.

A second source of concern is that the amount of electric energy consumed by internet infrastructures is going at the same pace as traffic growth. Energy consumption in telecommunication systems is already important because it impacts, sometimes heavily, on OpEx. Apart from being considered merely a cost factor, nowadays it is gaining much of attention *per se*, due to its implications to environmental issues like emissions of greenhouse gases and dependence on non-renewable energy resources. Indeed, as for cost, “energy efficiency” is going to be an additional design goal for network evolution, so that some consortia of manufacturers and operators are currently defining energy-related regulatory constraints (e.g., [1]).

The third issue involves how cost and energy is distributed across the infrastructure. The Internet is a huge, large, complex system, made of a multitude of networks featuring different tasks, thus different properties. A canonical hierarchical division is among three network segments: core (or backbone) networks, that interconnects regional and national hubs via high-capacity links; aggregation (or metro) networks, that interconnect urban centers belonging to a regional area via mid-capacity links; and access networks (also known as last-mile, or first-mile, or local loop), that connect final users to the Central Offices (CO), via lower-capacity links. Within access networks, a typical distinction is done between fixed access, in which final customers are gateways located in households, offices, buildings connected via copper or optical fiber links; and mobile access, in which final customers are mobile hand-held devices connected via radio (or wireless) links to intermediate Base Station nodes.

1.3 Access and Aggregation Networks

Within this structure, not all network segments contribute in the same amount to CapEx, OpEx and energy consumption, due to the differences in managed traffic (e.g., by different statistical multiplexing gains), used equipment, number of nodes, geographic density. As a matter of fact, and not as a first thought might suggest, the most relevant share comes from access networks, followed by aggregation, and at last by core. This is a consequence of the “economies of scale” that such kind of systems tend to exhibit. This trend, coupled with the traffic growth, definitely justifies the attention towards access and aggregation networks.

Focusing on the side of access, a great variety of technical solutions both in fixed and in mobile access networks allowed to increase the capacity and QoS experienced by each user. For instance, the sophisticated digital transmission techniques employed in DSL (Digital Subscriber Line) technologies enabled to better exploit the bandwidth available in the traditional copper pair and get rid of the limitation of dial-up and ISDN (Integrated Services Digital Network), to deliver for the first time fixed broadband access to households. Starting from the first standardized version, known as Asymmetric DSL (ADSL), such technology evolved into several variants featuring increasing bitrates in both downstream and upstream (e.g., ADSL2, ADSL2+), up to the “Very-high-bit-rate” versions (e.g., VDSL, VDSL2). More recently, a major upgrade has been standardized, denoted as G.fast (Fast Access to Subscriber Terminals), which makes use of a much larger bandwidth and innovative techniques (e.g., “vectoring”), to deliver up to 1 Gb/s for local loops shorter than 500 m. A similar process involving radio transmission technologies has been driving the evolution of mobile access, from 2nd-Generation, or 2G (e.g., GSM/GPRS/EDGE: Global System for Mobile communication General Packet Radio Service / Enhanced Data rates for GSM Evolution), to 3G (e.g., UMTS/HSPA: Universal Mobile Telecommunication System / High-Speed Packet Access).

Today, new access solutions are further enriching such panorama. The penetration of optical communication closer and closer to user premises seems to be the only way to evolve the fixed access in response to the growing data rates requested by users. The term Fiber-to-the-X (FTTX), where the X can stand for “Cabinet”, “Street”, “Building” or “Home”, generically denotes some architectural solutions in which optical fiber partially or fully substitutes the legacy copper infrastructure. Implemented technologies can range from point-to-point Ethernet links to point-to-multipoint networks, e.g., Active or Passive Optical Networks (AON, PON). The corresponding trends in mobile access are larger radio-frequency bandwidths and increased per-user spectral efficiency, that is obtained by means of more advanced multiple-access techniques like OFDMA (Orthogonal Frequency Division Multiple Access), multiple-antenna transmission/reception schemes, and cell densification (from macro- to micro-, femto- and pico-cells). They form the basis for 4G standards like LTE (Long Term Evolution) and LTE-Advanced.

Similar evolution trends are also acting on aggregation networks, which have been historically dominated by legacy transport technologies based on SDH/SONET (Synchronous Digital Hierarchy / Synchronous Optical NETworking) and switching technologies based on ATM (Asynchronous Transfer Mode). In more recent years, the adoption of optical-fiber networks based on Wavelength Division Multiplexing (WDM) appears to be crucial for traffic growth, together with the shift, currently in progress, towards Ethernet, OTN (Optical Transport Network) and IP/MPLS (Multi-Protocol Label Switching).

In general, a lot of separate access and aggregation technologies are going to coexist within next years. Although new technologies eventually replace obsolete ones, such process is slow and not always clean, because some degree of backward compatibility must be kept in order to serve also users who do not want to update their terminal equipment and/or to change their service level. Moreover, network operators might prefer upgrading partially their systems rather than change them completely, for cheaper and faster deployment. For instance some access solutions feature novel interfaces towards users, but with little or no change to the remaining part of the network that transports traffic between users and the rest of the Internet (usually known as backhaul network). Other causes of access technology diversification can be, for instance, geographical/environmental conditions that limit the installation of certain technologies, reuse of already existing infrastructure, nonuniform distribution of users' location, co-presence of competing operators, as well as different standardization drivers or law/regulation constraints.

1.4 Fixed/Mobile Convergence (FMC)

In the next years, this process is leading to an extreme fragmentation of the access/aggregation panorama. In particular, a clear separation is noticeable between the family of fixed access and the family of mobile access networks, which even exhibit contradicting design trends (e.g., centralization of fixed networks versus decentralization of mobile networks). The concept of "convergence", i.e., the progressive consolidation of many different technologies, architectures and protocols into unified systems, can play a very important role for network evolution, because of its potential benefits, like overall costs reduction, higher performance and lower power consumption. In its most general meaning, convergence is not a novel concept. For instance, it has been one of the causes of the success of the Internet, i.e., passing from multiple media-dependent communication technologies to a single, "converged" system. Another clear example of convergence is in the mobile hand-held devices, which now are capable of bringing many different services to the user, which were originally developed for dedicated devices.

Today, there are many different implementations of network convergence. In this work, we refer to the so called Fixed/Mobile Convergence (FMC). Although FMC can be generally applied in many different flavors and at various network levels, in this work we intend it as infrastructural convergence, in the context of access and aggregation. It means that a single network infrastructure becomes a shared facility devoted to the transport of both fixed and mobile access/aggregation traffic towards a higher-level gateway site, that can still be part of the aggregation network or also be directly a node of the core. Thanks to the mutualization of infrastructure and hardware, like cable plants, cabinets, equipment, sites and buildings, a better utilization of resources can be achieved. This improvement leads not only to a substantial reduction of deployment (CapEx) and operational (OpEx) costs, but also enables higher energy efficiency, because different active (i.e., power-supplied) devices can share housing facilities. For these reasons, there is a large consensus on the important, promising role that will played by the FMC paradigm within next years, for the development of Next-Generation Access and Aggregation Networks.

Among the different FMC technologies under research, optical networks based on WDM are the most promising. Indeed, they ensure a truly future-proof capacity provisioning, but also enable energy efficiency thanks to the employing of transparent wavelength router nodes. In fact, they are able to switch “lightpaths” (i.e., optical signals that traverse a path of fiber links using a single wavelength) directly in the optical domain, i.e., without performing power-consuming electronic traffic processing and optical/electro/optical (OEO) conversions. In addition, a common belief is that they will probably follow the same evolutionary steps of core networks, e.g., with more complex nodes architectures and the possibility of efficiently grooming sub-wavelength connections into single wavelengths, towards multi-stage and partially-meshed topologies. As a consequence, they will inherit some powerful features of core networks, like failure resilience, scalability and reconfigurability [2].

1.5 BBU Hotelling for Radio Access Networks

An FMC access/aggregation infrastructure is expected to support several mobile Radio Access Network (RAN) solutions. Among those, BBU hotelling (or hosteling) [3, 4] is a recent solution which radically changes the classical architecture of RANs. This technique takes advantage of the functional separation of a generic base station (BS) into two parts, the Base-Band Unit (BBU) and the Remote Radio Head (RRH). The BBU performs layer 1 (L1) digital processing of the baseband signals along with all functions of the upper layers, and interfaces with the backhauling network. The RRH interfaces with antennas' front/back-ends and performs remaining L1 functions, i.e., Digital-to-Analog (DA) / Analog-to-Digital (AD) conversion of the baseband signals, frequency up/down-conversion and power amplification.

The BBU Hotelling technique applies to a cluster of BSs and consists in geographically separating each BBU from its RRH, which remains located at the cell site (i.e., where the antennas are located). All BBUs are consolidated into a common location, called "Hotel". Each BBU/RRH pair exchanges a group of Digitized Radio-over-Fiber (D-RoF) baseband signals, one per radio access technology, per antenna, per direction (uplink, downlink). For these signals, some open transport interfaces have been already specified (e.g., CPRI [5], OBSAI [6]), which also define how the baseband signals are digitized (i.e., sampled and quantized), how they are multiplexed and mapped into a single data link and specific control/management protocols.

BBU consolidation, thanks to the sharing of backplanes, power, computational and maintenance resources of BBUs hosted in the same hotel, can potentially bring a valuable extent of costs and energy savings. Besides, it would also make more practically implementable advanced algorithms of coordinated multi-cell L2 scheduling and L1 signal processing, both proprietary and standard-based, e.g., for LTE/LTE-A, Cooperative MultiPoint transmission (CoMP) and Enhanced Inter-Cell Interference Coordination (eICIC). These techniques would allow a higher spectral efficiency, thus further reducing the total cost to system throughput ratio. The BBU hotelling technique, although ambitious and very promising, is actually seen as a first step towards the more general long-term concept of C-RAN (e.g., Centralized, or Cloud RAN), in which the radio access network functionality eventually will be detached by the cell/antenna sites and migrated/virtualized into a flexible pool of hardware/software resources placed higher over the network hierarchy [7, 8] .

1.6 Main Related Work

Keeping networks up to date with traffic growth and increasing QoS constraints requires massive investments by network operators, especially for fixed/mobile access and aggregation (i.e., backhaul) network, because they need much more infrastructural changes than in the core/backbone. Finding cost-efficient ways to evolve access and aggregation networks is a complex task because several environmental factors interact, like available time horizon, available network technologies, business ecosystem, geographical/regulation constraints. The EU FP7 project "COMBO" [9] aims at investigating this issue, by resorting on the concepts of infrastructural and functional convergence, in order to eventually define some FMC access and aggregation network architectures, as function of the environmental factors. Optical network technologies play a relevant role, because they are the only capable of fulfilling the requirements of FMC access/aggregation network, like high capacity due to Wavelength Division Multiplexing (WDM), large coverage areas due to the long optical reach, and cost efficiency due to the fact that in most cases intermediate nodes can be implemented as cheap passive devices. Specifically, Passive Optical Networks (PON) based on WDM are a promising FMC solution, enabling a flat access/aggregation integration directly towards the core network, as proposed in [10].

Optical networks also represent the ideal substrate for an effective implementation of the new Centralized Radio Access Networks (C-RAN) paradigm [11], in which Base Stations (BS) providing mobile access are split between BaseBand Units (BBU), which are centralized in shared premises denoted as "BBU hotels", and Remote Radio Heads (RRH), which are directly attached to antennas at cell sites. Centralizing BBUs into few hotels brings many benefits in term of cost efficiency and radio performance improvement, but requires an adequate network to transport the new "fronthaul" traffic, i.e., the digitized baseband signals exchanged by BBUs and RRHs [12, 13], in addition to the conventional backhaul traffic. In the short term, different existing technologies can be properly adapted to the possible mixes of fronthaul and backhaul. In the longer term, C-RAN can be combined with FMC PON-based solutions, as described for instance in [14].

PON-based access and aggregation is surely attractive as a disruptive FMC solution, but, in view of an evolutionary scenario, some recent works assume that a plurality of access technologies keep coexisting in the last-mile. Such heterogeneous access network segments extend up to some close nodes (usually, the first Central Office), where the traffic is multiplexed and sent to a single converged WDM aggregation network, as for instance in [15]. These alternative architectures are also rapidly gaining interest in the recent literature, since they exhibit a somewhat complementary approach to PON-based solutions. In fact, they prescribe a paradigm in which FMC applies to aggregation networks, that are extended closer to the final users, up to “penetrating into access”, while other PON-based FMC solutions describe an access “penetrating into aggregation”.

1.7 Contribution and Outline of the Thesis

As motivated above, Fixed/Mobile Convergence, BBU Hotelling and the paradigm of aggregation penetrating into access emerging from the more recent research have a great potential for the network evolution. Within the several research challenges that these topics can originate, in this thesis we investigate some FMC optical aggregation network architectures incorporating the concept of BBU hotelling.

Our research work starts with a survey of the technological features of BBU hotelling and implementation principles from a networking standpoint. Then an analysis of the energetic benefits that this technique can bring over currently aggregation infrastructure is performed. The research continues by introducing some possible aggregation FMC architectures and formulating a novel “BBU placement” network optimization problem aiming at energy efficiency, then the problem is generalized towards a more generic “cost” efficiency and larger network sizes, by introducing an heuristic optimization algorithm. Finally, a last evolution of the optimization model is presented, which deals with a more real-word scenario featuring the joint optimization of BBU and electronic switches placement, in a multifiber WDM aggregation network, up to the conclusion and some discussion about open research issues.

More in detail, these research steps are organized throughout the remainder of this thesis as follows.

- In Chapter 2, we focus on the role of BBU hotelling in converged network architectures. The main motivations behind this technique are described, and the critical drawbacks are detailed, mainly related to the transport of the new “fronthaul” traffic over the network infrastructure. A classification of the various architectural solutions for BBU hotelling is detailed, regarding BBU placement and implementation, and fronthaul transport.

- In Chapter 3, we present an energy-consumption comparison of different optical network architectures for both mobile backhauling and fronthauling, performed in collaboration with *Orange Labs Networks, Lannion.* The main focus is about how much energy savings are allowed in a macrocells based RAN, built on traditional mobile aggregation infrastructures with no fixed/mobile convergence, under the different combinations of choices for BBU placement and aggregation technologies that are described in the previous chapter. Finally, the energy results are combined with the analysis of the total length of fiber needed for each solution, to identify some basic tradeoffs.

- In Chapter 4, we continue the investigation of the benefits that can be achieved by BBU hotelling solutions in terms of energy consumption, from an overall network perspective. With respect to the previous chapter, we shift our focus to an optimization approach that can be applied on a more generic fixed/mobile converged network, in which the location of hotels in not a priori given, but constitutes one of the outputs of the decision. To do so, we define some architectures for a FMC aggregation network based on WDM, in which BBU hotels can be potentially located into any of the intermediate nodes of the aggregation infrastructure. Therefore, a novel network optimization problem, jointly involving BBU placement and traffic routing, is defined and investigated.

- In Chapter 5, the BBU placement problem introduced in the previous chapter for the optimization of energy consumption in a FMC WDM aggregation network is evolved towards two directions. First, a more generic optimization metric is considered, that jointly takes into account the dominating contributions of cost and energy consumption in such network. Second, a larger-scale aggregation network is considered, that comprises several stages of intermediate COs up to a single high-level Point of Presence (PoP) acting as edge node. As a consequence of the increased number of nodes, an ILP formulation would be too complex to be solved for realistic instances, therefore a greedy-based heuristic algorithm is proposed.

- In Chapter 6, we introduce a novel optimization problem that takes into account more degrees of freedom of WDM aggregation network design. Differently from the previously defined problem, a multifiber infrastructure is considered and a new node architecture is adopted, that allows both the insertion of internal traffic (locally generated) and the transit of external traffic. The transit traffic can be either purely bypassed by an optical switch, or be aggregated/groomed by installing an additional electronic switch, independently from the presence of a hotel. Therefore, the decision of installing electronic switches in each node and the multifiber network dimensioning are optimized, together with BBU placement and traffic routing. In this way, the model becomes more scalable to real-world larger networks and more interesting to investigate because of the non-trivial interaction among hotels/switches placement and traffic routing/grooming.
- In Chapter 7, final conclusions are drawn and some open issues are discussed.

BaseBand Unit (BBU) Hotelling: Network Architectures and Technologies 2

This chapter focuses on the role of BBU hotelling in converged network architectures. The main motivations behind this technique are described, and the critical drawbacks are detailed, mainly related to the transport of the new “fronthaul” traffic over the network infrastructure. A classification of the various architectural solutions for BBU hotelling is detailed, regarding BBU placement and implementation, and fronthaul transport.

2.1 Introduction

The current explosion of traffic generated by mobile devices (e.g., smartphones, tablets, USB dongles) requires some radical changes to existing mobile network technologies and architectures. Some changes are being already introduced, for instance adopting novel radio access technologies (e.g., LTE) to increase spectrum efficiency and deploying additional cells to serve high traffic density areas (e.g., micro-cells, small-cells). Several other improvements are under investigation by academic and industrial research. The “BBU hotelling” (sometimes written as “BBU hostelling”) is one of the ways to evolve mobile networks. It is not a single technique, but a family of technological and architectural concepts that radically change the way in which mobile networks are implemented. For this reason, in order to understand the motivations behind BBU hotelling and their potential benefits over conventional architectures, it is important to have a glance on how a mobile network is structured and how it works. Therefore, in Section 2 the fundamentals of mobile networks are presented, while in Section 3 details of the implementation of base stations are given, which lead to the definition of BBU hotelling and its main advantages in Section 4. The challenges of BBU hotelling, regarding the new "fronthaul" traffic, are described in Section 5. In Section 6, several possible BBU hotelling architectures are discussed.

2.2 Mobile Network

A typical mobile network can be divided into three parts: Radio Access Network, Backhaul Network and Core Network. The Radio Access Network (RAN) includes all and solely the systems performing radio-access related functions, i.e., directly managing radio transmission and reception towards/from mobile devices. Note that the standard that defines the network architecture and specifies functions, interfaces and protocols assigned to RAN nodes and mobile devices is called the Radio Access Technology (RAT) (e.g., WCDMA/HSPA, LTE, WiMax). The Backhaul Network performs traffic aggregation and transport between the RAN and the Core Network. For this reason, its architecture and implementation can be almost agnostic with respect to radio access and core architectures, so they are not covered by RAT standards, given that adequate transport capacity and QoS (e.g., latency) requirements are guaranteed. Finally, the Core Network is in charge of all remaining non radio-access related functions and acts as a gateway towards all other mobile and fixed networks, i.e., towards the Internet. Some core network functions and interfaces are standardized too, in most of cases accordingly to the adopted RAT.

The RAN directly interfaces to mobile devices (UE, User Equipment) via radio links established towards Base Stations (BS). Each BS manages the transfer of user and control data towards (downlink) and from (uplink) several UEs simultaneously, by means of physical-layer and multiple-access protocols, according to the so-called radio, or air, interface. Some higher-layer radio-access functions (e.g., radio resource control) can be either performed by other network nodes (e.g., Base Station Controllers, BSC, or Radio Network Controllers, RNC) that manage several BSs, or directly embedded into the BSs themselves. Each BS manages UEs belonging to a specific coverage area, denoted as “cell”, and the RAN also coordinates the procedures for user mobility, i.e., allowing UEs to move across adjacent cells (handovers), without losing data connection.

BSs are placed into premises denoted as “cell sites”, whose geographic coordinates are influenced by many different factors, most notably coverage, capacity planning and infrastructural/costs constraints. To save costs, a consolidated practice is implementing more than one BS into a single cell site, thus dividing the coverage area into up to three cells, denoted also as “sectors”. For the same reason, BSs of different RATs and different mobile operators often share the same cell site. A typical cell site consists in a tower or a mast, on top of which there are installed BS directional antennas (at least one per sector), and a cabinet, or shelter, where the remaining BS equipment is installed. The cabinet also hosts collateral systems which do not perform network functions, but ensure proper BS working. They typically consist in power supplying (AC/DC converters, backup batteries) and cooling systems (fans, air conditioning).

2.3 Evolving the Base Station: BBU and RRH

The fundamental enabler of BBU hotelling is a recent evolution on how BSs are physically implemented. Traditionally, the BS consisted in an all-in-one solution produced by a single manufacturer, performing all functions and interfacing to both the backhaul network and the antennas via coaxial cables. With time, manufacturers found more convenient to adopt a modular architecture, i.e., separating some functions in different subsystems. This trend eventually led to a well-consolidated solution, in which the BS is divided into two separate kinds of modules:

- *BaseBand Unit (BBU)*, which performs physical-layer digital processing of the baseband version of the radio signals, includes all functions of the upper layers, and interfaces with the backhauling network;

- *Radio Resource Head (RRH)*, which performs remaining physical layer functions, i.e., Digital-to-Analog (DA) / Analog-to-Digital (AD) conversion of the baseband signals, frequency up/down-conversion, power amplification and some measures on the received analog signal, and interface with antennas via coaxial cables.

The data exchanged between BBU and RRH consists in a digitized baseband version of the radio signals received and transmitted by the antennas, respectively for the uplink and downlink direction. This data communication is enabled by an ad-hoc transport interface established between BBU and RRH. Since the physical connection between them typically consists in an optical fiber link, such kind of transport is known as Digitized Radio-over-Fiber (D-RoF). Another widely used denomination is the “fronthaul”, as a contraposition to the backhaul.

As an alternative to different proprietary fronthaul interfaces, some consortia of manufacturers have defined public interfaces, of which the most established in the market are CPRI and OBSAI. It is important noting that such interfaces were initially defined for short-range connection of BBUs and RRHs located in the same cell site, and later they have been considered as natural solutions for fronthaul transport in BBU hotelling architectures. Both the standards define different fronthaul transport formats and protocols, covering all the configurations of the main radio access technologies. CPRI is currently the most commercially adopted solution, although it still has some vendor-specific implemented features, which prevent a full multi-vendor interoperability.

2.4 Advantages of BBU Hotelling

Splitting BSs into BBUs and RRHs and defining open public interfaces for fronthaul has some valuable benefits. In facts, it facilitates the design and implementation of evolved features, because each module can be separately updated, it also enables the interoperability between different manufacturers and opens the market to smaller newcomers. As a consequence, substantial costs savings can be achieved by both manufacturers and network operators. However, the most important innovation is that BBUs and RRHs are no more seen as different parts of the same BS, but as distinct network nodes. In fact, they can be independently placed over the infrastructure, and exploit the existent RAN to transport the fronthaul.

The BBU hotelling is a family of different techniques that take advantage of such physical separation between the BBU and the RRH. At a basic level, it consists in placing the BBUs of a set of cells no more in their respective cell sites, but centralizing them into a single shared site, denoted as "BBU hotel", or simply "hotel". This can be done in two main ways: by "stacking", i.e., purely grouping the BBUs; or by "pooling", i.e., re-implementing their functions in fewer devices that are designed to share ("pool") some of their hardware/software resources.

With respect to the traditional architecture, BBU hotelling solutions promise various kinds of advantages, regarding costs, energy consumption and radio performance [16].

2.4.1 Cost Reduction

Costs faced by network operators are typically classified as: Capital Expenditures (CapEx), paid once for purchasing sites and premises, deploying the support infrastructure and installing network equipment; Operational Expenditures (OpEx), paid continuously over time, to ensure the proper operation of the network, and including lease fees, energy bills, control/management, reparation and updating costs. Installing BBUs in a few centralized hotel sites rather than in every cell site can substantially reduce the required space and the installation times, therefore the relative purchase/rent cost.

In addition, the BBU centralization into hotels allows sharing some collateral subsystems (e.g., power supplying, cooling, air conditioning, control/monitoring servers, in-site interconnection backplanes), instead of deploying an independent subsystem for every cell site, thus gaining benefit from some "economies of scale". As a consequence, it gets less expensive to manage the maintenance, reparation of failures and hardware/software updates, because most of critical equipment is centralized in single hotel sites, which are often more easily and quickly accessible.

2.4.2 Energy Savings

The energy consumption of the RAN is of primary interest for operators, mainly because it heavily impacts on OpEx. In addition, it is also caused by the increasing concern towards environmental issues (reduction of the carbon footprint and less dependence from non-renewable energy sources). BBU hotelling has a great potentiality for reducing the energy consumption, as a direct consequence of the sharing of collateral subsystems into hotel sites. Differently from cell sites, whose equipment rooms are mostly obtained from existent premises used for other purposes, hence not energy-optimized, hotel sites can be conveniently adapted from already deployed telecom offices or cabinets, taking advantage of existent cooling and power management systems [17].

A further energy consumption reduction is expected by the BBU “pooling”, with respect to the simple “stacking”, because, thanks to an accurate optimization of the pool, an overall better resource utilization can be achieved. For instance, when the traffic load of some cells is low, fewer resources can be shared, thus reducing the energy consumption. It is also expected that the achievable amount of pooling gain increases when a bigger number of cells is managed by a single hotel, which could be a strong pushing towards massive BBU centralization (up to hundreds of BBUs per hotel).

2.4.3 Improved radio performance

The radio performance of a RAN, i.e., the average bitrate and Quality of Service (QoS) delivered to customers, strongly depends on the capability of the signal processing techniques used to cope with the hostility of the wireless channel, which is characterized by multi-user interference and harsh radio propagation conditions. To improve radio performance, the recent RATs (e.g., LTE, LTE-A) adopt multi-user and multi-cell joint processing schemes, that require exchange of end-user data and uplink/downlink channel state information among different BSs. This is one of the reasons why LTE introduces a logical interface (X2) to interconnect different BSs.

The BBU centralization makes available a virtually zero-latency backplane among all BBUs hosted in the hotel site, thus opening the doors to some advanced algorithms, such as multi-cell packet schedulers, Inter-Cell Interference Coordination (ICIC) and Coordinated MultiPoint transmission and reception (CoMP), which improve the radio performance by increasing the spectrum efficiency, especially for UEs located at cell edges. For the same reason, implementing BBU as pooling can further enhance the radio performance, because some processing functions can be implemented into single hardware parts, thus gaining more computational power and even less latency.

2.5 Challenges of BBU Hotelling: Fronthaul

Fronthaul traffic has radically different features compared to those of backhaul traffic, and it poses specific challenges to the transport network. There are three main critical characteristics.

2.5.1 High, constant bitrate

Fronthaul flows have a very high bitrate, on the order of units to tens Gb/s for a cell site, as resulting from aggregating flows of one or more antennas, sampling their baseband signals with bandwidths around tens of MHz, and quantizing their I (In-phase) and Q (Quadrature) components with about 10 to 20 bits per sample. The specific values depend on the fronthaul interface, RAT typology and configuration [5, 18]. For instance, an LTE sector configured as 2×2 MIMO with 20 MHz bandwidth requires approximately 2.5 Gb/s, which gives 7.5 Gb/s total fronthaul for a typical 3-sector cell site.

Normally, such bitrates do not scale with the varying traffic load condition of the cell, resulting in a fully non-elastic traffic. These features constitute a relevant problem for the traditional access infrastructures, which are designed to transport much lower bitrates and to massively exploit the statistic multiplexing gain (i.e., the fact that networks can be under-dimensioned with respect to the peak traffic values, because of their statistical variability). As a consequence, a practical and future-proof transport solution for fronthaul traffic is given by optical access networks, but this require a capillary rollout of optical fiber links, needing relevant investment by network operators.

As an attempt to mitigate this issue, some of the recent research is focusing on some alternative solutions. For instance, advanced compression schemes for fronthaul are under study, in order to reduce the bitrate by a factor 2 to 4 [19]. Also, thanks to the reduced bitrate, ways to transport the fronthaul over point-to-point microwave or millimeter wave links are being investigated [20].

2.5.2 Maximum End-To-End Latency

RAT standards specify strict timing conditions for some physical layer procedures between BS and UEs. Most of them explicitly pose bounds on the “BS latency”, i.e., the latency due to internal processing of radio frames by the BS. In BBU hotelling, the BS functions are actually “spread” between BBUs and RRHs potentially located very far apart from each other, therefore the “fronthaul latency”, i.e., the delay contribution due to the transport of fronthaul signals along the RAN infrastructure, has a relevant impact on the total latency budget inside the BS. As a result, given that the total BS latency budget is fixed by the standard and the internal processing delays of BBU and RRH are dependent on the specific hardware/software implementation, there is an upper limit on the tolerable fronthaul latency. Such latency can be further divided in two parts.

The first one is due to the adaptation of fronthaul signals into the RAN transport service and it is purely technology-dependent, e.g., caused by CPRI or OBSAI transmission/reception interfaces, and additional functions required by optional lower-layer transport technologies, e.g., buffering, reframing, mu/demultiplexing, error-correction.

The second contribution is due to the signal propagation along the RAN, and it constitutes a fundamental limitation for BBU hotelling, because it constrains the maximum geographical distance between hotel sites and controlled cell sites, with a relevant impact on architectural choices. To provide some illustrative numerical values, we consider the LTE Frequency Division Duplexing (FDD) radio interface. In this case, the tightest BS latency limitation is due to the timing of the synchronous uplink Hybrid Automatic ReQuest (HARQ). Specifically, once received the uplink data packet at radio frame number i, the BS must send back

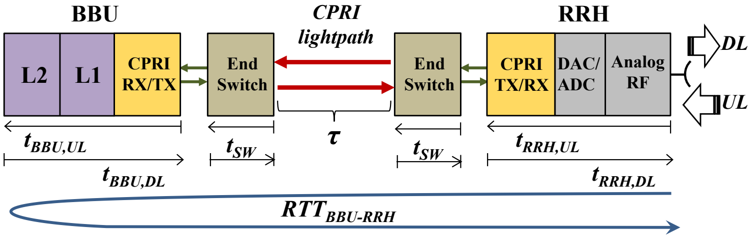

FIGURE 2.1: Delay contributions of $RTT_{\text{BBU-RRH}}$ along the fronthaul processing chain.

the correspondent ACK/NACK indication at frame number $(i+3)$ [21]. This gives exactly 3 ms of latency budget for the BBU-RRH round trip time $RTT_{\text{BBU-RRH}}$, i.e., the time difference between the complete reception of the data frame and the start of transmission of the ACK/NACK indication. All delay contributions having impact on this quantity are summarized in Fig. 2.1. Hence, the following must hold:

$$RTT_{\text{BBU-RRH}} = 2\tau + t_{\text{RRH,UL}} + t_{\text{BBU,UL}} + t_{\text{BBU,DL}} + t_{\text{RRH,DL}} + 4t_{\text{SW}} \leq 3\,\text{ms} \quad (2.1)$$

As a consequence, the maximum value of the one-way propagation delay along the fronthaul link (τ) depends on the remaining processing times. Regarding the internal BBU and RRH processing times in uplink (UL) and downlink (DL), there are no universal values, because they largely depend on the vendor implementation. By averaging over currently typical commercial equipment implementing CPRI interfaces, the following approximate values can be given:

$$t_{\text{RRH,UL}} \simeq 12\,\mu\text{s}; \quad t_{\text{BBU,UL}} \simeq t_{\text{BBU,DL}} \simeq 1.3\,\text{ms}; \quad t_{\text{RRH,DL}} \simeq 21\,\mu\text{s} \quad (2.2)$$

The processing time introduced by the transport end node (t_{TR}) also depends on the specific transport technology. For instance, typical values for OTN wrappers can be in the range 28 to 41 μs . Considering a basic fronthaul implementation, in which the fronthaul is transported “as it is”, for instance as dedicated CPRI links, there is no latency contribution caused by the fronthaul transport over a lower-layer technology (i.e., $t_{\mathrm{TR}} = 0$).

Therefore, the previous numbers lead to a value of maximum one-way fronthaul latency equal to about 184 ms, which is equivalent to a maximum fronthaul link distance about 37 km (assuming the classical value of 5 μs / km for fiber propagation delay). However, in case of OTN transport such values reduce to 128 to 102 μs, which correspond to a range 20 to 26 km of maximum distance. These values are in accord with the typical range of 20 to 40 km reported in the recent literature [16]. However, since a precise characterization of such quantity is extremely important for BBU hotelling deployment, further investigation is also being performed by standardization bodies.

2.5.3 Strict QoS Requirements

Digitized baseband radio signals are very sensitive to received bits errors, mismatches between transmitter and receiver clock frequencies and random variations of the instantaneous received clock phase (jitter). In facts, differently from traditional BSs, in which a single clock generator feeds both co-located BBU and RRH, in BBU hotelling the fronthaul transports not only the baseband data in both directions, but also the clock signal, generated at the BBU, towards the RRH. In order to meet the requirements imposed by RAT standards on accuracy and stability of radio interface, BBUs and RRHs must be precisely synchronized to each other. As a consequence, fronthaul has strict QoS requirements, in terms of Bit Error Rate (BER), frequency and phase synchronization. As an example, CPRI specifies that BER must not be greater than 10^{-12} and that the jitter introduced by the CPRI link can contribute of a quantity not greater than 2 parts per billion (ppb) to the BS frequency accuracy budget [5]. While such requirements are satisfied in a BBU hotelling scenario in which fronthaul is transported over dedicated point-to-point optical links automatically meeting the CPRI specification, they are critic whenever fronthaul exploits an underlying transport technology. In principle, the constraint on BER is less problematic, because it can be met by adjusting the transmission power budget, in order to increase the received Signal-to-Noise Ratio

(SNR), and/or adopting more powerful Forward Error Correction (FEC) algorithms. It is worth noting that such algorithms generally increase the signal processing latency, thus reducing the maximum tolerable fronthaul latency. The synchronization constraints are more critical. For instance, they make unfeasible the transport over Ethernet today, even though Ethernet standards are continuously evolving to include advanced carrier-grade QoS features in next future (e.g., Carrier Ethernet Transport [22]). The transport over Optical Transport Network (OTN) is seen as a more feasible solution, because it defines a fully synchronous optical signals hierarchy, with embedded advanced features for mu/demultiplexing and controlling several fronthaul flows into a single optical signal. For this purpose, the ITU-T G.709 standard specifies a set of recommendations for mapping CPRI flows into various OTN signals [23]. However, the effects on synchronization are still not completely investigated and they are currently scope of the next standard releases.

2.6 RAN Architectures Based on BBU Hotelling

There are various proposed RAN architectures, employing BBU hotelling principles through different flavors. To get some ordered survey of the main solutions, we propose a classification by some key features, namely the BBU placement, the fronthaul transport solution, the BBU implementation. To limit the number of possible architecture variations and at the same time present the more realistic cases, we consider a pure LTE and macro-cell based coverage scenario. The same considerations can be extended to other RATs, with no complication.

2.6.1 Classification on BBU placement

We classify four main options to place BBUs across a RAN infrastructure, that reflect the progressive evolution from a traditional BS (i.e., not employing BBU hotelling), towards most advanced hotelling solutions [24]. The four architectures are depicted in Figures 2.2, 2.3, 2.4, 2.5, and described in the following.

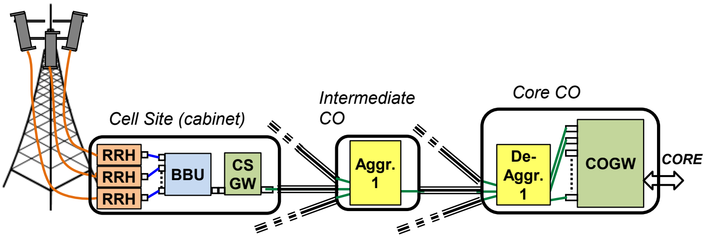

Figure 2.2: BBU placement: *All-in-one BS* (No hotelling).

2.6.1.1 All-in-one BS (No Hotelling)

This (Fig. 2.2) is the traditional widely-deployed macro-cell RAN architecture. The BS is placed in the cell site, so no BBU hotelling is performed. It can be implemented either as a single form factor device or, more frequently, as separated BBUs and RRHs, both located in the same cell site cabinet. Since they are connected by dedicated short-range D-RoF links, there is no fronthaul transport needed over the RAN. The RRHs are connected to respective transmission and reception antennas via dedicated coaxial cables, which carry the radio-frequency analog signals. Their length can be up to some meters or tens of meters, depending on the typology of cell site and distance between the antennas and the equipment cabinet. As a result, they experience a non-negligible power loss, which must be properly taken into account by increasing the output power of RRHs (approximately 3 dB of margin is reported as average value in literature [25]).

The backhaul traffic generated by every BBU in the cell site is aggregated by a Cell Site GateWay (CSGW), and sent towards the core, through the remaining RAN portion. Traditionally, such gateway implemented legacy transport technologies, e.g., SDH/SONET signals, but currently a relevant shift to Ethernet transport solutions is taking place, pushed by their lower costs and higher capacities.

2.6.1.2 Distributed BS (No Hotelling)

This (Fig. 2.3) is a first preliminary step towards BBU-RRH splitting, serving as a basis for all BBU hotelling solutions, though

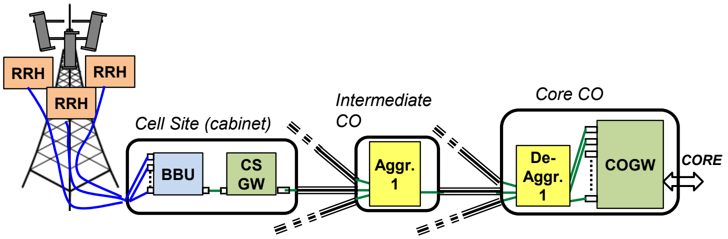

FIGURE 2.3: BBU placement: *Distributed BS* (No hotelling).

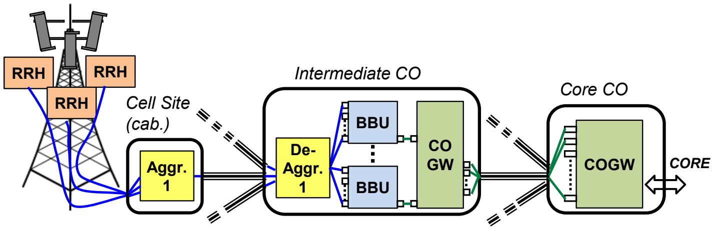

FIGURE 2.4: BBU placement: *BBU hotel at First CO.*

it is not a hotelling architecture. In fact, BBUs remain in cell site cabinets, while RRHs are moved apart from cabinets and directly attached to antennas. To do this, RRHs must be implemented as Remote Radio Heads, i.e., stand-alone devices (embedding their own power and cooling subsystems) designed for operating in outdoor environmental conditions. Placing RRHs close to antennas enable to reduce their power consumption with respect to the All-in-one BS, because the power margin for coaxial cable losses is much lower, virtually negligible. In addition, the equipment located into the site cabinet is partially reduced, with consequent benefits due to the smaller required space occupancy.

2.6.1.3 BBU Hotel at First CO

This (Fig. 2.4) is the basic BBU hotelling solution. Through the RAN infrastructure, each cell site is directly connected to a first-hop CO, here denoted as "first" CO. BBUs of different cell sites are placed in their first COs, which become hotel sites,

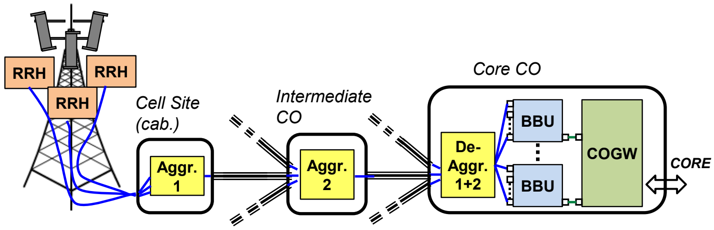

Figure 2.5: BBU placement: *BBU hotel at Higher-level CO.*

while RRHs are remotized to antennas (as in the Distributed BS case). The RAN portion between cell sites and first COs is used to transport the fronthaul, with different possible solutions, depending on fiber availability and required costs.

For instance, instead of dedicating a separate point-to-point connection for each fronthaul flow, it is possible to multiplex into few optical links several flows generated by a cell site, towards the first CO. A more detailed discussion of fronthaul transport is reported in the dedicated subsection. A very important consequence of this architecture is that the cell site cabinet space occupancy is greatly reduced, while the CO can more efficiently manage a large number of hotels, resulting in relevant benefits on energy consumption and costs savings, as introduced previously.

2.6.1.4 BBU Hotel at Higher-Level CO

BBUs can also be placed at higher-level COs (Fig. 2.5), instead of at first COs. With respect to the previous case, the architecture of cell sites and hotel sites is unchanged, but there can be different solutions for the fronthaul transport. In fact, first COs become intermediate “transit” nodes for fronthaul, hence their implementation is an additional degree of freedom, depending on how the fronthaul is aggregated towards higher-level COs.

In general, this architecture allows higher concentration of BBUs into single hotels, that leads to more energy and costs savings (reduction of the number of hotel sites) and potentially higher pooling gains. For these reasons, such solution has a great relevance about the concept of Fixed-Mobile Convergence (FMC), because the reduction of the number of hotels can be seen as part of the more general reduction of the number of active nodes, also known as “node consolidation”, which is one of the key enablers of FMC.

In this case, this makes possible to implement intermediate COs as completely passive nodes, with great reduction of energy and costs. Also the fronthaul latency constraint plays a critical role, more critical than in the previous first-CO hotelling architecture, because of the higher distances between cell sites and hotel sites. This could make this solution infeasible in some scenarios featuring more spread geographical distribution of COs, for instance rural coverage cases.

2.6.2 Classification on Fronthaul Transport

In BBU hotelling architectures, the fronthaul can be transported over the RAN by using existing network infrastructure at different levels, e.g., cable, wavelength, sub-wavelength (bitstream). To do so, some transport technologies can be used (e.g., OTN, Ethernet). A distinction is made between passive and active solutions, where passive means that no additional energy-consuming equipment is needed at end-points or intermediate nodes, and active indicates the converse. Passive solutions lead to lower operational costs, because they require less maintenance and they are more robust against failures. In the following we present a general classification of the main transport solutions, but it is implied that more cases can be obtained by mixing together some of these solutions.

2.6.2.1 Dedicated Point-to-Point

In this case, fronthaul is transported by Point-to-Point (PtP) dedicated links, possibly reusing existing RAN infrastructure elements, as trenches, ducts and multi-fiber cables. Although any fronthaul interface can be used, already defined public standards are preferred (mostly, CPRI). If intermediate nodes are present, PtP connection are routed via passive interconnection fabric, namely Optical Distribution Frames (ODF). With these assumptions, this method is classified as passive. The main drawback is in the high number of required fibers (each carrying a separate flow for each RRH-BBU CPRI port pair).

2.6.2.2 Passive WDM

Fronthaul flows are transmitted on separate wavelength channels, by means of transceivers that operate according to the Wavelength Division Multiplexing (WDM) technology. This makes possible to multiplex several wavelengths into a few fibers, via a passive WDM multiplexer placed in each cell site. At the hotel site, the incoming wavelengths are separated by passive demultiplexers and sent to separate CPRI ports of BBUs, equipped with WDM transceivers. Among the different WDM technologies, Coarse WDM (C-WDM) is regarded as the most practical, because it is best suited for outdoor equipment (it does not require temperature control) and has lower costs, with respect to Dense WDM (D-WDM) solutions. With C-WDM, up to about 16 wavelengths can be multiplexed into a single fiber. Since the maximum number of RRHs in a cell site (considering several sectors, RATs, antennas) is in most cases smaller than this, it is possible to aggregate the whole cell site fronthaul into a single fiber. This solution exhibits relevant advantages with respect to dedicated PtP, because the amount of fiber is reduced with no impact on energy consumption. A drawback is that a proper inventor management is needed, in order to align wavelengths of end-points transceivers for each fronthaul flow.

2.6.2.3 TDM-over-WDM (OTN)

In this solution, several fronthaul flows are end-to-end mapped into wavelengths, by Time Division Multiplexing (TDM), over WDM. One of the best ways to do this is employing the OTN (Optical Transport Network) technology. At both cell sites and hotel sites, there are OTN mu/demultiplexers (commercially denoted also as “wrappers”), which transport fronthaul flows over the OTN signal hierarchy, i.e., CPRI signals are mapped into OTN low-level containers, which are multiplexed into high-layer signals and transmitted on different wavelengths.

Differently from the previous passive case, this solution can easily manage D-WDM transport, in which each fiber can carry up to about 40 to 90 wavelengths, even with bidirectional transmission. Moreover, it is automatically endowed with control and management functions, without having to resort to external monitoring devices. In presence of intermediate nodes, the fronthaul infrastructure becomes a complete OTN network, in which cross-connect devices can be added, in order to get more advanced transport functions, e.g., reconfiguration of routes, protection with redundant paths. They can be implemented either as electronic switches (i.e., OTN wrappers), or all-optical switches, e.g., Arrayed Wavelength Grating (AWG) or Optical Add/Drop Multiplexers (OADM), whose energy consumption is much lower.

Unfortunately, the drawback of such active architecture is the relevant additional energy consumption due to OTN devices and, particularly, the much higher costs. These features make it unattractive for operators, at least for a short-term deployment and for smaller RAN instances. Nevertheless, they are promising future-proof solutions in the long-term.

2.6.3 Classification on BBU Implementation

2.6.3.1 BBU Stacking

This is the basic case, in which original BBUs simply are placed in centralized hotel sites, rather than cell sites, without changes on their implementation. Thanks to the centralization into the hotel site, collateral systems, i.e., cabinets, racks, power supplying (voltage transformers, rectifiers, back-up batteries), cooling (ventilation or air conditioning), backplanes, aggregation gateways, can be re-implemented in a way such that relevant energy and costs savings can be achieved.

If the association between each BBU and its cell site is preserved, the BBUs re-placement does not modify their number, so their cost and energy contribution is unchanged. However, commercial BBUs have a capacity, in terms of maximum number of CPRI flows, therefore maximum number of managed RRHs, which is designed to serve maximum-configuration cell sites, such that a single BBU is sufficient for every cell site, for each RAT.

This leads to a capacity underutilization, because some CPRI ports are not used for typical cell configurations. In case of BBU stacking, the 1:1 association between BBUs and cell sites can be relaxed (within certain limits), thus achieving a higher BBU capacity utilization. The consequence is that the total number of BBUs can be reduced, with further cost and energy gains.

2.6.3.2 BBU pooling

In this architecture, BBUs placed at hotel sites are radically different from BBUs at cell sites. They can be implemented in several ways (for instance, maintaining a modular structure, or as monolithic devices), but the main feature is that they share some portion of their hardware resources. For this reason, we refer to the BBUs in a single hotel as a single BBU “pool”. Resource pooling can be of two kinds: static or dynamic.

Static pooling occurs when a processing function is performed by a single resource element, rather than replicating it across many separate elements. For instance, some low-level physical layer functions requiring intensive vector-based processing (typically joint multi-cell and multi-user algorithms) are often performed by dedicated hardware. By consolidating such hardware, it is possible to improve computational performance and in some cases reduce energy consumption. The pooling is static, because these functions are fixedly assigned to hardware, at the design stage.

Dynamic pooling consists in allocating computational resources “on demand” for processing signals of different cells. To some extent, resources usage can be adapted to the load of each controlled cell, i.e., reserving more resources to high-loaded cells with respect to low-loaded ones. In this way, differently from a traditional BBU exhibiting almost constant energy consumption as function of the load, the BBU pool energy can better scale with it.

Adapting the RAN consumption to the traffic load is seen as the most promising way of improving energy efficiency. For non-pooling based BBUs (i.e., traditional no-hotelling and pure BBU stacking) “sleep-states” techniques are being investigated. They consist in entirely switching off some cells (both BBUs and RRHs) when they are off-peak and providing them coverage through adjacent cells. BBU pooling can provide a smoother adaptation, thanks to dynamic pooling, without sacrificing coverage, because critical RRHs can be kept in on-state while correspondent BBU resources are drastically reduced.

Pooling techniques can be further differentiated by their adaptation time scale. For instance, by slow reconfiguration of pool resources (order of minutes or hours) it is possible to adapt to predictable periodic variations of traffic load occurring with daily or weekly periods, denoted as “tidal effect”. For instance, a well-known tidal effect is due to the almost complementary traffic load patterns experienced by residential and office areas, during the day. If the network planning is that each BBU manages both kinds of areas, relevant pooling gains can be achieved.

On the other side, by fast reconfiguration of the pool (orders of seconds, or comparable with radio frames) it is possible to adapt resources to quasi-instantaneous unpredictable traffic variations, for instance, caused by user behaviors and/or packet schedulers that operate according to the instantaneous channel condition. Even higher pooling gains are expected in this case, of course at the expense of much more complex BBU hardware.

2.6.3.3 BBU Virtualization

This approach is envisioned as the future evolution of BBU hotelling, commonly denoted as Cloud RAN (C-RAN). BBU processing functions are fully virtualized over a distributed hardware/software platform de-centralized over several hotel sites, hence the term “cloud”. In this way, the resources assigned to each cell are no more statically located inside a fixed hotel, but can be dynamically re-assigned to different hotels, thus implementing a kind of inter-hotel BBU pooling [26, 27].

Such paradigm can be seen as a particular case of a more general concept, denoted as Network Function Virtualization (NFV) [28]. For this architecture, an advanced underlying RAN transport infrastructure is required, which allows inter-hotel communication with very low latencies and high reliability, and online traffic reconfiguration. BBU virtualization is regarded as a way to further reduce RAN costs, because the virtualized platform can be built over general-purpose commodity equipment, instead of specialized hardware, potentially increasing the competition in the market and enabling multi-vendor interoperability.

The strong point of BBU virtualization is enabling full load adaptation and balancing of BBU processing. For instance, if some cells are heavily loaded, more resources can be reserved, possibly involving more than one hotel location. Similarly, during off-peak periods, the whole fronthaul of a certain area can be re-routed towards a single hotel site and remaining hotels can be switched off. From these considerations, even bigger energy gains are expected with respect to single-hotel pooling solutions.

However, differently from pooling, for which some commercial solutions are starting to be deployed, BBU virtualization is still a pure research topic.

2.7 Conclusion

This chapter has introduced the main motivations behind the novel BBU hotelling technique, together with its potential advantages and critical issues for implementation in radio access networks. The main issue comes from the introduction of a new kind of network traffic, called fronthaul, which has radically different features respect to backhaul traffic and enforces strict transport requirements. Then, a survey of BBU hotelling architectural solutions has been presented, with suitability for the implementation on fixed/mobile converged access and aggregation networks. How and where to place the BBUs and which transport technology has to be used for fronthaul are important questions in the design of future converged network architectures. Some of such questions will be analyzed and answered throughout the following chapters.

Energy Comparison of Different Mobile Backhaul and Fronthaul Architectures

3

In this chapter, we present an energy-consumption comparison of different optical network architectures for both mobile backhauling and fronthauling. The main focus is about how much energy savings are allowed in a macro-cells based RAN, built on traditional mobile aggregation infrastructures with no fixed/mobile convergence, under the different combinations of choices for BBU placement and aggregation technologies that are described in the previous chapter. Finally, the energy results are combined with the analysis of the total length of fiber needed for each solution, to identify some basic tradeoffs.

3.1 Introduction

Since energy efficiency represents a true sore spot for mobile access networks, optical communication technologies play a crucial role, because they can guarantee high capacity and long reach, with a substantial energy consumption reduction. Moreover, the novel C-RAN concept has been gaining attention for the last years [16], which is mainly enabled by the BBU hotelling technique.

The enhanced flexibility of C-RAN promises not only significant energy consumption savings, but also improved radio performance, reduced CapEx, OpEx, footprint, installation/intervention times [29]. To fully exploit such potential, a drastic re-engineering of the access infrastructure is needed, which will be one key of future fixed/mobile converged networks [30]. However, in practice, it is likely to be rather a gradual process.

Therefore, if an operator decides to implement C-RAN in a short term, exploiting the already existing infrastructure (i.e., outside plant and telecom offices) and with available commercial technologies, how much energy savings are allowed by C-RAN? This work tries to answer such question, focusing only on network-level architectural choices, i.e., excluding the study of radio performance improvements that C-RAN would make possible. To do so, several possible mobile backhaul and fronthaul architectures are compared, in a real-world infrastructural scenario and over a RAN based on macro-cells.

3.2 Considered Backhaul and Fronthaul Architectures

We consider an existing dark-fiber aggregation infrastructure featuring a tree topology with a three-level hierarchy: Cell Sites (CS) composed by equipment cabinets and antennas, Intermediate Central Offices (ICO) and the Core Central Office (CCO), which acts as a point of presence towards the core network. Depending on the placement of BBUs and RRHs, we consider the four cases that have been described in Section 2.6.1 of the previous chapter, specifically adapted for the scenario of this analysis.

These four cases are depicted in Fig. 3.1 and summarized in the following.

- *a-BS (All-in-one BS)*: BBUs and RRHs are enclosed in the same cabinet at CS, analog radio signals are sent to antennas via coaxial feeders (pure backhauling RAN);

- *d-BS (distributed BS)*: BBUs remain in the CS cabinet, but RRHs are installed next to antennas, CPRI is exchanged via dedicated short fibers (pure backhauling RAN);

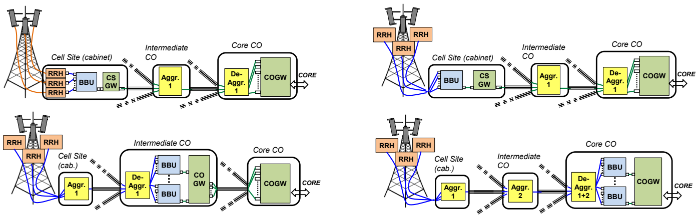

FIGURE 3.1: BBU/RRH placement cases (top-left: a-BS; top-right: d-BS; bottom-left: FH-I; bottom-right: FH-C).

- *FH-I (FrontHaul to Intermediate)*: FH is up to ICOs, where BBUs are stacked, backhaul is up to the CCO (mixed fronthauling/ backhauling RAN);
- *FH-C (FH to Core)*: FH is up to the CCO that hosts BBUs (pure fronthauling RAN). Wherever several backhaul or fronthaul flows are collected in a single site, aggregation can be performed towards higher levels. For backhaul, we assume end-to-end aggregation performed by Ethernet gateways, respectively: Cell Site GateWay (CSGW) and Central Office GateWay (COGW), for both Intermediate and Core COs.

For fronthaul, we differentiate among three kinds of aggregation.

- *PtP*: no aggregation, all input links from the lower level are Point-to-Point connected to output links towards the higher level (typically, through a passive optical distribution frame);
- *WDM*: inputs at different wavelengths are packed into few output fibers, by means of passive wavelength division multiplexing;
- *OTN*: inputs are time-domain multiplexed into few higher-bitrate signals, through an Optical Transport Network (OTN) device.

In our analysis, we consider the most significant combinations of BBU/RRH placement and aggregation. The rationale

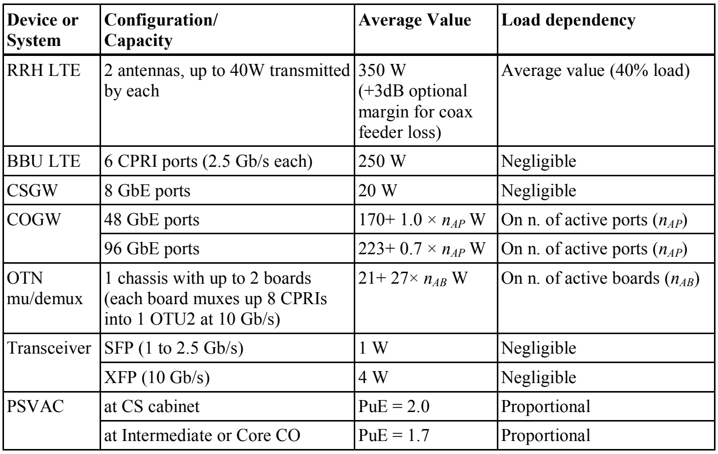

Device or System	Configuration/ Capacity	Average Value	Load dependency
RRH LTE	2 antennas, up to 40W transmitted by each	350 W (+3dB optional margin for coax feeder loss)	Average value (40% load)
BBU LTE	6 CPRI ports (2.5 Gb/s each)	250 W	Negligible
CSGW	8 GbE ports	20 W	Negligible
COGW	48 GbE ports	170+ 1.0 × n_{AP} W	On n. of active ports (n_{AP})
	96 GbE ports	223+ 0.7 × n_{AP} W	On n. of active ports (n_{AP})
OTN mu/demux	1 chassis with up to 2 boards (each board muxes up 8 CPRIs into 1 OTU2 at 10 Gb/s)	21+ 27× n_{AB} W	On n. of active boards (n_{AB})
Transceiver	SFP (1 to 2.5 Gb/s)	1 W	Negligible
	XFP (10 Gb/s)	4 W	Negligible
PSVAC	at CS cabinet	PuE = 2.0	Proportional
	at Intermediate or Core CO	PuE = 1.7	Proportional

Figure 3.2: Energy consumption values.

is that, whenever possible, intermediate offices should be passive, to reduce energy and costs (central office consolidation). By naming each combination as “Placement/Aggr.1/Aggr.2”, we consider the following eight architectures: a-BS/PtP, d-BS/PtP, d-BS/WDM, FH-I/PtP, FH-I/WDM, FH-I/OTN, FH-C/PtP/PtP, FH-C/WDM/PtP.

3.3 Energy Consumption Model

For each architecture, the total energy consumption is computed as the sum over all active devices and systems. All the energy models are detailed in the table of Fig. 3.2. The power values are averaged over typical commercial devices. Some devices (i.e., COGW and OTN equipment) can be properly configured for partially adapting to load, leading to a baseline-plus-proportional consumption. For each active site, collateral Power Supplying, Ventilation and Air Conditioning (PSVAC) systems are accounted by their Power usage Efficiency (PuE), i.e., the ratio of the total power to the one consumed by network equipment.

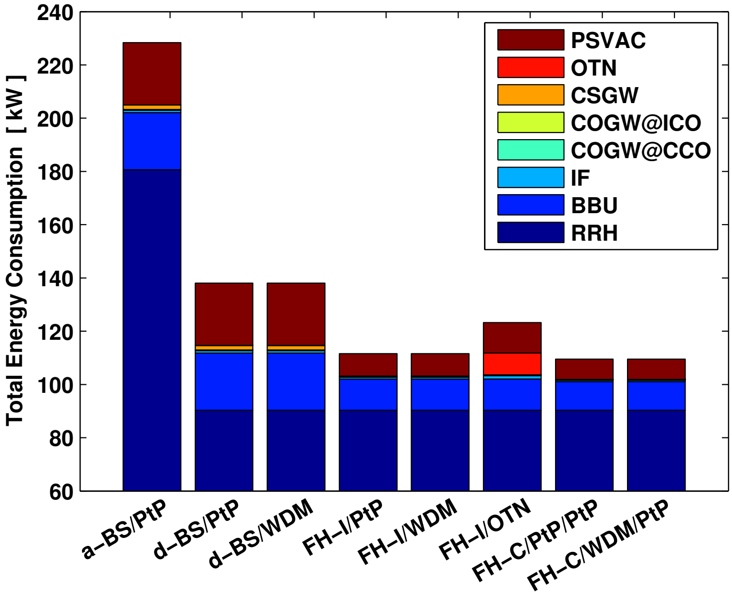

FIGURE 3.3: Total Energy Consumption, with breakdown of different contributions, for all the architectures.

3.4 Results and Discussion

As a real-world case study, we analyze the architectures over a portion of the aggregation infrastructure owned by the operator *Orange*. Specifically we consider a 15-km square coverage area, centered in one of the top-ten French cities (as representative of a typical urban area), and comprising 86 CSs, 13 ICOs and one CCO. We assume that the number of available fibers per link is unlimited.

Fig. 3.3 shows the total energy consumption for all the specified architectures, with a breakdown of the different contributions. In all cases the contribution of RRHs is largely dominant, which justifies the current research efforts aimed at improving the energy efficiency of such devices.

With respect to the a-BS architecture, which needs an average additional 3 dB margin to balance the RF power losses introduced by coaxial feeders (around 30 m long), all other architectures benefit from a halved RRH energy contribution, with a very relevant effect on the total consumption.

The energy consumed by BBUs is the second relevant contribution. Passing from a-BS and d-BS, to FH-I and FH-C architectures, such contribution progressively decreases, thanks to the concentration of several BBUs into a single hotel. This leads to a more efficient utilization of the CPRI input ports, thus reducing the total number of BBUs needed for managing the same number of RRHs. This gain becomes more relevant when the hotel size increases, i.e., when more BBUs are concentrated, that happens when the fronthaul extends towards higher levels of the aggregation infrastructure.

We identify this as a first, important energetic benefit resulting from the so called “BBU consolidation”. The third important energy contribution comes from the collateral PSVAC systems. A substantial reduction can be seen, from a-BS and d-BS, to FH-I and FH-C. It is caused by the combination of both the reduction of per-site consumption (mainly driven by the lower number of BBUs) and the higher efficiency (lower PuE) resulting from moving network equipment from small cell site cabinets to large telecom office rooms. We see this as another strong energetic motivation for adopting fronthaul architectures.

The remaining contributions due to active aggregation equipment are almost negligible for backhaul (i.e., Ethernet gateways), even though a slightly higher contribution can be noted for fronthaul aggregation in the OTN architecture. This is a consequence of how energy consumption of aggregation equipment scales as a function of the transported traffic. For a-BS and d-BS architectures, energy consumption approximately scales as the number of CSs, because it is dominated by the CSGWs, each of them exhibiting a constant energy contribution. For the OTN architecture, it scales as the number of BBU-RRH pairs, because each pair requires a separate CPRI flow, so a separate pair of optical ports, and energy consumption of OTN devices has a contribution which is proportional to the number of used ports. Hence, we can observe here that, in the considered scenario, the choice between passive and active aggregation is more critical for fronthaul than for backhaul.

We finally remark that, even though they do not exhibit any

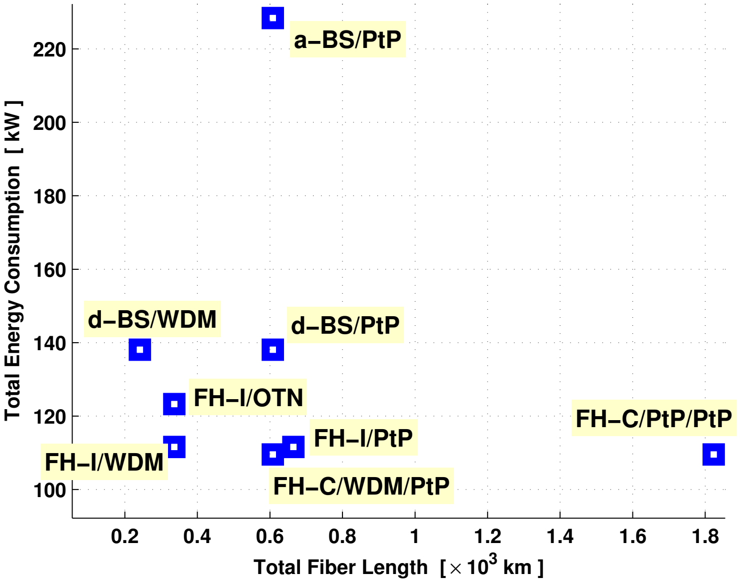

FIGURE 3.4: Total Fiber Length vs. Total Energy Consumption, for all the architectures.

difference in terms of energy consumption, some architectural variants (i.e., PtP vs. WDM) still have a very different impact on infrastructural cost. To investigate also a “resource-consumption” metric, we compare the architectures in terms of the total fiber length, i.e., the sum of the length of fibers used in every link. Results are shown in Fig. 3.4, on a fiber length versus energy plane. For all architectures, the superiority of the WDM variant with respect to PtP is now obvious, because of the substantial lower number of fiber required in case of WDM aggregation. Such difference becomes more relevant for FH-C, because of the longer distances of the fronthaul transport. The analysis also evidences that a-BS/PtP and d-BS/PtP are Pareto-dominated (i.e., for each, there is at least another architecture that guarantees both lower energy and lower total fiber length), therefore they should never be taken into consideration, at least for the considered joint metrics.

3.5 Conclusion and Future Evolution

In view of a gradual evolution towards C-RAN, we have shown that a first-step fronthaul implementation on a macro-cell based RAN can lead to valuable energy savings (up to 40 to 50%), thanks to an intelligent re-placement of RRHs and BBUs, which is already feasible over existing infrastructures with available commercial technologies.

However, this is only the beginning. The next step consists in evolving the hotel equipment, from many stacked and independent BBUs, to a single optimized device in which hardware resources are shared among all pooled BBUs. A first substantial reduction of the baseline consumption comes from the consolidation of some subsystems into a single one (e.g., power, fans).

Moreover, the BBU pool can be designed to allow processing resources to be dynamically allocated to different cells/RRHs, thus achieving a smoother adaptation to traffic load. Relevant amount of pooling gains can be obtained by exploiting statistical variations of cells' load over time and space, based on the "tidal effect", i.e., the almost complementary traffic load patterns experienced by different clusters of cells (typically belonging to residential vs. office areas), on a daily period.

Orange internal studies have shown that the tidal effect can lead to nearly 50% pooling gains, even with small pool sizes ($<$ 20 cell sites), especially in dense urban areas. Some other studies indicate that gains up to 80% can be obtained if resources are dynamically allocated on a much shorter time basis, thus getting a full load adaptation. Since the pooling gain proportionally translates to an energy consumption reduction, most of the current C-RAN research efforts go towards such directions.

Energy-Efficient BBU Placement over a WDM Aggregation Network 4

In this chapter, we continue the investigation of the benefits that can be achieved by BBU hotelling solutions in terms of energy consumption, from a overall network perspective. With respect to the previous chapter, we shift our focus to an optimization approach that can be applied on a more generic fixed/mobile converged network. To do so, we define some architectures for a FMC aggregation network based on WDM, in which BBU hotels can be potentially located into any of the intermediate nodes of the aggregation infrastructure. Therefore, a novel network optimization problem, jointly involving BBU placement and traffic routing, is defined and investigated.

4.1 Introduction

Energy efficiency is gaining more and more importance as a fundamental driver of the technological progress, for both economic and environmental reasons. During the last decade, thanks to the prominent role that the Internet has gained in our society, the Information and Communication Technology (ICT) have been experiencing an impressive evolution. As a collateral consequence, the impact on human energy consumption has been evolving at the same rate. Today, it is estimated that ICT accounts for roughly 1–2% of the worldwide energy consumption [31].

Currently the majority of energy consumption of the Internet is consumed in fixed and mobile access and backhaul (i.e., aggregation) networks, and it is forecast to be dominant also in the next years [32]. This represents a relevant issue for network operators, which already have to cope with old infrastructures exhibiting low cost efficiency and capacity bottlenecks. Therefore, such networks must evolve in a disruptive way, in which not only costs reduction, but also energy efficiency must play a central role as design metric.

In general, as pointed out in [33], the WDM technology enables to implement energy efficient network architectures directly at the design stage. An example of this approach is given in [34], where four IP-over-WDM architectures for core/backbone networks are analyzed, showing that the highest energy efficiency is obtained by avoiding electronic processing at intermediate nodes, thanks to the use of transparent architectures, which enable transit traffic to be "bypassed" directly in the optical domain (optical bypass).

Shifting the focus on access networks, in [35] the authors show that optical access solutions based on Passive Optical Networks (PON) are more energy efficient with respect to classical point-to-point active optical networks, in terms of consumption per average access bit rate. For this reason, WDM-PON solutions appear as the ideal candidate for next-generation optical access/backhaul networks. Differently from core/backbone networks, most of ongoing research regarding WDM-PONs focuses on energy savings that are achievable by properly exploiting network adaptation to instantaneous traffic load variations. For instance, in [36] the authors propose some power management schemes that allow to get about 40% of energy savings, by dynamically putting in sleep mode some Optical Line Termination (OLT) cards during low-traffic periods of the day.

When dealing with FMC networks, similar considerations can be combined with architectures that are optimized for the particular features required by mobile backhaul networks. For instance, in [37] the energy efficiency of a converged wireless/optical access network is evaluated, where the wireless part coverage is given by LTE femtocells and the optical backhaul consists of PON technology. The results show that relative higher energy savings (around 10 to 27%) can be achieved by utilizing low-power modes in urban-suburban scenarios. In [38] some energy efficient design schemes and bandwidth-allocation mechanisms are investigated for long-reach converged fiber/wireless networks. Specifically, an evolved dynamic bandwidth allocation scheme is proposed for an architecture based on low-energy states ONUs, that leads to a decrease of the energy consumption up to about 0.3 kWh per day and per ONU belonging to the optical back-end (relative reduction of 25%), with respect to conventional schemes.

As first contribution of this chapter, we introduce and describe a FMC aggregation optical network based on WDM, in which the concept of BBU hotelling is employed. As second one, we devise for such architectures a novel energy efficient BBU placement optimization problem. A preliminary work has been carried out in [24], where a simple multi-stage WDM-PON is considered as FMC network and the number of deployed hotels is minimized. In this chapter, the aggregation network is optimized by minimizing an ad-hoc performance indicator, the Aggregation Infrastructure Power (AIP), and taking into account that some nodes can perform optical bypass and/or active (electronic) traffic aggregation of sub-wavelength traffic demands. Finally, to quantify the tradeoff among BBU consolidation, optical bypass and active aggregation, we propose three different network architectures and compare them in terms of AIP and other defined metrics.

The remainder of this chapter is organized as follows. In Section 4.2, the network scenario is introduced, three different network architectures are presented and the power consumption model is described. In Section 4.3, the energy efficient BBU placement problem is described and formalized as an Integer Linear Programming model (ILP), for each architecture. In Section 4.5, the case study is introduced and the simulation results are presented and commented. In Section 4.6, the conclusion is given.

4.2 Network Architecture and Power Model

In this section, we start by giving a description of the basic features of the considered WDM aggregation network. Then, we stress the importance of the Maximum CPRI Route Length constraint,

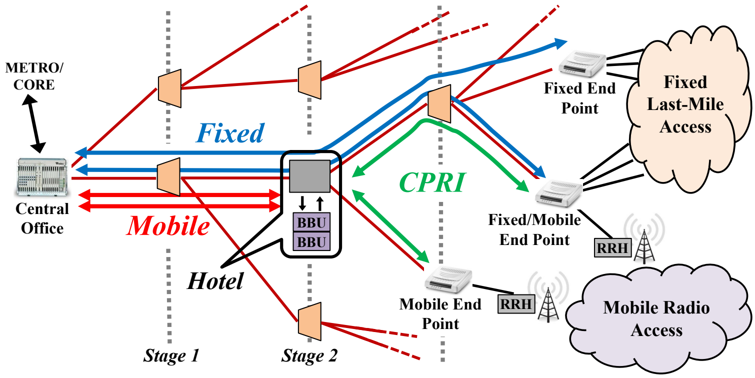

Figure 4.1: Illustrative example of a FMC WDM aggregation network, employing BBU hotelling.

which is fundamental for our problem. We continue by describing in detail three proposed architectures for BBU hotelling on such network. We conclude by detailing the power consumption model which is used in the formalization of the problem.

4.2.1 FMC Aggregation Network Scenario

As depicted in Fig. 4.1, our scenario is a metropolitan area that is covered and served by a FMC WDM aggregation network, whose terminal nodes are the “end points” and the Central Office (CO). The end points act as demarcation points between the aggregation network and the access network. We assume three typologies: Fixed End Points, which collect the traffic of various kinds of last-mile fixed access technologies (e.g., xDSL, TDM-PON); Mobile End Points, which collect the traffic generated by cell (or antenna) sites; and Fixed/Mobile End Points, which aggregate both kind of traffic. As a consequence, such end points represent the clients which require to the FMC network the end-to-end transport service toward the CO. Their geographical location is considered given and independent from our problem.

We consider an already deployed infrastructure of optical fiber links which provides connectivity between each end point and the CO, through a set of intermediate nodes placed at given locations. Each couple of nodes is connected by one pair of monodirectional fiber links, providing a certain capacity divided among a number of WDM channels, or wavelengths, in both directions. We give no further assumption on the given topology, i.e., the fiber infrastructure may be a common multi-stage tree (also known as branch-and-tree), or a multi-stage ring (ring-and-spur), as well as whatever hybrid partially meshed topology.

This choice is due to the fact that, although today typically standard-reach tree-based topologies are considered for the implementation of BBU hotelling architectures, in next future, to ensure greater coverage areas, longer-reach networks might be considered as well. For these kinds of networks, the resiliency to links and node failures has more importance, therefore path- and node-redundant topologies are considered as more valid alternatives. Being the fiber infrastructure given, the degree of freedom is the implementation of nodes, depending on where BBUs are placed and which traffic requests must be satisfied.

In general, a FMC aggregation network offers a transport service for both mobile backhaul and fixed IP traffic. However, since in our case BBUs are in general detached from corresponding cell sites, there are actually two classes of mobile backhaul: IP mobile backhaul, which is the well-known packet traffic between the central office and each BBU uplink interface; and D-RoF backhaul, or fronthaul, consisting in constant bitrate (CBR) CPRI flows among each BBU and the respective RRH located at cell site.

For simplicity, from now on we refers to the three classes of transported traffic respectively as: Fixed, Mobile, CPRI (see Fig. 4.1) For all these kinds of traffic, we assume a common Ethernet-over-WDM transport, therefore each node terminating traffic is equipped with an Ethernet switch featuring both long-reach point-to-point WDM interfaces and short-reach interconnection interfaces. In such a way, switches are also capable of multiplexing different IP and CPRI flows into few wavelengths, pushing toward a reduction of the number of used WDM interfaces and consequently energy savings.

4.2.2 Maximum CPRI Route Length

Differently from IP backhauling, CPRI traffic is much less tolerant to jitter accumulated along the transmission chain[1]. For this reason, we impose that each CPRI flow must be routed on a single end-to-end lightpath, therefore it can not be split among parallel paths and it can not be processed by intermediate electronic switches.

Another important feature of CPRI traffic it that each flow can tolerate up to a maximum value of end-to-end latency, as a consequence of the timing constraints imposed on radio interface physical layer procedures. A detailed discussion of such feature has been presented in 2.5.2. When measured between the transporting WDM circuit endpoints, such value depends on the processing time budget available for the ACK/NACK response of a radio frame (standard dependent), and the elaboration delays due to DAC/ADC, L1/L2, CPRI interface adaptation, and Ethernet switches processing (technology dependent).

Specifically, considering for instance the LTE FDD radio interface, there is an upper limit of 3 ms [21] to the BBU-RRH round trip latency ($RTT_{\text{BBU-RRH}}$), defined as the sum of all uplink (UL) and downlink (DL) processing times (t) of BBU and RRH, plus the processing times of switches at both endpoints ($4t_{\text{SW}}$), plus the round trip propagation delay (2τ)(see Fig. 2.1 for a schematic depiction of such delay contributions.):

$$\begin{multline} RTT_{\text{BBU-RRH}} = 2\tau + t_{\text{RRH,UL}} + t_{\text{BBU,UL}} + t_{\text{BBU,DL}} \\ + t_{\text{RRH,DL}} + 4t_{\text{SW}} \leq 3 \text{ ms} \quad (4.1) \end{multline}$$

In this framework, all these processing delays are fixed, because purely technology-dependent. Therefore, the maximum $RTT_{\text{BBU-RRH}}$ directly translates to a maximum admissible one-way propagation delay for every CPRI flow ($\tau \leq T_D$), or, equivalently, to a maximum length of its route ($L \leq L_D$). In such a way, a Maximum CPRI Route Length constraint must be explicitly taken into account in our optimization problem, whenever physical lengths for all links are given.

---

[1] The transport of CPRI flows over Ethernet is a very challenging task today. However, upcoming standard evolutions (e.g., Synchronous Ethernet) and technical advances go towards such direction.

Values of L_D which typically range from 20 to 40 km are reported in the recent literature [16]. We remark that, when BBU hotelling is applied in our network scenario, the practical values might also be smaller, due to the fact that CPRI flows are transported over the Ethernet layer. Also, physical channel impairments introduced along the optical propagation can pose further severe limitations to CPRI flows. We argue that, since CPRI flows consist of digital data, the channel impairments basically have impact on the BER, so they can be properly compensated by more complex signal processing, leading to increased processing times, thus lower values of L_D.

4.2.3 The Three Network Architectures

In the view of energy analysis, we classify all nodes into two categories: active nodes, which are equipped with at least one power consuming device, and passive nodes. By definition, terminal nodes and intermediate nodes equipped with one or more BBUs are active, because they terminate, route and process traffic in the electrical domain. Intermediate nodes which do not terminate any traffic are made up only of transparent wavelength routers and mu/demultiplexers, hence they are passive. Their specific implementation depends on factors that are out of the scope of the present study (e.g., cost budget, degree of reconfigurability). For sake of simplicity, we model them as capable of routing any wavelength to any output.

According to the ways BBUs are placed and active intermediate nodes can be implemented, we consider three different network architectures.

1. *Bypass.* Each BBU can be placed in any node. For intermediate nodes, this implies they become active and also equipped with an electronic switch that terminates and aggregates all traffic destined to the BBU hotel. Remaining transit traffic is bypassed directly in the optical domain (the node is optically transparent to transit traffic).

2. *Opaque.* Each BBU can be placed in any node. For intermediate nodes, this implies they become active and equipped with an electronic switch that terminates and aggregates all

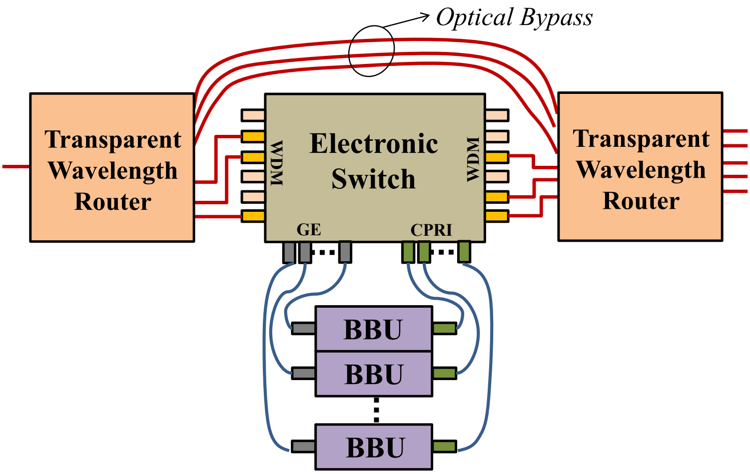

Figure 4.2: Implementation detail of an active intermediate node, i.e., equipped with a BBU hotel.

passing traffic, included traffic which is not destined to the hosted BBUs (the node is optically opaque to all traffic).

3. *No-Hotel.* Hotelling is not performed, i.e., each BBU is placed in its cell site. Therefore, all intermediate nodes are passive and active aggregation can be only performed end-to-end.

The implementation detail of an active intermediate node is depicted in Fig. 4.2. An electronic switch is equipped with both long-reach interfaces, i.e., WDM, and short-reach interfaces, i.e., Gigabit Ethernet (GE) and CPRI. The incoming lightpaths are mu/demultiplexed by wavelength routers and enter the switch via the WDM interfaces. IP traffic destined to a hosted BBU is extracted and sent over the GE port, while the D-RoF (fronthaul) is collected by the switch from the CPRI interfaces, to be mapped into one or more output lightpaths. In case of Bypass architecture, transit traffic flows can be optically bypassed, by directly connecting the two wavelength routers. In case of Opaque architecture, all incoming lightpaths are terminated and processed by the switch.

4.2.4 Power Consumption Model

For the energy efficiency analysis of the proposed FMC architectures, we introduce as performance indicator the Aggregation Infrastructure Power (AIP). It is defined as the sum of the power contributions of solely the equipment that can be considered as part of the aggregation infrastructure. Specifically, the AIP includes two kinds of contributions.

The first one comes from all aggregation network devices, i.e., end switches (placed at the CO and at the end points of the network) and intermediate switches (placed in intermediate hotel nodes). The second contribution is caused by the "housing" equipment, which does not perform network functions, but is needed by the infrastructure for guaranteeing the proper operational conditions (mostly, power supplying and cooling systems). The contributions by mobile access devices (BBUs and RRHs), fixed access devices (e.g., DSLAMS, PON OLTs) and CO gateways (interfacing with core network) are not included in the AIP, because they are not part of the aggregation network, but clients of it.

For all the switches, a common power model is adopted, namely a universal scalable switch model. The consumption is modeled as the sum of a baseline (traffic-independent) term, plus a variable quantity which is linearly proportional to the number n_p of ports configured as active.

$$P_S = P_{S,0} + K_S \cdot n_p \tag{4.2}$$

Only long-reach interfaces (WDM) are considered in the port count (n_p), with a proportionality coefficient equal to $K_S = 15\,\mathrm{W}\,/\mathrm{Port}$, which includes the term due to the WDM transceiver and an additional allowance for the smaller consumption of short-reach interfaces (GE, CPRI) [1]. The baseline is fixed to $P_{S,0} = 50\,\mathrm{W}$, including also the term due to embedded cooling systems (typically small fans).

Regarding the power consumption of housing systems, we assume different models, according to the typology of node. Fixed End Points can be deployed as compact premises which do not require external housing resources (beyond that included in the baseline consumption), due to the relatively small size of hosted network equipment, hence we assume zero housing consumption for this kind of nodes. For the same reason, also Mobile End Points with no located BBUs do not require housing power (RRHs are stand-alone devices directly located at the antenna mast). Mobile End Points equipped with BBUs (conventional BS with no hotelling) exhibit a housing contribution which is basically dominated by the BBU. A value equal to $P_{H,\mathrm{BBU}} = 600\,\mathrm{W}$ is assumed, by considering a typical cell-site BBU under mid-to-high load condition, which consumes in average 750 W, and with a PuE (Power Usage Effectiveness) around 1.8.

The housing power model of intermediate hotel nodes is assumed as the sum of a baseline contribution and a variable term.

$$P_{H,I} = P_{H,I,0} + K_{H,I} \cdot n_{\mathrm{BBU}} \tag{4.3}$$

The baseline contribution is $P_{H,I,0} = 500\,\mathrm{W}$, needed for the housing of the shared facility in which the BBU hotel and the intermediate switch are located [39]. According to the power model of switches and their relatively lower power consumptions (with respect to BBUs), intermediate switches can benefit from the shared power supplying system of the facility and, since they are typically equipped with built-in cooling fans, their required housing power can be considered as embedded in the 500 W value and fairly independent from their traffic load. The variable term is proportional to the number of hosted BBUs (n_{BBU}), via the coefficient $K_{H,I} = 100\,\mathrm{W}\,/\,\mathrm{BBU}$.

It is worth noting that, for each hosted BBU, a considerably smaller amount of power is needed for housing, with respect to the stand-alone case, because the devices can benefit from the shared housing resources whose consumption is already taken into account in the baseline term. The described housing model take into account the “scale” gain due to the consolidation of network elements into few nodes, which is one of the principles pushing towards infrastructural convergence.

Finally, the housing contribution of the CO is not considered in our scenario, because the aggregation switch and the optionally colocated BBUs are placed into an existent equipment room, with a negligible impact on the overall housing power. As a consequence, their power consumption is fixed and independent from the network architectures that we intend to compare. A summary of all the described power model parameters with their values is given in Table 4.1.

Table 4.1: Summary of power model parameters and their values.

Description	**Parameter**	**Value**
Switch: baseline	$P_{S,0}$	50 W
Switch: variable	K_S	15 W /Port
Housing of stand-alone BBU	$P_{H,\mathrm{BBU}}$	600 W
Housing of interm. hotel: baseline	$P_{H,I,0}$	500 W
Housing of interm. hotel: variable	$K_{H,I}$	100 W / BBU

4.3 Energy Efficient Optimization Problem

Establishing which of these FMC aggregation architectures is most energy efficient is not a trivial task. In fact, the following observations can be done:

- the No-Hotel architecture has the smallest number of active nodes (only terminal nodes), but with respect to others, it offers no energy savings deriving from the consolidation of multiple BBUs into single hotels;
- for every intermediate hotel node, the power consumption due to electronic processing in the Opaque case is greater than in the Bypass case, because a greater amount of traffic must be terminated and processed. However, a more efficient aggregation of the traffic towards the CO can be performed, leading to a potential reduction of the number of wavelengths, with a consequent reduction of power.

Therefore the question is: in such FMC network, how and how much do BBU consolidation, optical bypass and active aggregation in intermediate nodes impact on overall energy efficiency? To answer this, for each of the three proposed architectures, we define an optimization problem as follows.
Given: the network topology (nodes connectivity and links lengths), the set of traffic demands, the links capacity (number of wavelengths), the traffic capacity (bitrate) of each wavelength, the maximum number of BBUs that can be hosted in each intermediate active node (hotel capacity), the Maximum CPRI Route Length (L_D), and the power consumption contributions relative to network devices and housing systems;
decide: the placement of each BBU (for Bypass and Opaque

architectures), together with the Grooming, Routing and Wavelength Assignment (GRWA) of all traffic demands;
to minimize: the Aggregation Infrastructure Power (AIP), as previously defined.

The defined problem can be interpreted as a non-trivial combination of two well-known network optimization problems, namely the Grooming, Routing and Wavelength Assignment (GRWA) [40], and the Facility Location Problem (FLP) [41]. However, here the two problems are not independent, since the placement of BBUs influences both the source/destination nodes of mobile and CPRI traffic and consequently their routing. In addition, the different transport requirement for CPRI flows makes the problem more complex and valuable.

4.3.1 ILP Formulation

For each defined network architecture, we formalize the optimization problem through an Integer Linear Programming (ILP) formulation. It operates on a double-layer network graph, composed by a set N of nodes and two sets, E and V, of directed arcs between couples of nodes. E defines the lower-layer topology of physical fiber links (m, n) laid from node m to node n. V defines the upper-layer topology of virtual links (i, j) representing end-to-end connectivity between nodes equipped with electronic switches, by means of lightpaths established from i to j.

A generic IP (mobile or fixed) traffic demand from a source to a destination node is routed in the upper layer over a sequence of virtual links. Each virtual link provides a certain capacity that can be shared among other requests, thanks to aggregation performed by switches endpoints. Every virtual link is routed in the lower layer over physical links as a number of lightpaths. To reduce end-to-end packet jitter, the routing over virtual links is assumed to be non-bifurcated, i.e., a switch can not split the connection among different virtual links. On the other hand, the routing over physical links can be bifurcated in more lightpaths, because the jitter due to different optical paths can be neglected for this class of traffic.

For CPRI connections, since they are subject the Maximum CPRI Route Length constraint, we introduce a kind of delay-limited lightpaths, or D-lightpaths. Each D-lightpath has a propagation delay not bigger than a maximum value, derived from the CPRI constraint. Because of the strict synchronization jitter requirements, we impose that every CPRI flow can be directly routed only over a single D-ligthpath, thus it can be neither bifurcated, nor processed by intermediate switches. A very important feature of the proposed model, which pushes towards convergence, is that D-lightpaths are not exclusively reserved for CPRI flows, in fact the spare capacity of D-lightpaths can also be filled by IP requests, improving the overall aggregation efficiency. To reduce the model complexity, the physical paths of all D-lightpaths are pre-computed, and the routing of CPRI flows simply consists in associating a D-lightpath with each flow (as for a path formulation).

For sake of simplicity, we assume that each connection request is symmetric, i.e., it requires the same traffic in the downstream and upstream direction. Therefore, we can formalize and solve the problem considering only one direction (e.g., downstream) and the results are straightforwardly replicated in the other one. The case of asymmetric requests can be derived from this formulation with no difficulty.

4.3.2 Input Data

- N is the set of nodes, partitioned into: the central office $\{\text{CO}\}$, the subset N_I of intermediate nodes, the subset N_U of network end points. Fixed End Points belong to $N_F \subseteq N_U$, Mobile End Points belong to $N_M \subseteq N_U$, Fixed/Mobile End Points belong to $N_M \cap N_F$.

- E is the set of physical links, indexed by (m, n), with $m, n \in N; m \neq n$.

- V is the set of virtual links, indexed by (i, j), with $i, j \in N; i \neq j$.

- R is the set of connection requests, partitioned into: the subset $R_{\text{IP},M}$ of IP mobile requests, the subset $R_{\text{IP},F}$ of IP fixed requests, the subset R_{CPRI} of CPRI requests. For each request $r \in R$, $k(r)$ is the associated customer, $t(r)$ is the requested traffic demand.

- Λ is the set of wavelengths, indexed by λ.
- C is the traffic capacity of each wavelength.
- B is the BBU capacity of each active node.
- L_D is the Maximum CPRI Route Length, thus the maximum length of D-lightpaths.
- l_{mn} is the length of physical link (m, n).
- $H_{\text{IN}}(n) = \{m \in N | (m, n) \in E\}$ is the set of incoming adjacent nodes of n, with respect to the physical topology.
- $H_{\text{OUT}}(n) = \{m \in N | (n, m) \in E\}$ is the set of outgoing adjacent nodes of n, with respect to the physical topology.
- $H^V_{\text{IN}}(i), H^V_{\text{OUT}}(i)$ are the equivalent sets defined for the virtual topology.
- $\mathcal{M}$ is a very big positive number.

4.3.3 Pre-computed Data

All the D-paths, i.e., physical paths of total length not exceeding L_D, are derived from the physical topology. For each pair of end nodes (i, j), with $i \neq j$, a number Ω_{ij} of D-paths is obtained, which are indexed by $\omega \in \{1, \dots, \Omega_{ij}\}$. Each triplet (i, j, ω) univocally identifies a D-path, and the set $Q_{ij\omega} \subseteq E$ contains all the physical links composing its route. The set of all enumerated D-paths is denoted by D. Therefore, the following holds:

$$\sum_{(m,n) \in Q_{ij\omega}} l_{mn} \leq L_D; \;\; \forall (i, j, \omega) \in D \tag{4.4}$$

4.3.4 Decision variables

- $P^{ij}_{mn\lambda} = 1$, if virtual link (i, j) comprises a lightpath routed over physical link (m, n), on wavelength λ (binary).
- $v_{ij\lambda}$ = number of established lightpaths on wavelength λ composing virtual link (i, j) (integer).

- $Y_{ij}^r = 1$, if the IP request r is routed over virtual link (i, j) (binary).
- $\overline{v}_{ij\omega\lambda} = 1$, if D-ligthpath (i, j, ω, λ) is established, i.e. routed over D-path (i, j, ω), on wavelength λ (binary).
- $\overline{Y}_{ij\omega\lambda}^r = 1$, if the CPRI request r is routed over D-lightpath (i, j, ω, λ) (binary).
- $x_i^k = 1$, if the BBU of cell site $k \in N_M$ is placed at node $i \in N$ (binary).
- $w_i = 1$, if a hotel is installed in intermediate node $i \in N_I$, i.e., with at least one hosted BBU.

4.3.5 Aggregation Infrastructure Power Minimization

The objective is the minimization of the Aggregation Infrastructure Power (AIP), which can be written as:

$$\begin{aligned} \min \Bigg\{ & P_{S,0} \cdot \big(1 + |N_U|\big) + \big(P_{S,0} + P_{H,I,0}\big) \cdot \sum_{i \in N_I} w_i \\ & + P_{H,\mathrm{BBU}} \cdot \sum_{k \in N_M} x_k^k + K_{H,I} \cdot \Bigg(|N_M| - \sum_{k \in N_M} x_k^k\Bigg) \\ & + K_S \cdot \sum_{i \in N} \sum_{\lambda \in \Lambda} \sum_{j \in H_{\mathrm{OUT}}^V(i)} \Bigg(v_{ij\lambda} + \sum_{\omega \leq \Omega_{ij}} \overline{v}_{ij\omega\lambda}\Bigg) \\ & + K_S \cdot \sum_{i \in N} \sum_{\lambda \in \Lambda} \sum_{j \in H_{\mathrm{IN}}^V(i)} \Bigg(v_{ji\lambda} + \sum_{\omega \leq \Omega_{ji}} \overline{v}_{ji\omega\lambda}\Bigg) \Bigg\} \end{aligned} \tag{4.5}$$

The four terms of the first two lines are respectively: the total baseline consumption of switches located at the CO and at Fixed and Mobile End Points (constant), the baseline power of switches and housing systems installed in intermediate hotel nodes, the housing contribution due to BBUs placed at Mobile End Points, and the sum of housing contributions in intermediate hotel nodes, proportional to the number of BBUs. The last two lines quantify the switches traffic-proportional power contributions for all nodes originating and terminating lightpaths.

4.3.6 Common constraints for all architectures

Routing of virtual links' lightpaths over physical links

$$\sum_{m\in H_{\mathrm{IN}}(n)} P^{ij}_{mn\lambda} - \sum_{m\in H_{\mathrm{OUT}}(n)} P^{ij}_{nm\lambda} = \begin{cases} -v_{ij\lambda} & \text{if } n = i \\ v_{ij\lambda} & \text{if } n = j \\ 0 & \text{otherwise} \end{cases}$$

$$\forall n \in N,\ (i,j) \in V,\ \lambda \in \Lambda \quad (4.6)$$

Routing of IP fixed requests over virtual links

$$\sum_{j\in H^v_{\mathrm{IN}}(i)} Y^r_{ji} - \sum_{j\in H^v_{\mathrm{OUT}}(i)} Y^r_{ij} = \begin{cases} -1 & \text{if } i = \mathrm{CO} \\ 1 & \text{if } i = k(r) \\ 0 & \text{otherwise} \end{cases}$$

$$\forall r \in R_{\mathrm{IP},F}, i \in N \quad (4.7)$$

Routing of IP mobile requests over virtual links

$$\sum_{j\in H^v_{\mathrm{IN}}(i)} Y^r_{ji} - \sum_{j\in H^v_{\mathrm{OUT}}(i)} Y^r_{ij} = \begin{cases} x_i^{k(r)} & \text{if } i \neq \mathrm{CO} \\ x_i^{k(r)} - 1 & \text{if } i = \mathrm{CO} \end{cases}$$

$$\forall r \in R_{\mathrm{IP},M}, i \in N \quad (4.8)$$

Routing of CPRI requests over D-lightpaths

$$\sum_{\lambda\in\Lambda} \sum_{\omega\leq\Omega_{ij}} \overline{Y}^r_{ik(r)\omega\lambda} = x_i^{k(r)}$$

$$\forall r \in R_{\mathrm{CPRI}},\ i \in N | i \neq k(r) \quad (4.9)$$

Wavelength occupation (capacity) of physical links

$$\sum_{(i,j)\in V} P^{ij}_{mn\lambda} + \sum_{(i,j,\omega)\in D|(m,n)\in Q_{ij\omega}} \overline{v}_{ij\omega\lambda} \leq 1$$

$$\forall (m,n) \in E,\ \lambda \in \Lambda \quad (4.10)$$

Capacity of D-lightpaths

$$\sum_{r\in R_{\mathrm{CPRI}}} t(r) \cdot \overline{Y}^r_{ij\omega\lambda} \leq C \cdot \overline{v}_{ij\omega\lambda}$$

$$\forall (i,j,\omega) \in D,\ \lambda \in \Lambda \quad (4.11)$$

Capacity of virtual links

$$\sum_{r\in R_{\mathrm{IP}}} t(r)\cdot Y_{ij}^{r} + \sum_{r\in R_{\mathrm{CPRI}}}\sum_{\lambda\in\Lambda}\sum_{\omega\leq\Omega_{ij}} t(r)\cdot \overline{Y}_{ij\omega\lambda}^{r} \leq$$

$$C\cdot\sum_{\lambda\in\Lambda}\big(v_{ij\lambda}+\sum_{\omega\leq\Omega_{ij}}\overline{v}_{ij\omega\lambda}\big); \qquad \forall(i,j)\in V \quad (4.12)$$

BBU Capacity of active nodes

$$\sum_{k\in N_M} x_i^k \leq B; \qquad \forall i\in N \quad (4.13)$$

One BBU for each cell site

$$\sum_{i\in N} x_i^k = 1 \qquad \forall k\in N_M \quad (4.14)$$

Identification of hotel nodes

$$\frac{1}{\mathcal{M}}\sum_{k\in N_M} x_i^k \leq w_i \qquad \forall i\in N \quad (4.15)$$

4.3.7 Extra constraints for Opaque Architecture

$$\sum_{m\in H_{\mathrm{OUT}}(n)} P_{nm\lambda}^{ij} + \sum_{m\in H_{\mathrm{IN}}(n)} P_{mn\lambda}^{ij} \leq \mathcal{M}\cdot(1-w_n)$$

$$\forall\lambda\in\Lambda,\,(i,j)\in V,\,n\in N_I\,|\,n\neq i\wedge n\neq j \quad (4.16)$$

$$\overline{v}_{ij\omega\lambda} \leq (1-w_n)$$

$$\forall\lambda\in\Lambda,\,(m,n)\in E,\,(i,j,\omega)\in D\,|\,(m,n)\in Q_{ij\omega}\wedge n\neq j \quad (4.17)$$

$$\overline{v}_{ij\omega\lambda} \leq (1-w_n)$$

$$\forall\lambda\in\Lambda,\,(n,m)\in E,\,(i,j,\omega)\in D\,|\,(m,n)\in Q_{ij\omega}\wedge n\neq i \quad (4.18)$$

4.3.8 Extra constraints for No-Hotel Architecture

$$x_i^k = \begin{cases} 1 & \text{if } i = k \\ 0 & \text{if } i \neq k \end{cases} \qquad \forall i \in N,\ k \in N_M \tag{4.19}$$

Equations from 4.6 to 4.15 are common to all the three defined network architectures. Eq. 4.6 enables the routing of lightpaths composing virtual links over the physical topology, by imposing the flow balancing at each node. Eq. 4.7 and 4.8 enable the routing of IP demands over the virtual topology, taking into account that the hotels are not only an outcome of the optimization process, but also source/destination nodes for mobile flows. Eq. 4.9 enables the routing of each CPRI flow over a single D-lightpath, thus implicitly imposing the Maximum CPRI Route Length constraint. Eq. 4.10 ensures that each wavelength in every physical link can be occupied by at most one lightpath. Eq. 4.11 and 4.12 limit the maximum traffic that can be carried by each D-lightpath and each virtual link, where the spare capacity of D-lightpaths (i.e., not used for CPRI traffic) can be exploited by virtual links. Eq. 4.13 limits the maximum number of BBUs that each hotel can host. Eq. 4.14 enforces that each cell site (Mobile End Point) must be associated to exactly one BBU, thus one hotel. Eq. 4.15 is used to identify hotels as nodes which host at least one BBU.

Equations from 4.16 to 4.18 are specific for the Opaque architecture. They are needed to prevent any lightpath from transiting through any active node. Specifically, eq. 4.16 imposes that, for every intermediate active node $(n|w_n = 1)$, any lightpath belonging to a virtual link (i, j) not originated or terminated in n $(n \neq i \wedge n \neq j)$ can not be physically routed over one of its outgoing or incoming links. Eq. 4.17 and 4.18 act similarly for D-lightpaths, in fact if the route of the D-lightpath (i, j, ω, λ) includes at least one physical link passing through the active node n, i.e., (m, n) or $(n, m)|(m, n) \in Q_{ij\omega}$, then such D-lighpath can not be established $(\overline{v}_{ij\omega\lambda} = 0)$.

Eq. 4.19 is specific for the No-Hotel architecture, and it simply forces all BBUs to be placed at their cell sites, without performing hotelling.

4.4 Case Study

As a case study, we estimate the Aggregation Infrastructure Power (AIP) of the proposed network architectures on a multi-stage tree topology. This is expected to be the dominant kind of topology for the typical coverage areas which are considered in the case study. Moreover, by a proper design, a multi-stage tree network could enable intermediate optical transparent switches to be implemented as cyclic Arrayed Waveguide Gratings (AWG), which are currently the most considered solutions for a practical deployment, because of their lower costs with respect to the more advanced Optical Add/Drop Multiplexers (OADM). Several multi-stage tree network instances are randomly generated, while keeping fixed the number of stages (3). As input, the spatial coordinates of Fixed and Mobile End Points are obtained by uniformly scattering their locations over a square coverage area, according to predefined geographical density parameters. Such parameters are differentiated among three typical geotypes: Dense-Urban, Urban and Rural ([42], [43]), and summarized in Table 4.2. In all cases a small percentage of end points (around 20%) is assumed to be both fixed/mobile.

The infrastructural planning process is simulated by a hierarchical clustering procedure which runs the k-means algorithm, with a city-block (or Manhattan) distance metric. For each stage, it returns the location of intermediate aggregation nodes, such that the average distance towards customer nodes is minimized. The number of clusters is fixed to 4 and 2, respectively for the second and third aggregation stage. The procedure allows us to obtain the network instance, i.e., its topology and the lengths of all physical links.

For each network instance, the traffic demands are generated by the following assumptions. Each Mobile End Point collects the traffic of a single macro cell site, providing LTE radio coverage of 3 sectors with 20 MHz bandwidth and 2 × 2 MIMO configuration. Therefore, the corresponding CPRI demand is computed as 6.29 Gb/s [5], including only the CPRI control overhead and not the 8B/10B line coding, which is not necessary in our CPRI-over-Ethernet transport scenario. The corresponding IP demand is a random value with uniform distribution in the range of 300–750 Mb/s (typical mobile backhaul bitrates for such cell site con-

Table 4.2: Case Study parameter values for the three geotypes: Dense-Urban (DU), Urban (U) and Rural (R).

Geotype	**DU**	**U**	**R**
Cell Sites Density (per km^2) [42, 43]	10	3.5	0.15
Households Density (per km^2) [42, 43]	4000	900	300
N. of Mobile End Points ($\|N_M\|$)	30	20	10
N. of Fixed End Points ($\|N_F\|$)	15	7	25
Hotel Capacity (B)	30	20	10

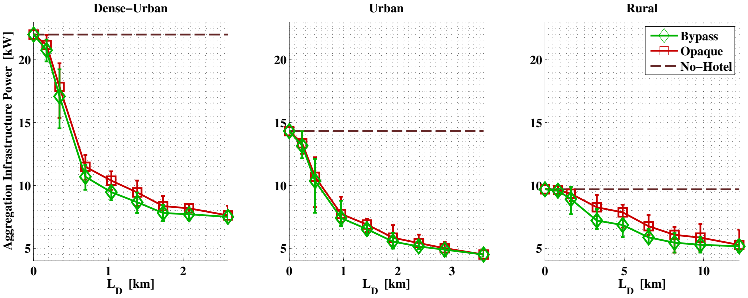

FIGURE 4.3: Aggregation Infrastructure Power (AIP) vs. L_D, for all geotypes.

figuration). The IP demand of each Fixed End Point is a random value with uniform distribution between 10 and 20 Gb/s. Such values are obtained by considering respectively the minimum and maximum density of households for each geotype, an average value of 800 aggregated households per Fixed End Point, a peak access rate equal to 100 Mb/s per household, and a factor 1/5 of statistical multiplexing gain. All traffic demands are generated as integer multiples of a basic granularity, equal to the gross bitrate of the OC-3 signal of the SONET/SDH hierarchy (155.52 Mb/s).

In all cases, the number of wavelengths per link is 40 and each wavelength has capacity equal to 64 OC-3s (about 10 Gb/s). In the following, each value is computed by averaging the results obtained by around 30 independently generated instances, for each geotype and network architecture, and varying the Maximum CPRI Route Length (L_D) and the hotel capacity (B).

4.5 Numerical Results and Discussion

Fig. 4.3 shows the average values of the Aggregation Infrastructure Power (AIP) (i.e., the resulting objective value of eq. 4.5) for the three defined geotypes and network architectures, as a function of L_D. There are also shown the errorbars computed as the distances towards the minimum and maximum values of the data sets. For all geotypes, Bypass and Opaque architectures clearly outperform the No-Hotel. In fact, while in the No-Hotel case the power consumption is independent from L_D, because it is only influenced by the efficiency of routing and aggregation which are performed end-to-end between the CO and the end points, in the two hotelling architectures there is substantial reduction of power consumption. This means that the energy cost that is paid to make some intermediate nodes active is overcompensated by the gain obtained by consolidating BBUs into a few number of nodes. As intuitively expected, the improvement with respect to the No-Hotel case approaches zero for low values of L_D, because BBUs are pushed towards end points and thus no consolidation can be achieved. By increasing L_D, relevant savings of the AIP are observed, until about 60–65% for Dense-Urban and Urban, and 40% for the Rural case. The smaller relative saving in this last case can be justified by the lower geographical density of mobile and fixed access traffic, with consequently a smaller volume of aggregated traffic and fewer BBUs to consolidate. Nevertheless, in all cases the results confirm the promising energy-reduction opportunity offered by the proposed BBU hotelling FMC architectures.

Focusing on the two hotelling architectures, it can be observed that the Bypass never performs worse than the Opaque, thus confirming the fact that the optical bypassing strategy always leads to the higher energy savings. The difference becomes slightly more relevant for medium values of L_D, because of the higher volume of traffic passing through intermediate nodes, which allows to bypass a higher number of lightpaths. For the highest values of L_D, most of BBUs are consolidated in the CO, therefore the two architectures perform similarly. However, in all cases the difference between Bypass and Opaque is much less relevant than the difference with respect to the No-Hotel (not greater than approximately 15% of relative difference). This raises the suspect that the tradeoff between optical bypass and active traffic aggregation plays only a

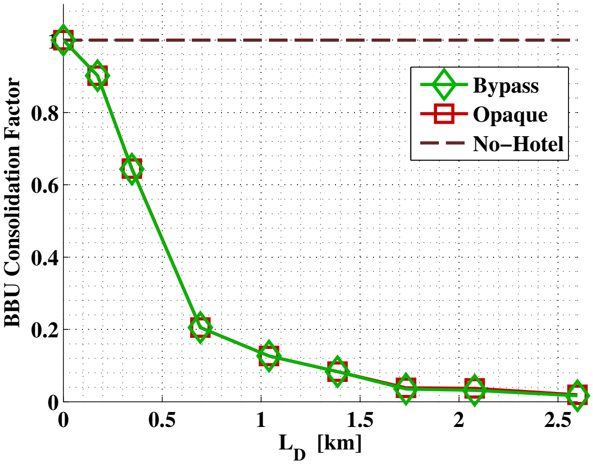

Figure 4.4: BBU Consolidation Factor vs. L_D, for Dense-Urban geotype.

marginal role in the energy efficiency, for the analyzed case study.

To further investigate this aspect, we analyze two additional performance indicators. The first one is the BBU Consolidation Factor (BCF), defined as the number of BBU sites (i.e., hotels and cell sites hosting their BBUs), normalized to the number of Mobile End Points. The second indicator is the Average Lightpath Utilization (ALU), which quantifies the efficiency of active traffic aggregation performed by electronic switches. It is defined as the ratio between the sum of all routed traffic demands and the total lightpath capacity (i.e., the number of established lightpaths multiplied by the wavelength capacity).

Fig. 4.4 and 4.5 show the defined quantities, as a function of L_D, for the Dense-Urban geotype and the considered network architectures. For the two hotelling architectures, the BBU Consolidation Factor exhibits large variations, ranging from 1 (lowest L_D) when all BBUs are placed at their cell sites; to $1/|N_M|$, which is the best value, achieved when all BBUs are consolidated into a single site, namely the CO. It can also be observed that the BBU Consolidation Factor is virtually unaffected by the specific hotelling architecture (Bypass or Opaque). Conversely, the Aver-

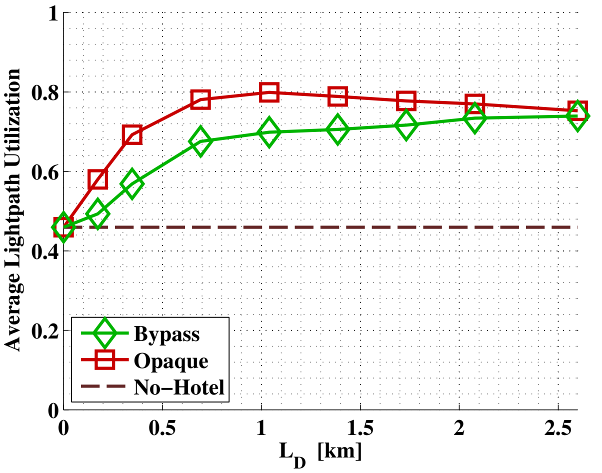

FIGURE 4.5: Average Lightpath Utilization vs. L_D, for Dense-Urban geotype.

age Lightpath Utilization shows different values. As expected, the addition of intermediate switches in the hotelling architectures enables a more efficient traffic aggregation with respect to the No-Hotel case. Also, the Opaque performs slightly better than the Bypass, but with a relative maximum difference of about 15%. Considering the power amount that can be saved for each less established lightpath ($2 \cdot K_S = 30\,\mathrm{W}$), it can be inferred that the energy savings due to the more efficient aggregation are not relevant in these cases.

We interpret such results as a strong confirmation that the BBU consolidation is the dominant principle which pushes toward valuable energy savings for the proposed architectures, in the analyzed case study. The possibility of performing optical bypass and active aggregation in intermediate nodes plays only a marginal role in the Aggregation Infrastructure Power. Since the same conclusion can be drawn from the analysis of the Urban and Rural geotypes, we do not report the corresponding figures for sake of space.

Finally, to study the effects due capacity limitation of hotels (i.e., when $B < |N_M|$), we consider a sensitivity analysis over

Table 4.3: Maximum AIP increase when the hotel capacity (B) is scaled to 2/3 and 1/3 of the original value.

Geotype	**2/3 capacity scaling**		**1/3 capacity scaling**	
	Bypass	*Opaque*	*Bypass*	*Opaque*
Dense-Urban	3.3%	6.9%	11.1%	26.1%
Urban	11.6%	16.5%	25.4%	40.0%
Rural	4.3%	8.7%	16.7%	28.5%

different values of B. All simulations have been re-run by scaling them down to approximately 2/3 and 1/3 of their original values, for each geotype. The scaled values are, respectively, 20, 10 for Dense-Urban; 13, 7 for Urban; 7, 4 for Rural. The main effect that can be expected from scaling down the hotel capacities is that the achievable BBU consolidation is also limited by the saturation of the capacity of some hotels. The reduced BBU consolidation causes a slight increase of the AIP with respect to the original capacity, which becomes more and more relevant when BBUs are moved towards the CO (higher L_D). To analyze such effect, in Table 4.3 we report the percentage of the maximum increase of the AIP, with respect to the original capacity value, for each hotelling architecture and geotype. As expected, for smaller hotel capacity there is a larger increase of power, because of the lower BBU consolidation. Also, the difference is more evident for the Opaque architecture than for the Bypass one. Anyway, in almost all cases, such increases can be considered as not critical for the analyzed case study, because they could be justified in cases when it is be more cost convenient to have some hotels with reduced capacity (e.g., because of lower lease fees).

4.6 Conclusion

In this chapter, we have presented three different architectures of FMC WDM aggregation networks: Bypass, Opaque and No-Hotel. We have proposed an energy efficient BBU placement optimization problem, whose objective is the minimization of a specific performance indicator, the Aggregation Infrastructure Power (AIP). A fundamental role in the optimization is played by the maximum route length of CPRI flows between Mobile End Points (i.e., cell sites) and the respective hotel nodes.

The case study analysis has been carried on randomly generated multi-stage tree topologies, representative of three different geotypes: Dense-Urban, Urban and Rural. The simulation results show that a relevant reduction of the AIP is achieved by both hotelling architectures (Bypass and Opaque), with respect to the No-Hotel. Typical values are around 60–65% of relative saving for Dense-Urban and Urban, and 40% for the Rural geotype.

Moreover, it is confirmed that BBU consolidation is the dominant principle which pushes toward valuable energy savings for the proposed architectures, in the analyzed case study, while optical bypass and active aggregation in intermediate nodes play only a marginal role.

This research topic opens the doors to many possible further lines of investigation. One of them is the extension to a dynamic traffic scenario, in which the optimization model also embeds the statistical information about daily traffic variation patterns. Another promising investigation in such scenario is about integrating per-node energy efficiency strategies (e.g., sleep modes) in order to dynamically adapt the network configuration to the varying traffic conditions. A further step is about the study of long-reach networks featuring more complex topologies, which also require a deeper analysis of the wavelength routing capabilities of intermediate optical devices.

A Heuristic Algorithm for the BBU Placement Problem 5

In this chapter, the BBU placement problem introduced in the previous chapter for the optimization of energy consumption in a FMC WDM aggregation network is evolved towards two directions. First, a more generic optimization metric is considered, that jointly takes into account the dominating contributions of cost and energy consumption in such network. Second, a larger-scale aggregation network is considered, that comprises several stages of intermediate COs up to a single high-level Point of Presence (PoP) acting as edge node. As a consequence of the increased number of nodes, an ILP formulation would be too complex to be solved for realistic instances, therefore a greedy-based heuristic algorithm is proposed.

5.1 Introduction

As seen in the previous sections, BBU hotelling in centralized radio access networks brings several benefits in terms of costs, energy consumption and network performance improvements. However, it also imposes strict requirements to the RAN infrastructure, particularly regarding the transport of fronthaul traffic between

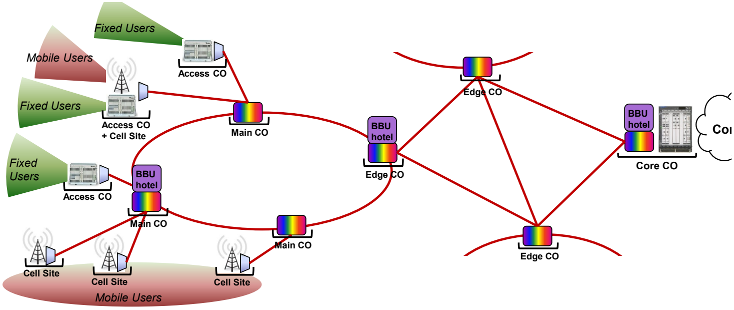

FIGURE 5.1: Illustrative example of a large-scale FMC WDM aggregation network, based on BBU hotelling.

BBUs and RRHs. Therefore, to fully enable BBU hotelling potentialities, not only BBU implementation, but also the underlying RAN must be properly designed and optimized. Many different optimization problems originate, covering various technological and architectural issues (e.g., [30, 44, 45]). In this section, we present a network optimization problem, featuring the placement of BBUs over a fixed/mobile converged WDM network acting at the same time as a backhauling/fronthauling RAN for mobile traffic and metro/aggregation for fixed traffic.

5.2 General Network Architecture

The considered large-scale aggregation network offers transport of traffic between a single Point of Presence (PoP) (identified as a Core CO), which is the edge node towards the core network, and two types of clients spread across a metropolitan or regional area: Central Offices (CO), which are the Fixed clients collecting the backhaul of various kinds of last-mile fixed access technologies (e.g., xDSL, TDM-PON), and Cell Sites (CS), which are the Mobile clients collecting backhaul/fronthaul traffic of mobile users. Some CO and CS may be co-located, and in this case they are Fixed/Mobile clients.

An illustrative example of such network is shown in Fig. 5.1.

The physical infrastructure consists in mono-fiber optical links connecting the PoP with network clients, through some intermediate nodes. As physical topology, the most straightforward is a multistage tree, but more complex topologies are possible as well, for instance, including rings and meshed connections. The network is based on WDM, i.e., fiber spectrum is divided in a number of independent optical signals at different wavelengths. The capacity of each wavelength is shared by multiple traffic flows, by means of electrical-domain multiplexing schemes (typically based on TDM). For this reason, at the PoP and at network clients, where optical signals are terminated via “colored” transceivers, there are also electronic switches that perform mu/demultiplexing of different flows in/from the wavelength signals.

Besides increasing the system capacity, WDM allows to route some traffic flows directly in the optical domain, i.e., without performing optical/electro/optical signal conversion. To do this, intermediate nodes are equipped with optical wavelength routers, i.e., devices that are capable of cross-connecting wavelengths at certain inputs towards different outputs, without performing wavelength conversion (e.g., AWG, OADM). Installing optical wavelength routers instead of electronic switches in intermediate nodes allows to greatly reduce energy consumption and operational costs for such nodes.

5.2.1 BBU Placement

The proposed WDM network serves not only as converged network for fixed and mobile backhaul, but also as infrastructure for a centralized radio access network, by means of BBU hotelling. This means that every network node becomes eligible for hosting one or more BBUs, so the BBU associated with a RRH can be placed in any node along the path from the CS to the PoP. Considering a generic BBU, there are three cases for its placement: at the CS, at the PoP, or at one of intermediate nodes. Placing the BBU at its CS is equivalent to the distributed BS architecture, i.e., an intra-CS optical connection transports the D-RoF signals between the BBU, located at the CS cabinet, and the RRHs, directly attached to antennas on the tower. In this case, the CS does not require fronthaul transport from the network. If the BBU is placed at the PoP, the network transports fronthaul up to the CS cabinet, where the switch extracts the D-RoF signals to be sent towards the RRHs. Therefore, the CS does not require backhaul transport from the network.

If the BBU is placed at an intermediate node, the same network transports both backhaul, towards the PoP, and fronthaul, towards the CS. The node becomes hotel and it can host collateral equipment which requires higher energy and operation/management costs than non-hotel intermediate nodes. In order to save switch processing resources (thus, energy consumption), some of the transit traffic can be directly routed by wavelength routers, without passing through the electronic switch. This technique is known as "optical bypass".

5.2.2 Traffic Routing

As a result of the BBU placement, in general three types of traffic coexist in the aggregation network: Fixed backhaul between COs and the PoP, Mobile backhaul between CSs or hotels and the PoP, and CPRI (fronthaul) between CSs and hotels. Differently from fixed and mobile backhaul, which is packet-based, asynchronous and latency-tolerant, CPRI consists in constant-bitrate flows, which require accurate synchronization and a constrained end-to-end latency. In order to effectively route such different traffic flows over the same network, a two-layer transport hierarchy is adopted.

At the lower layer the basic element is the "lightpath" which is a wavelength-routed circuit established between two nodes equipped with optical line terminations, which only traverses optical wavelength routers along its path. Every lightpath is uniquely identified by its wavelength and physical path, and provides a fixed capacity equal to the wavelength signal bitrate. All lightpaths whose propagation delay (i.e., total length) is not bigger than a threshold value are denoted as delay-limited lightpaths, or "D-lightpaths". Each CPRI flow is end-to-end transported over a single D-lightpath. Preventing the flow to be split among or traverse multiple lightpaths is necessary to meet fronthaul synchronization and latency requirements. For this purpose, the delay threshold of D-lightpaths is calculated as the fronthaul maximum latency minus the processing delays introduced by electronic equipment at lightpaths terminal nodes. Of course, more than one CPRI flow can be multiplexed into the same D-lightpath.

In our assumption, several lightpaths, including also D-lightpaths, can be established from a source to a destination node, in general routed over different paths and wavelengths. They constitute a single "virtual link" between these nodes, with capacity equal to the sum of capacities of all component lightpaths. Many backhaul flows can be transported in a virtual link, via inverse multiplexing techniques, i.e., splitting them in any proportion across multiple component lightpaths. Making backhaul flows splittable at lower layer enables more flexible routing, but different lightpaths delays increase transport latency due to buffering and reordering. However, unlike fronthaul, this does not critically impact on backhaul requirements. Virtual link capacity can be actually shared among backhaul and fronthaul, because some component D-lightpaths are filled with CPRI flows. The spare capacity of D-lightpaths, plus the capacity of regular lightpaths, is used for backhaul.

5.3 The BPTR Optimization Problem

We propose a network optimization problem, denoted as BPTR (BBU Placement and Traffic Routing), which addresses two main questions: where to place the BBU for each CS, and how to route traffic flows with the objective of the minimizing a suitable global "cost" function. The right cost function would result from a proper mix of CapEx, OpEx and energy consumption, but it is difficult to characterize it in a precise manner, as it largely depends on specific implementation details (e.g., used technologies) and other features which are difficult to generalize (e.g., economic ecosystem surrounding the network operator). In the following numerical evaluation, we choose a generic cost characterization, which is based on the following observations.

- Placing BBUs in the PoP has zero cost, because its premises are already optimized for hosting network systems, so its cost is virtually unaffected by additional hotelling equipment.

- Placing BBUs in any intermediate node has a fixed cost, say A, independent of how many BBUs are hosted, because the fixed costs due to the hotel installation should be dominant.

- Placing a BBU in its CS has cost equal to A, too. This comes from the assumption that a CS cabinet is similar to other intermediate sites. In fact if it is equipped with the BBU, it requires approximately the same CapEx/OpEx (that dominate over energy). If the BBU is not installed, the cabinet hosts only the terminal electronic switch, with a negligible cost.

- Establishing a lightpath has another cost, say B. It includes the cost for activating the pair of WDM interfaces but it can be dominated by other costs related to "resource utilization" (for instance if dark fiber capacity is bought from another operator, in units of used lightpaths). For this reason, giving B an absolute value is not trivial, but we can safely assume that it is much smaller than A (say, $B = A/100$).

Therefore, the BPTR problem can be described as follows:
given: the network topology (nodes connectivity and links lengths), the set of traffic demands, links capacity (number of wavelengths), traffic capacity (bitrate) of each wavelength, maximum number of BBUs that can be hosted in each intermediate node (hotel capacity), and the maximum fronthaul length;
decide: the placement of each BBU, and the routing of traffic requests, that includes lightpaths establishment (path and wavelength assignment) and mapping of flows into lightpaths;
to minimize: the global cost, which, following from the previous observations, can be written as:

$$Z = A\,(n_{HI} + n_{BC}) \; + \; B\,n_{EL} \tag{5.1}$$

where n_{HI} is the number of hotels placed at intermediate nodes, n_{BC} is the number of CSs equipped with their BBU (i.e., featuring the regular BS architecture), and n_{EL} is the number of established lightpaths.

It is important noting that, since we do not explicitly assign practical cost values to parameters A and B, the value of the global cost Z has not a practical meaning too. Instead, its definition embeds two distinct optimization functions, which feature two different minimization priorities. The primary one is the number of "BBU sites" $(n_{HI} + n_{BC})$, while the secondary one is the number of established lightpaths n_{EL}. This means that a solution is better than another one either is it features a smaller number of BBU sites (regardless the number of lightpaths), or if they both have the same number of BBU sites, but the number of lightpaths is smaller. Therefore, only the values assumed separately by these two sub-function are actually meaningful, in order to analyze or compare different solutions.

The defined problem can be interpreted as a non-trivial combination of two well-known network optimization problems, namely the Grooming, Routing and Wavelength Assignment (GRWA) typical of core/backbone WDM network [40], and the Facility Location Problem (FLP) [41]. Both GRWA and FLP are known to be NP-hard, therefore our defined problem is NP-hard too, because it contains them as subproblems. In practice, it means that all known algorithms can find the optimal solution with complexity that grows exponentially as a function of the size of the network instance. This prevents us to obtain exact solutions for realistic network sizes, so we must resort on heuristic algorithms, which exhibit reasonable complexity, but do not guarantee to find the optimal solution. In the following, we propose a heuristic formulation of the problem.

5.4 A Heuristic Greedy Algorithm for BPTR

Considering the fact that the cost of housing a BBU at CS dominates compared to the cost of establishing lightpaths, the focus of proposed greedy heuristic is on FLP. The strategy to minimize the global cost consists of the following principles (in decreasing order of priority): *(a)* to open as few intermediate hotels as possible; *(b)* to locate as many as possible BBUs at intermediate hotels; and *(c)* to activate as few lightpaths as possible. Fronthaul has more priority in routing, with respect to backhaul, because it cannot be split among different wavelengths and it is delay-sensitive, so it exclusively requires D-lightpaths.

5.4.1 Notation and Input Data

- N is the set of network nodes, partitioned into: the PoP, the subset N_I of intermediate nodes, and the subset of clients, which can be Fixed ($\in N_F$), i.e. COs; Mobile ($\in N_M$), i.e., CSs; or both ($\in N_F \cap N_F$), i.e., CSs collocated with COs.

- E is the set of physical links, a length is associated to each physical link.
- V is the set of virtual links, indexed by (i, j), with i,j $\in N, i \neq j$.
- R is the set of connection requests, partitioned into: Fixed backhaul (R_F), Mobile backhaul (R_M), and Fronthaul, e.g., CPRI (R_C). For each request, $r \in R$, $t_{(i)}^{(r)}$ is the traffic demand of node i.
- Λ is the set of wavelengths.
- C is the traffic capacity of each wavelength.
- L_D is the maximum fronthaul length, i.e. the maximum length of D-ligthpaths.

5.4.2 Heuristic Subroutines

- *Initialization*: For each node pair (i, j), $i \neq j$, a number $\Omega_{(i,j)}$ of D-paths (i.e., paths from i to j with physical length not exceeding L_D) is calculated). For each mobile client $i \in N_M$, a first-fit decreasing bin packing heuristic [46] is solved to calculate $NUM_{(FH)}^{((i))}$ = the number of D-lightpaths that the client needs for fronthaul if its BBU is placed at a hotel. For each $i \in N_M$ the Residual Fixed Traffic is defined as: $RFT^{(i)} = \sum_{r \in R_F} t_{(i)}^{(r)} - \lfloor t_{(i)}^{(r)} \rfloor$, i.e., the amount of fixed traffic modulo the single wavelength capacity.
- *BestFeasible*: Given a candidate hotel node j, a mobile client i is said to be reachable from j if there are enough D-lightpaths for accommodating the total fronthaul traffic requested by i. The set of reachable mobile clients, denoted as $P_C^{(j)}$, is computed for each candidate hotel j. The candidate hotel with the highest cardinality $|P_C^{(j)}|$ is selected and feasibility of the routing of the traffic associated with reachable mobile clients is checked. If the routing is not feasible, then the next highest-cardinality candidate hotel is checked, until a feasible routing is found. In this case, the relative candidate hotel is opened.

- *Stop*: Break if no other candidate hotel can host at least two BBUs.
- *Groom&Route*: D-lightpaths for fronthaul are established, choosing, among free D-lightpaths, those with the smallest hop-distance and randomly assigning the wavelegths. If $RFT^{(i)}$ is less than the spare capacity of established D-lightpaths, then it is groomed in those D-lightpaths up to the hotel node. All groomed RFTs and mobile traffic demands associated with $P_C^{(k)}$ are further groomed together and routed from the PoP to the hotel. For each mobile client i in $P_C^{(k)}$, the remaining fixed traffic, i.e., the RFTs which have not been groomed plus all the traffic which fits into an integer number of wavelengths, is routed from PoP to i. The routing phase is performed in such way that any active node is bypassed. Physical layer routing is implemented by shortest path.
- *Update*: Remove k and $P_C^{(k)}$from N.
- *RemainedBBUs&Routing*: BBUs of remaining mobile clients are placed at CSs. Traffic of these mobile nodes and of the nodes with pure fixed traffic $(N_F - N_M)$ are routed from PoP to client nodes.

5.5 Heuristic Scheme

```
begin
  Initialization {
    extract all D-paths
    calculate $NUM_{FH}^{(i)} \vee i \in N_M$
    calculate $RFT^{(i)} \vee i \in N_M$
  }end Initialization
  do{
    $(k, P_c^{(k)}) = \text{BestFeasible}(N, \text{Free D-lightpaths}, \text{Free lightpaths}, NUM_{FH}^{(\forall i)})$

    Stop $(|P_c^{(k)}|)$
    Routing of
        $P_c^{(k)} = \text{Groom\&Route}(P_c^{(k)}, \text{Free lightpaths}, \text{Free D-lightpaths})$
```

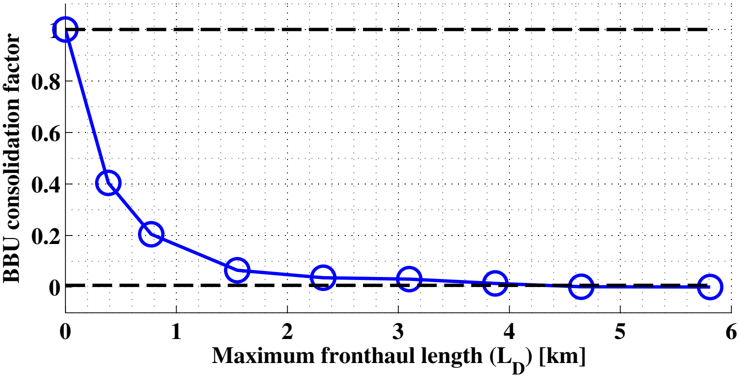

Figure 5.2: Dense-Urban geotype: BBU consolidation factor vs. maximum fronthaul length (*dashed top line:* fully distributed case; *dashed bottom line:* fully centralized case)

```
        $(N) = Update(N, P_c^{(k)}, k)$
    }end do
    $(RemainedRouting, BBUsatCSs) = \text{RemainedBBUs\&FinalRouting}(N)$

end begin
```

5.6 A Case Study for the BPTR

As a case study for numerical evaluation, we performed multiple optimization runs over synthetic network instances based on multistage tree topology, which are randomly generated according to the model summarized in the following (a more detailed description can be found in Section 4.4). Three typical geotypes are considered: “Dense Urban”, “Urban”and “Rural”, featuring different values of the size of the coverage area and of the spatial densities of network clients (CSs and COs). Clients’ coordinates are randomly scattered over the coverage area, according to a uniform distribution. A multistage tree graph is constructed by performing a hierarchical clustering, via the k-means algorithm with city-block (also known as “Manhattan”) distance metric. The number of stages is fixed to 3 and the numbers of clusters for each stage are such that the corresponding split ratio is approximately constant. After this process, the physical topology and all links’

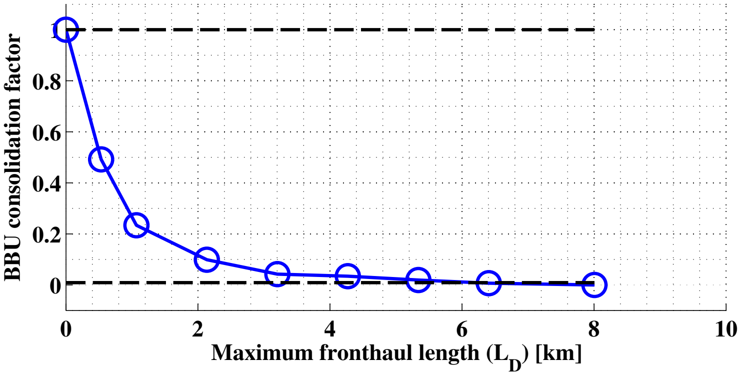

FIGURE 5.3: Urban geotype: BBU consolidation factor vs. maximum fronthaul length (*dashed top line:* fully distributed case; *dashed bottom line:* fully centralized case)

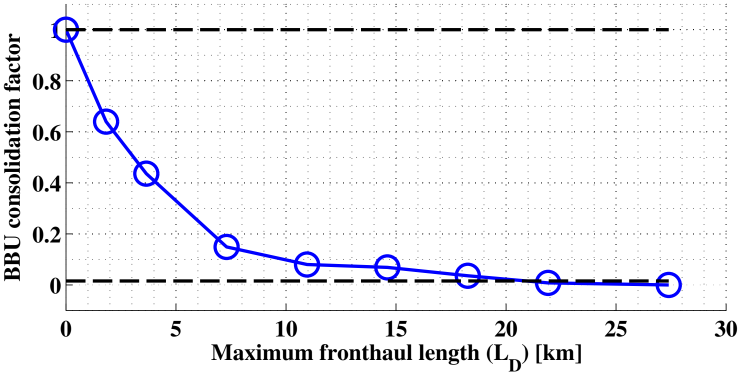

FIGURE 5.4: Rural geotype: BBU consolidation factor vs. maximum fronthaul length (*dashed top line:* fully distributed case; *dashed bottom line:* fully centralized case)

lengths of the instance are obtained.

The traffic associated with each network client is generated according to the following assumptions. Traffic requests are rounded to integer multiples of the gross bitrate of the SONET/SDH OC-3 signal (155.52 Mb/s), hence, they can be quantified in terms of “OC-3” units. Each CS provides LTE coverage of 3 sectors, with 20 MHz carrier bandwidth and 2x2 MIMO configuration. Using the CPRI interface, the required fronthaul bitrate can be computed as 6.29 Gb/s (41 OC-3s), not including the CPRI line code overhead, which is not needed if fronthaul is mapped on an underlying transport network. Required mobile backhaul bitrate is uniformly distributed in the range 300-750 Mb/s (2-5 OC-3s), which is typical for such macro CS configuration. Each CO collects fixed access traffic from approximately 800 households, each requiring a typical bitrate uniformly distributed from 12.5 to 25 Mb/s. Therefore, each CO requires fixed backhaul bitrate uniformly distributed in the range 10-20 Gb/s (64-128 OC-3s). In the following results, the number of wavelengths per link is 80 and each wavelength has a capacity equal to 64 OC-3s (around 10 Gb/s). Each plotted point results from the average of 10 randomly generated instances, according to the previously described model and the BPTR problem input parameters.

Figures 5.2, 5.3 and 5.4 show a metric that quantifies how much BBU are consolidated into shared hotels, as opposed to being placed into CSs. It is denoted as “BBU consolidation factor”, and it is defined as the primary sub-function, namely the number of “BBU sites” (i.e., the sum of the number of intermediate hotels and the number of CSs equipped with their BBU $(n_{HI} + n_{BC})$), normalized to the number of CSs (n_{CS}). The metric is shown as a function of the maximum fronthaul length (L_D), in km. As a comparison benchmark, there is also shown the BBU consolidation factor in the two extreme cases: fully distributed (maximum) and fully centralized (minimum).

For all geotypes, it is evident that the metric rapidly decays when the maximum fronthaul length (L_D) increases, because higher values of L_D allow BBUs to be placed at highest nodes of the multistage tree hierarchy (i.e., closer to the PoP), thus enabling a bigger number of CSs to be controlled by a single hotel site. The case $L_D = 0$ reduces to the conventional no-hotelling network architecture, in which all BBUs are located in respective CSs and no fronthaul is transported (zero BBU consolidation), i.e., the metric assumes the highest value, equal to the number of CSs to which the network provides coverage. It can be observed that the metric reduces to 20% of the maximum value after L_D reaches a few kms (around 1km for Dense-Urban and Urban scenarios,

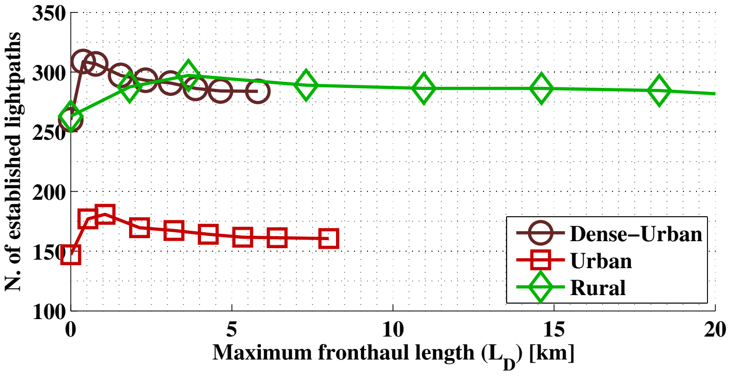

FIGURE 5.5: All geotypes: Number of established lightpaths vs. maximum fronthaul length.

around 6km for the Rural scenario). It also gets the minimum value, which is 0 (maximum BBU consolidation), corresponding with all BBUs placed in the PoP premises, by less of 10 km for Dense-Urban and Urban and around 25 km for the Rural scenario. Such values are considerably below the classical 40 km limit, hence they indicate not only that BBU hotelling is a feasible solution for the proposed network architecture, but also that relevant amount of latency budget is available for adding more complexity to the transport architecture. This strongly justifies, as a possible next step, the investigation of long-reach architectures, that feature regional coverage areas (characterized by a mixture of different geotypes) and distances from the PoP to clients up to hundreds of km. In this case, the longer optical distances require more advanced signal processing in order to overcome propagation impairments, with an impact on the available fronthaul latency budget, therefore the achievable BBU consolidation.

Fig. 5.5 shows secondary sub-function contribution, namely the number of established lightpaths (n_{EL}). For every geotype, it can be observed that such metric stays almost constant for all values of L_D. The only remarkable difference is the case $L_D = 0$, in which, as intuition suggests, the architecture reduces to a pure no-hotelling network, hence the total amount of traffic significantly reduces as there is no fronthaul traffic to be transported.

5.7 Conclusion

In this chapter, we have proposed a generalization of the energy-efficient BBU placement problem, that employs a generic cost/energy optimization metric and can be applied to larger-scale FMC WDM aggregation networks. Because of the increased size of considered network instances, we have tackled such problem by means of a greedy-based heuristic algorithm. Finally, a numerical illustrative evaluation has been carried on over randomly generated multistage tree networks, for different geographic scenarios. Results agree with the same obtained for the energy-efficient problem, and clearly indicate that BBU consolidation into fewer hotel nodes is the key principle for reducing the overall cost/energy and justifies the evolution towards converged architectures, also for larger-scale aggregation networks.

Joint BBU and Electronic Switches Placement in a Multifiber WDM Aggregation Network

6

In this chapter, we introduce a novel optimization problem that takes into account more degrees of freedom of WDM aggregation network design. Differently from the previously defined problem, a multifiber infrastructure is considered and a new node architecture is adopted, that allows both the insertion of internal traffic (locally generated) and the transit of external traffic. The transit traffic can be either purely bypassed by an optical switch, or be aggregated/groomed by installing an additional electronic switch, independently from the presence of a hotel. In addition, we consider different transport methods for fronthaul, exhibiting different multiplexing and routing constraints. Therefore, the decision of installing electronic switches in each node and the multifiber network dimensioning are optimized, together with BBU placement and traffic routing. In this way, the model becomes more scalable to real-world larger networks and more interesting to investigate because of the non-trivial interaction among hotels/switches placement and routing/grooming of both backhaul and fronthaul traffic.

6.1 Introduction

The valuable benefits that a FMC WDM aggregation network can bring in terms of cost and/or energy consumption have been studied and demonstrated in the previous chapters. The BBU placement problem assumes a big relevance for these networks because, from one hand, the BBU consolidation principle pushes hotels to be placed far away from managed cell sites, but for the other hand the strict requirements of fronthaul traffic put a limit on the maximum distance between hotels and cell sites. Depending on the specific aggregation infrastructure supporting BBU hotelling, there are additional constraint, e.g., on traffic multiplexing and routing or physical topology, that make this problem interesting. In the previously introduced optimization model, we studied the interaction of BBU placement with lightpath routing in a WDM network and traffic grooming by both end and intermediate electronic switches.

In this chapter, we extend this model, in order to make it more suitable to real-world network scenarios. In particular, the changes introduced with respect with the previous BBU placement problem are detailed in the following.

- The WDM network architecture is multifiber, i.e., optionally more than one fiber can be used in each physical link. In this way, the network can be scaled to sizes that are more typical in aggregation, like extra-metropolitan or regional.

- The placement of electronic switches is decoupled from the placement of BBU hotels, meaning that switches are not necessarily combined with hotels. This relaxation would make possible to optimize the grooming of aggregation traffic in a more efficient way without directly impacting on BBU hotelling.

- There is no more strict distinction between end-nodes and intermediate nodes of the aggregation infrastructure. Therefore, each CO is assumed to be equipped with both access gateways, multiplexing different first-mile access technologies, and aggregation switches.

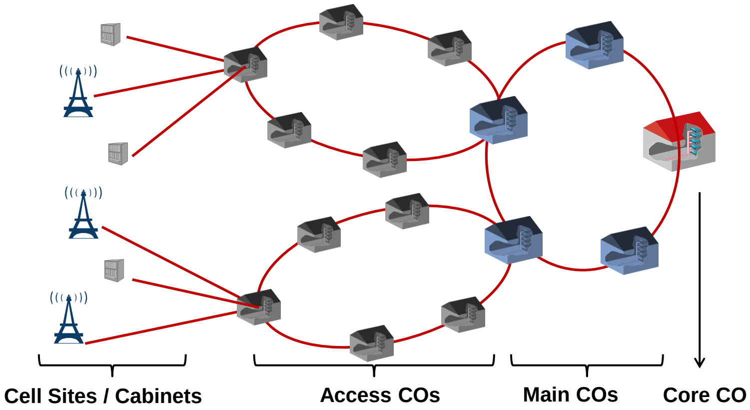

FIGURE 6.1: Example of a hierarchical aggregation network scenario

- There are two different option for fronthaul transport, namely: *OTN aggregation* and *Overlay*, which are described in the following sections.

6.2 Network Architecture

As for the previously defined architecture, the WDM aggregation network transports fixed and mobile traffic from end nodes and the edge node toward the core, denoted ad Point of Presence (PoP). End nodes can be either COs, collecting fixed access traffic from a corresponding fixed access service area, or Cell Sites (CSs), collecting mobile traffic of the served cell/sector areas, in form of backhaul or fronthaul, respectively if they host their BBUs or not, or colocated CO/CSs.

Any kind of hierarchical structure among COs can be present, that distinguish them in relation to the aggregation stage they belong to. An example is depicted in Fig. 6.1, in which at stage 0 there are Access Cabinets (Access sub-COs), at stage 1 Access COs, at stage 2 Main COs, and at stage 3 a single Core CO, wich acts as PoP. In a similar way, a stage can be defined for each CS, depending on the CO it is connected to or co-located with.

The network is multifiber, i.e., links between adjacent nodes

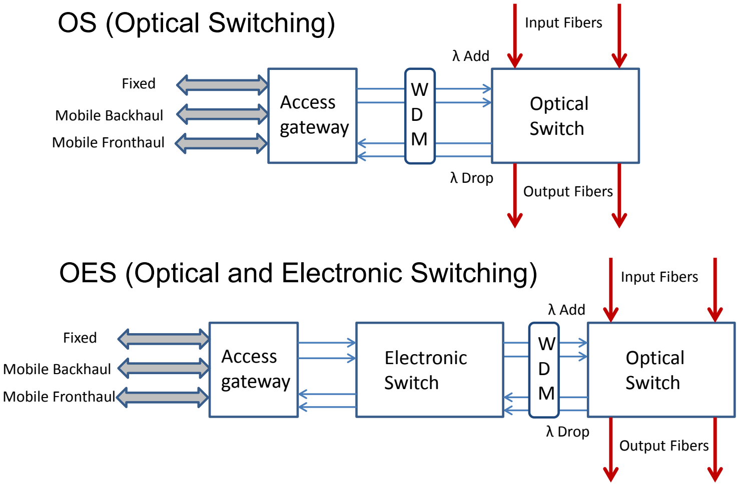

Figure 6.2: Architectures of OS (Optical Switching) and OES (Optical and Electronic Switching) nodes.

consist in a number of fiber pairs, where each fiber in each pair gives connectivity towards one of the two directions. The capacity of each fiber is divided in a number of WDM channels, or wavelengths. Over such network, not only BBU hotels but also electronic switches can be placed, independently from each other.

The traffic managed by each node can be classified between: *internal* traffic, i.e., fixed, mobile, or fronthaul originated from the service area of the same node; and *external* traffic, i.e., originated from any other nodes, thus exclusively transiting. The node switching architecture is the most relevant innovation with respect to the previous model, hence it is described more in detail. As depicted in Fig. 6.2, there are two kinds of nodes:

- *OS (Optical Switching)*: An access gateway collects and packs the internal traffic into one or more gray multiplex flows. Such signals are converted to colored wavelength (WDM) signals and sent towards an optical switch, that performs wavelength switching among add/drop ports and input/output fibers directly in the optical domain, i.e., without electronic/optical conversion.

- *OES (Optical and Electronic Switching)*: The gray multiplex flows generated by the access gateway enter an electronic switch, that is capable of extracting, permuting and repacking tributary subflows among different multiplex signals on different ports, i.e., performing traffic grooming. Then, the obtained "groomed" traffic flows are converted to colored wavelength signals and sent to the optical switching as in the OS case. It is worth noting that the access gateway and electronic switch can be implemented as a single physical device.

To make these figures clearer, the optional BBU hotel is not depicted, but its presence is implied by the presence of internal fronthaul traffic in substitution of the backhaul.

6.3 Traffic Multiplexing and Routing

There are three classes of traffic: Fixed, Mobile (backhaul) and Fronthaul (or CPRI). Each traffic request is associated with a specific node and a traffic class, and requires a given transport capacity in both downstream and upstream directions. The transport is done in a two-layer fashion: in the upper layer, each requests is a tributary for one or more multiplex flows at a fixed line rate, in the lower layer, such flows become wavelength signals that are transparently routed along paths, thus forming the so called wavelength-paths, or "lightpaths".

Requests can be multiplexed and/or demultiplexed among different lightpaths, but this operation can be performed by access gateways for internal traffic in OS/OES nodes, or electronic switches for external traffic exclusively in OES nodes. Each request is end-to-end routed over a sequence of lightpaths established between the origin node and optionally some intermediate OES nodes, up to the final destination. Specific constraints on traffic routing are applied, depending on the class.

Fixed and Mobile backhaul traffic is natively packed-based (e.g., IP), exhibiting some degree of tolerance on absolute delay, differential delay and out-of-order delivery. Therefore, we assume that relative requests can be arbitrary splitted among parallel lightpaths (i.e., with common starting and ending nodes), and they can traverse multiple OES nodes.

Fronthaul traffic is way more restrictive, because it is natively circuit-based, with hard synchronization requirements and a maximum transport round-trip time due to the fact that radio-interface physical layer procedures strictly limit the maximum delay between reception of UL frames and transmission of corresponding DL replies (more details are give in Section 2.5.2). We consider two kinds of transport strategies for fronthaul traffic: *"OTN (Aggregation)"* and *"Overlay"*.

In the *OTN* case, both fronthaul and fixed/mobile backhaul are transported over a common OTN layer. This makes possible to multiplex them together into the same wavelengths. Fronthaul flows can traverse intermediate OES nodes, which in this case are equipped with OTN switches or "wrappers", provided that the node-processing extra latency contribution is subtracted from the latency budget available for fiber propagation. Note that at least two switches are traversed, one for ingress, the other for egress of fronthaul. As current switching technologies (e.g., Ethernet, or OTN) feature processing delays that are too far from the requirements of fronthaul transport, we assume "low-latency" switches, on-purpose tailored for fronthaul applications, adding a delay of $t_{SW} = 20\,\mu\text{s}$. The possibility of performing fronthaul splitting among parallel lightpaths requires a more detailed analysis of the impact of extra delays and how to meet the strict synchronization requirements, which is out of the scope of this work. Hence, in this architecture we assume that fronthaul flows can not be splitted among parallel lightpaths.

In the *Overlay* case, fixed/mobile backhaul goes over a separate electronic layer (OTN or even Ethernet), while fronthaul is directly transported over dedicated wavelengths. Therefore, each fronthaul flows is assigned to a single lightpath end-to-end, and entirely occupies a whole wavelength, even if it requires a lower capacity, so no multiplexing with any other flow is possible. Since there is no electronic transport, the only fronthaul latency contribution is due to propagation.

6.4 Optimization Problem Definition

Over the introduced FMC WDM aggregation network, we define the new "Joint BBU and Electronic Switch Placement Problem"

as follows.
Given: the network topology (nodes connectivity and links lengths), the number of available fibers for each link, the number of wavelengths available for each fiber, the line rate (capacity) of each wavelength signal, the set of traffic requests, the Maximum Fronthaul Route Length (L_D) or equivalently its maximum delay, and the network cost contributions;
decide: the placement of each BBU, the placement of electronic switches (i.e., which nodes are OES), the Grooming, Routing and Wavelength Assignment (GRWA) of all traffic requests;
to minimize: the total network cost, defined as the sum of four terms: cost of Hotels, cost of Electronic Switches, cost of Fiber, cost of Lightpaths.

6.4.1 ILP Formulation

The optimization problem is mathematically modeled by means of an Integer Linear Programming (ILP) formulation. Differently from the previous problem modeling, here a "path" formulation is employed. A set of paths for each pair of nodes is pre-computed, where each path is a sequence of consecutive links from the start node to the end node. Explicitly enumerating the possible paths makes more straightforward to introduce more complex routing constraint, that leaves room for further improvement of the model (e.g., limitations caused by physical-layer impairments, typical of networks spanning larger geographical areas).

Paths are the building blocks for the hierarchical topological structure of the network. Each path, associated with a single wavelength, constitutes a lightpath, established between nodes equipped with access gateways and/or electronic switches. Being the network multifiber, in general more than one lightpaths over the same path and wavelength can be established. The maximum number depends on how many fibers are available on each link composing the path.

All the parallel lightpaths that can be established between the same node pairs constitute a higher-level routing element, denoted as "virtual link". Virtual links enable a powerful modeling of the traffic splittability, because each of them constitutes a single circuit whose capacity is equal to the sum of the capacities of the established lightpaths on that virtual link. In this way, fixed and mobile requests are routed over sequences of consecutive virtual links, without explicitly deciding which lightpaths are associated with each request. This allows to reduce the number of variables, thus the model complexity.

Fronthaul requests, which are not splittable, are routed over sequences of lightpaths, where in general each lightpath should be explicitly assigned in terms of path, wavelength and fiber. However, since different fronthaul flows can not be multiplexed into the same lightpath, there will always be at most one flow in each lightpath. Thanks to this restriction, we are able to avoid an explicit assignment of lightpaths, by associating each fronthaul flow to a path, thus partially reducing the number of variables.

For sake of simplicity, we assume that each connection request is symmetric, i.e., it requires the same traffic in the downstream and upstream direction. Therefore, we can formalize and solve the problem considering only one direction (in this case, upstream) and the results are straightforwardly replicated in the other one.

6.4.1.1 Input Sets

- N is the set of nodes, partitioned into: the PoP $\{o\}$, and the subset N_U of users (COs and CSs). Each user is respectively Fixed ($\in N_F$), or Mobile ($\in N_M$), if it has at least a corresponding request.
- E is the set of physical links, indexed by e.
- P is the set of (precomputed) paths, indexed by p, and such that the following subsets can be identified:
  - P_{i*}: paths starting from node i;
  - P_{*j}: paths ending to node j;
  - $P_{ij} = P_{i*} \cap P_{*j}$: paths from node i to node j;
  - $P_i = P_{i*} \cup P_{*i}$: paths starting from or ending to node i;
  - P_v: paths belonging to virtual link v;
  - P^e: paths passing through link e;

- V is the set of virtual links, indexed by (i, j), i.e., all pairs of nodes $i, j \in N$, with $i \neq j$, such that the sets V_{i*}, V_{*j}, V_{ij} and V_i are defined similarly as for P.
- Λ is the set of wavelengths, indexed by λ.
- R is the set of connection requests, indexed by r, and such that the following subsets can be identified:
  - R_n^F: Fixed requests of node n;
  - R_n^M: Mobile requests of node n;
  - R_n^C: Fronthaul (CPRI) requests of node n;
  - $R_{F \cup M}$: all Fixed and Mobile requests;
  - R_C: all Fronthaul requests;

6.4.1.2 Input Parameters

- l_p is the length of path p, or equivalently expressed as (propagation) delay.
- l_{EL} is the processing delay introduced by each electronic switch, or equivalently expressed as length.
- L_D is the maximum fronthaul delay, or equivalently expressed as length.
- c_r is the capacity of request r (note that since fronthaul is not splittable among different lightpaths, there must hold: $c_r \leq C$, $\forall r \in R_C$, while no limitation applies for other requests).
- C is the capacity of each wavelength signal, (i.e., of each lightpath).
- K is the maximum number of fibers available for each link.

6.4.2 Decision variables

- $y_v^r = 1$, if Fixed or Mobile request $r \in R_{F \cup M}$ is routed over virtual link v (binary).

- $\overline{y}_p^r = 1$, if Fronthaul request $r \in R_C$ is routed over path p (binary).
- $u_{p\lambda}$ = number of established lightpaths on path p and wavelength λ (integer).
- f_e = number of used fibers on link e (integer).
- $x_i^n = 1$, if the BBU of node n is placed at node i (binary).
- $w_i = 1$, if a hotel is placed in node i, i.e., with at least one hosted BBU (binary).
- $z_i = 1$, if an electronic switch is places at node i, i.e., it is OES (binary).

6.4.3 Objective Function

The objective is the minimization of the network cost, defined as the sum of four terms: cost of hotels, cost of electronic switches, cost of fiber and cost of lightpaths.

$$\min \left\{ A_{HOT} + A_{EL} + A_{FIB} + A_{LIG} \right\} = \alpha_{HOT} \sum_{i \in N} w_i + \alpha_{EL} \sum_{i \in N_U} z_i + \sum_{e \in E} \alpha_{FIB,e} f_e + \alpha_{LIG} \sum_{p \in P, \lambda \in \Lambda} u_{p\lambda} \quad (6.1)$$

where the corresponding α parameters are the unitary costs of: a hotel, an electronic switch, a used fiber in link e and an established lightpath.

6.4.4 Constraints

Routing of Fixed requests

$$\sum_{v \in V_{*i}} y_v^r - \sum_{v \in V_{i*}} y_v^r = \begin{cases} -1 & \text{if } i = n \\ 1 & \text{if } i = o \\ 0 & \text{otherwise} \end{cases} \quad \forall n \in N_F,\ i \in N,\ r \in R_n^F \quad (6.2)$$

Routing of Mobile requests

$$\sum_{v \in V_{*i}} y_v^r - \sum_{v \in V_{i*}} y_v^r = \begin{cases} -x_i^n & \text{if } i \neq o \\ 1 - x_i^n & \text{if } i = o \end{cases}$$
$$\forall n \in N_M,\ i \in N,\ r \in R_n^M \quad (6.3)$$

Routing of Fronthaul (CPRI)

$$\sum_{p \in P_{*i}} \overline{\mathrm{y}}_p^r - \sum_{p \in P_{i*}} \overline{\mathrm{y}}_p^r = \begin{cases} x_i^n & \text{if } i \neq n \\ x_i^n - 1 & \text{if } i = n \end{cases}$$
$$\forall n \in N_M,\ i \in N,\ r \in R_n^C \quad (6.4)$$

Capacity of virtual links

$$\sum_{r \in R_{F \cup M}} c_r\ y_v^r + \sum_{r \in R_C, p \in P_v} c_r\ \overline{\mathrm{y}}_p^r \leq \sum_{p \in P_v, \lambda \in \Lambda} C\ u_{p\lambda}$$
$$\forall v \in V \quad (6.5)$$

One fronthaul request per lightpath

$$\sum_{\lambda \in \Lambda} u_{p\lambda} \geq \sum_{r \in R_C} \overline{\mathrm{y}}_p^r$$
$$\forall p \in P \quad (6.6)$$

Maximum Fronthaul Delay

$$\sum_{p \in P} (l_p + l_{EL})\ \overline{\mathrm{y}}_p^r + l_{EL}(1 - x_n^n) \leq L_D$$
$$\forall n \in N_M,\ r \in R_n^C \quad (6.7)$$

Identification of Electronic Switch nodes

$$y_v^r \leq z_i;\ \forall n \in N_F,\ i \in N - \{o, n\},\ v \in V_i,\ r \in R_n^F \quad (6.8)$$

$$y_v^r \leq z_i + x_i^n;\ \forall n \in N_M,\ i \in N - \{o\},\ v \in V_i,\ r \in R_n^M \quad (6.9)$$

$$\overline{\mathrm{y}}_p^r \leq z_i + x_i^n;\ \forall n \in N_M,\ i \in N - \{n\},\ p \in P_i,\ r \in R_n^C \quad (6.10)$$

One BBU for each mobile node

$$\sum_{i \in N} x_i^n = 1 \qquad \forall n \in N_M \tag{6.11}$$

Identification of hotel nodes

$$w_i \geq x_i^n \qquad \forall n \in N_M,\, i \in N \tag{6.12}$$

Maximum number of fibers for each link

$$\sum_{p \in P^e} u_{p\lambda} \leq f_e \leq K \qquad \forall e \in E,\, \lambda \in \Lambda \tag{6.13}$$

Eq. 6.2, 6.3 and 6.4 enable the routing of fixed and mobile requests over virtual links of and fronthaul over lightpaths, taking into account that the hotels are not only an outcome of the optimization process, but also source/destination nodes for mobile flows. Eq. 6.5 ensures that the sum of the capacities of requests routed over a virtual link does not exceed the aggregated capacity of the virtual link. Eq. 6.6 avoids multiplexing among fronthaul requests by forcing each of them to be routed in a separate lightpath from the others, so that the total number of fronthaul requests in each path can not be greater than the number of established lightpaths on such path. Eq. 6.7 enforces the maximum fronthaul delay constraint, by taking into account both the fiber propagation delay (l_p) and the processing delay contributions of traversed electronic switches (l_{EL}), whose number is equal to the number of consecutive paths, plus one. Note that the factor $(1 - x_n^n)$ is necessary in order to disable this constraint in case the BBU is located at its cell site and fronthaul is not transported. Eq. 6.8, 6.9 and 6.10 are needed to identify which nodes are OES, i.e., equipped with electronic switches. For each node, they force to zero all the external traffic routed over virtual links originating from or terminating to such node, if it is OS instead of OES. For internal traffic (that passes through the access gateway) and in any case in the node is OES, the constraint is not active. Eq. 6.11 enforces that each mobile node is associated with exactly one BBU, thus one hotel. Eq. 6.12 is needed to identify hotels as nodes which host at least one BBU. Finally, Eq. 6.13 is a bottleneck constraint that linearizes the computation of the number of used fibers in each link, i.e., $f_e = \max_{\lambda \in \Lambda} \sum_{p \in P^e} u_{p\lambda}$. It also limits it to maximum number of fibers available for each link.

6.5 Case Study 1: Topology and Traffic

To obtain numerical results, we choose to investigate more general network topologies. This represents an evolution with respect to multistage tree topologies that have been studied for the previous optimization model. The reason is that, in view of a more real-world model assumptions (for instance, dealing with a multifiber network and larger network sizes), path- and node-redundant topologies are a more realistic feature, because the network needs a certain degree of resiliency to links and node failures. However, the complexity of the ILP model makes almost unfeasible to obtain results for practical sizes with available time and computational power. Therefore, in this work we try to stay within a reasonable tradeoff between feasibility of the problem solving and meaningful network instances.

In this first case study, the considered instances are based on a 2-stage ring and spur topology with three stages, in which a big ring at the higher stage comprising the PoP is connected to a number of smaller rings at the intermediate stage, which are connected to trees at the lower stage.

Topological instances are randomly generated by means of an algorithm adapted by that of 4.4. At first, a number n_N of nodes is uniformly scattered over a square coverage area, according to predefined geographical density parameters specific for COs and CSs. Then, given as input n_N, the number of trees n_T and the number of intermediate rings n_R, an algorithm computes the node connectivity and link lengths in order to sub-optimally minimize the total length. In detail, a hierarchical clustering based on the k-means algorithm with Euclidean distance metric is used to extract nodes belonging to each tree and ring. For each tree, the root is chosen as the closest node to the cluster centroid, while, for each ring cluster, links are obtained by a heuristic minimum-length Hamiltonian cycle computation routine. Finally, the node closest to the center of the area is chosen as the PoP. An example of 2-stage ring and spur topology instance obtained by such procedure

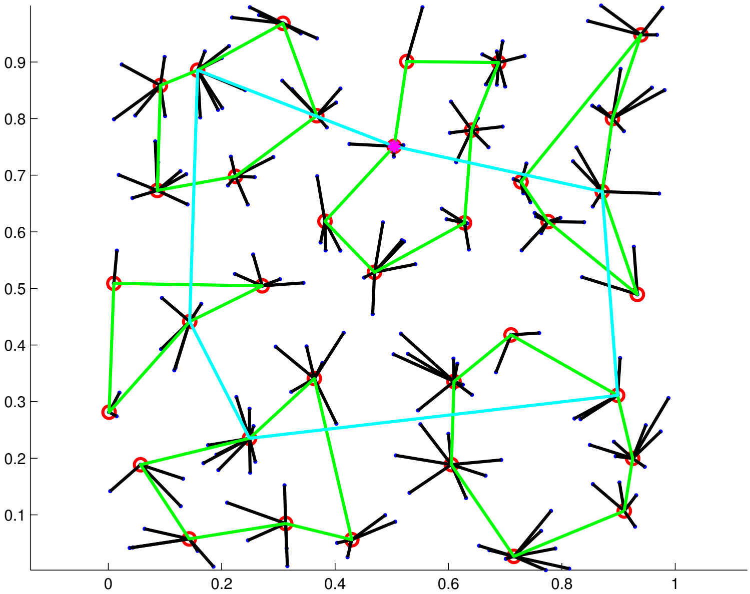

Figure 6.3: Example of 2-stage ring and spur topology instance obtained by the hierarchical clustering.

is given in Fig. 6.3.

To compute link lengths, a scaled Euclidean distance is adopted, that takes into account the overhead in link length due to typical urban planning constraints. Since for such 2-stage ring and spur topology there are at most 8 different (loopless) paths between any node pair, we precompute all the paths by means of a *k-shortest paths* algorithm, with $k = 8$. Finally, traffic requests are randomly generated for each node, by deciding if it is a fixed (COs) or mobile (CSs) node, and according to specific capacity distribution for each traffic class.

For the case study here presented, we consider an average urban scenario featuring a square area of side 3.5 km, with $n_N = 18$ nodes, consisting of 4 COs and 16 CSs, in accordance to density parameters indicated in Table 4.2 ([42, 43]). Several random topology instances are obtained by setting $n_T = 6$ trees and $n_R = 2$ small rings. several traffic request matrixes are generated

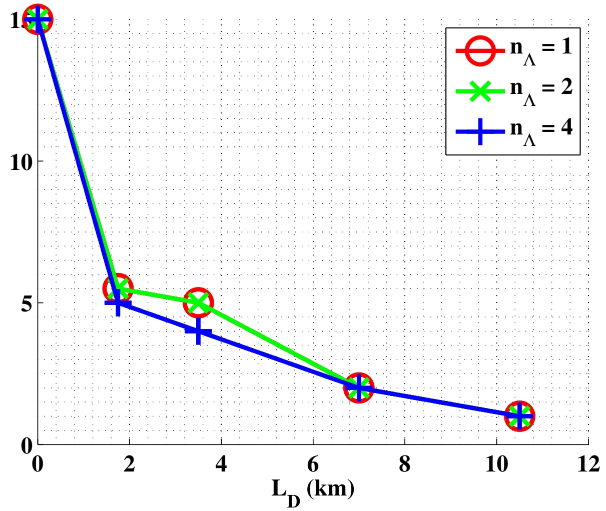

FIGURE 6.4: Number of Hotels, for $K = 10$ and *ES enabled* case.

according to the distributions detailed in Section 4.4. The cost function parameters are set as: $\alpha_{HOT} = 100$, $\alpha_{EL} = 0.1$, $\alpha_{FIB,e} = 0.1$, $\alpha_{LIG} = 0$. This choice gives priority to the minimization of hotel cost, while the cost of each used fiber-link is equal to the cost of electronic switches, and cost of established lightpath is set to zero (considered negligible). In the following, we assume that the processing delay introduced by each traversed electronic switch is negligible ($l_{EL} = 0$), so that it does not decrease the available budget for maximum fronthaul delay. The number of wavelengths varies in $n_\Lambda \in \{1, 2, 4\}$ and the maximum number of fibers available for each link varies in $K \in \{10, 8, 6\}$. Each network instance is optimized under two cases: *ES enabled*, in which electronic switches can be placed in network nodes; and *ES disabled*, in which no electronic switches are used, by fixing $z_i = 0$ for all nodes i.

6.6 Case Study 1: Numerical Results and Discussion

Fig. 6.4 shows the average number of hotels as a function of L_D, for varying n_Λ, $K = 10$ and *ES enabled* case. The curve indicates the obtainable degree of BBU consolidation, which ranges from its obvious maximum value (equal to the number of mobile nodes, i.e., CSs) when $L_D = 0$, in which no hotelling is performed and all BBUs are placed at their CSs; to the minimum value, approaching 1, for increasing L_D, in which the PoP is the only hotel node. The curve exhibits a rapid cliff for lower values of L_D, indicating the effects of the tree-like aggregation infrastructure, which makes available more and more paths from CSs to intermediate nodes candidates to be hotels. Such results fully agree with corresponding ones obtained in previous chapters, thus confirming the validity of optimizing the BBU placement in aggregation networks. It is also worth noting that the degree of consolidation keeps constant under variations of n_Λ, K and *ES enabled/disabled* (results for other values of K and in *ES disabled* case are identical, therefore they have been omitted). This can be interpreted as a predictable effect of the much larger cost associated to hotels, with respect to other costs. However, we have reasons to suspect that the very small considered instance (chosen for computational complexity issues) so far hides some interesting interactions between hotel placement, electronic switching placement and traffic routing, that can be evident in much larger (and realistic) instances, with more complex topologies.

For the *ES enabled* case, the average number of electronic switches is reported in Fig. 6.5, for $n_\Lambda = 1$ and varying K. The figure shows that the number of installed electronic switches ranges in a relatively small interval between around 1 and 3.5. The curves exhibit a minimum around middle values of L_D and two increasing trends. Towards higher values of L_D, the number of electronic switches slightly increases, indicating that the network becomes closer to capacity saturation, as a consequence of the increased amount of high-bitrate fronthaul traffic due to BBU consolidation. Also, the dependency on the value of K is negligible. Towards lowest values of L_D, up to zero, surprisingly the number of electronic switches exhibits a relevant peak, which is also virtually unaffected by K. We interpret this as a side effect due to some non-trivial interaction between installation of electronic switches and network resource utilization that deserves some further investigation (currently in progress).

Fig. 6.6 shows the average number of fiber-links, defined as the sum, over all links, of the number of used fibers per link, for both *ES enabled* and *ES disabled* cases. It is strongly evident the difference between the two cases for electronic switching. For *ES enabled*, the quantity grows almost linearly and does not show relevant differences by varying K. This is due to the overall grooming performed by installed electronic switches, that enables a more efficient packing of traffic into lightpaths, thus reducing the number of needed fibers. For *ES disabled*, the lack of grooming by electronic switches placed at intermediate nodes causes more fibers to be used. Moreover, a clear saturation effect is noticeable for lower K (in this case, $K = 6$) and higher L_D, meaning that in almost all links the number of used fibers is maximum and no further links can be used.

The analysis of the number of established lightpaths, depicted in Fig. 6.7, shows an opposite trend with respect to the number of fiber-links. For the *ES enabled* case, although electronic switches allow to increase lightpath utilization and partially reduce the number of established lightpaths, they generate an even bigger number of them, which explains the growing trend towards higher L_D. On the opposite for the *ES disabled* case, the growth is limited by fiber availability (i.e., K) and by the lightpath-blocking features of the network, which depends on the number of wavelength and the topology.

All the previously analyzed results are relative only to mono-wavelength ($n_\Lambda = 1$) networks, because in case of more wavelengths the computation times increase considerably and do not allow to obtain sufficient sets of data with meaningful average values in reasonable times. Nevertheless, also the basic mono-wavelength case has highlighted some trends and issues that are currently under investigation, in order to design a heuristic algorithm for such optimization problem.

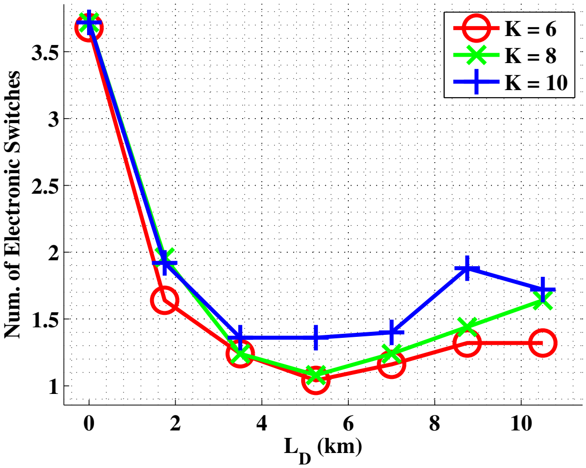

Figure 6.5: Number of Electronic Switches, for $n_\Lambda = 1$ and *ES enabled* case.

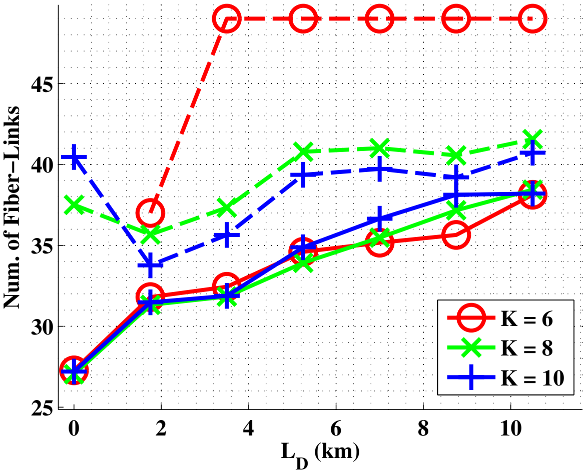

Figure 6.6: Number of Fiber-Links, for $n_\Lambda = 1$ (dashed: *ES disabled*; continue: *ES enabled*).

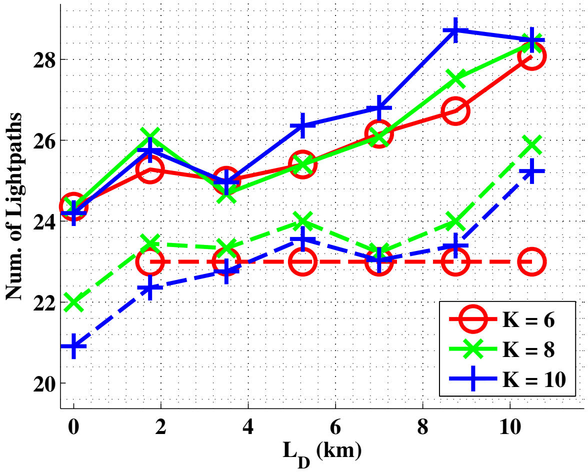

FIGURE 6.7: Number of established Lightpaths, for $n_\Lambda = 1$ (dashed: *ES disabled*; continue: *ES enabled*).

6.7 Case Study 2: Topology and Traffic

In this second case study, we consider “ring and spur” topologies with three stages, similarly to the previous one. However, we now make use or real-world parameters (summarized in Tab. 6.1), which define three different geotypes: Urban, Suburban and Rural. For traffic matrixes, we assume the same characterization of the previous test case.

Differently from the previous case, we now assume a prioritized multi-objective optimization, in which the different terms of the cost function are sorted according to a “priority” order (also denoted as lexicographic optimization). Such priority is explicitly imposed in the objective function by properly assigning relative cost coefficients. Specifically, there are two choices: *minHotel*, in which the primary function is the minimization of the number of hotels and the secondary one is the minimization of the number of used fiber-links; and *minFiber*, in which the two priorities are inverted. In both cases, the minimization of the number of electronic switches has the lowest priority.

In the following results, we assume that installing electronic

Table 6.1: Parameters of the three geotypes: Urban (U), Suburban (S) and Rural (R), per Main CO area.

Geotype	**U**	**S**	**R**
Coverage area (km^2)	15	142	615
Num. of COs (n_{CO})	3	6	11
Num. of CSs (n_{CS})	23	29	31
CO density (per km^2)	0.193	0.042	0.018
CS density (per km^2)	1.5	0.2	0.05

switches is always possible (i.e., *ES enabled*), and we focus our analysis on the comparison between the two fronthaul transport options, *OTN* and *Overlay*. Each numerical value is obtained by averaging over 25 random instances generated according to the previous parameters.

6.8 Case Study 2: Numerical Results and Discussion

Figs. 6.8–6.9 show the ratio R between the number of hotels and the number of CS as a function of the maximum fronthaul round-trip time (τ_{RTT}). The metric R quantifies the degree of BBU consolidation, where $R = 1$ indicates no consolidation (each BBU hosted in the respective CS) and $1/n_{CS}$ indicates the highest degree of consolidation (all BBUs in a single hotel). Starting with Fig. 6.8 (*minHotel*), we can see that R significantly decreases for higher τ_{RTT} since a higher τ_{RTT} allows to route fronthaul flows over longer paths. Comparing the *OTN* (left column) to the *Overlay* (right column) case, we can clearly see that the the *Overlay* case achieves much lower values of R (higher BBU consolidation), suggesting that the higher traffic multiplexing enabled by OTN does not compensate for the increased latency due to OTN encapsulation. The columns in Fig. 6.8 depict also the breakdown of the BBU placement over the different stages, showing that, for higher values of τ_{RTT}, BBUs are placed at higher stages, i.e., closer to the Core CO. Very similar results are shown in Fig. 6.9, where the main objective function is the minimization of number of fibers (*minFiber*), meaning that the achievable degree of BBU consolidation is not affected by changing the objective function (fibers). This result is counter-intuitive as *minFiber*

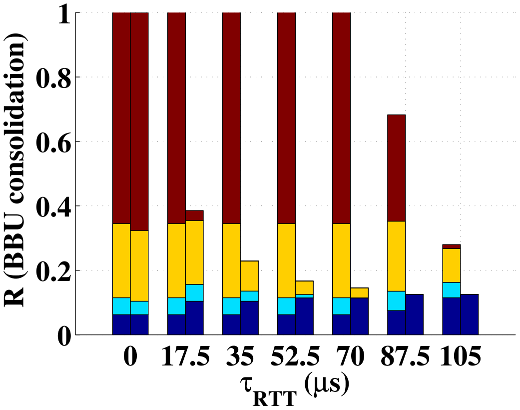

FIGURE 6.8: Per-stage (s) BBU consolidation (R), for Urban geotype, *minHotel*, $n_Lambda = 4$, $K = 10$ (left column: *OTN*; right column: *Overlay*)

should promote much tighter traffic aggregation and discourage BBU hotel consolidation (as fronthaul traffic uses much higher capacity). This result can be explained by the fact that, in our case study, the offered traffic is quite low and requires limited capacity, so the number of fibers can be easily minimized without affecting BBU consolidation. In contrast, if we now decrease the capacity of the network (see Fig. 6.10, where each fiber supports only a single wavelength), we can clearly see that the optimized solution results in a much higher value of R (i.e., less consolidation), especially for large values of τ_{RTT}.

Finally, Fig. 6.11 shows that R strongly depends on the three geotypes (urban, suburban and rural), as a consequence of their different network sizes. Note that the difference between *Overlay* and *OTN* becomes less evident only for larger budget of latency (beyond 200 μs), where the impact of additional switching latency is lower.

In conclusion, using our novel design model, we have shown that the *Overlay* case always ensures a better BBU consolidation than *OTN*, while the amount of available transport resources has negligible impact if the priority is minimizing the number of hotels.

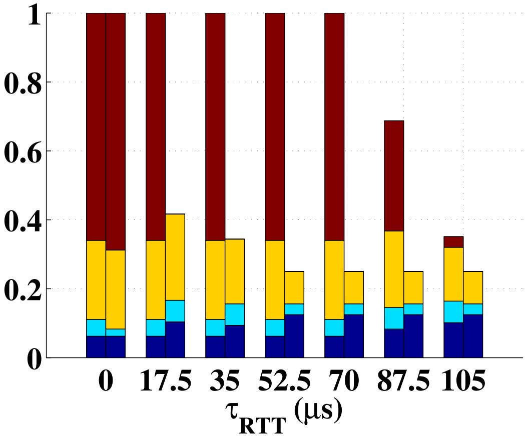

FIGURE 6.9: Per-stage (s) BBU consolidation (R), for Urban geotype, *minFiber*, $n_Lambda = 4$, $K = 10$ (left column: *OTN*; right column: *Overlay*)

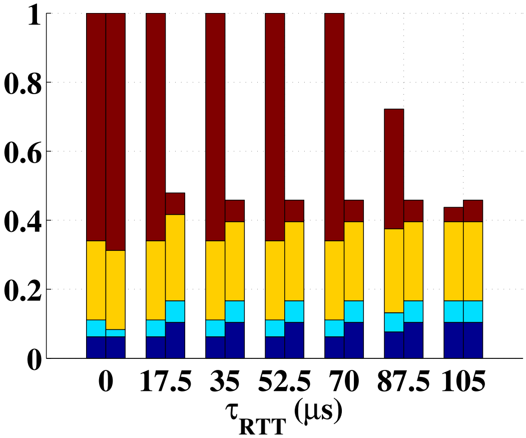

FIGURE 6.10: Per-stage (s) BBU consolidation (R), for Urban geotype, *minFiber*, $n_Lambda = 1$, $K = 6$ (left column: *OTN*; right column: *Overlay*)

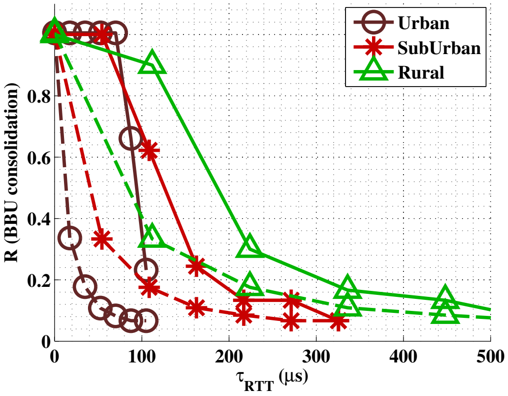

Figure 6.11: BBU consolidation (R), *minHotel*, $n_Lambda = 4$, $K = 6$; for all the geotypes (continue line: *OTN*; dashed line: *Overlay*)

Conclusion 7

In this thesis, we have investigated the energy and cost effectiveness of architectures for next-generation access and aggregation networks, based on two promising principles: *1)* Fixed/Mobile Convergence, i.e., the concept of designing and optimizing networks as “a whole” resorting to infrastructure and equipment sharing; *2)*, BaseBand Unit Hotelling, i.e., the new radio access network paradigm in which BaseBand Units are centralized into hotels and the new fronthaul (CPRI) traffic is introduced into the network.

In Chapter 2, we made a survey of the technological features of BBU hotelling and implementation principles from a networking standpoint.

In Chapter 3, we pursued an analysis of the energetic benefits that this technique can bring over currently aggregation infrastructure (Chapter 3). Results indicates that a first-step fronthaul implementation on a macro-cell based RAN can lead to valuable energy savings (up to 40 to 50%), thanks to an intelligent re-placement of RRHs and BBUs, which is already feasible over existing infrastructures with available commercial technologies.

In Chapter 4, we presented three different architectures of FMC WDM aggregation networks: Bypass, Opaque and No-Hotel. Then we proposed an energy efficient BBU placement optimization problem, whose objective is the minimization of a specific energetic performance indicator, the Aggregation Infrastructure Power (AIP). A fundamental role in the optimization is played by the maximum route length of CPRI flows between cell sites and the respective hotel nodes. Results for multistage tree networks showed values of relative AIP savings around 60 to 65% for Dense-Urban and Urban, and 40% for the Rural scenarios. They also confirmed that BBU consolidation is the dominant principle which pushes toward valuable energy savings for the proposed architectures, in the analyzed case study.

In Chapter 5, the same problem has been generalized towards a more generic “cost” efficiency and larger network sizes, by introducing an heuristic optimization algorithm based on a greedy scheme. Results agree with the same obtained for the energy-efficient problem, and clearly indicate that BBU consolidation into fewer hotel nodes is the key principle for reducing the overall cost/energy and justifies the evolution towards converged architectures, also for larger-scale aggregation networks.

In Chapter 6, a last evolution of the optimization model has presented, which deals with a more real-word scenario featuring the joint optimization of BBU and electronic switches placement, in a multifiber WDM aggregation network. So far, the increased complexity of the problem makes possible only to investigate very small instances. Nevertheless, some interesting trends appear, regarding the interaction between hotel and electronic switches placement, that is in some aspects non-trivial. For this reason, further investigations are needed to clarify such behavior and predict their impact on more realistic network sizes (that would presumably require an ad-hoc heuristic).

Although the presented research is in a initial stage, we believe that the defined optimization problems are enough general and powerful to serve as basis for many possible extensions, in order to study some interesting open issues that emerged during this work.

For instance, the proposed models assume a static or very-slowly varying traffic characterization. A possible extension is about a dynamic traffic scenario, in which the optimization model also embeds the statistical information about daily traffic variation patterns (i.e., the “tidal effect”). Within such scenario, it may be interesting to integrate per-node energy efficiency strategies (e.g., ON/OFF, "sleep" or low-power modes) in the network optimization, in order to dynamically adapt the configuration of some nodes to the varying traffic conditions. Also, dynamic traffic reconfiguration originates a tradeoff between required functionality and cost of transparent wavelength routers, whose relationship with the BBU placement would be interesting to study.

Another promising research line regards more advanced BBU implementations at hotel sites, based on resource pooling, that makes the processing of fronthaul more scalable with variations of cells load. Integrating such technologies in our optimization models would add a new constraint on the association between cell sites and hotel nodes, because the benefits of resource pooling is expected to depend on how the clusters of controlled cells are formed.

On the same line, BBU pooling is one of the candidate technologies enabling a practical implementation of coordinated multicell processing for mobile access, e.g., Cooperative MultiPoint (CoMP), and enhanced Inter-Cell Interference Coordination (eICIC). For non-hotelling (distributed) RAN, a competing solution consists in adopting a centralized controller that aggregates backhaul traffic and manages coordination data of a cluster of cells.

Since such data are subject to some range of latency/bandwidth constraints (depending on the implemented coordination features), there is a tradeoff between placing the controller close enough to cell sites to satisfy such constraints, and move it far to cover larger cell clusters and get more efficient joint processing. Such tradeoff quite resembles that of BBU placement, so a similar optimization problem can be devised, with potentially interesting insights on the evolution of mobile access network architectures in the next future.

Finally, some considerations can be made about the next 5th-Generation (5G) mobile access, whose roll-out is scheduled by the early 2020s. Even though the 5G standardization process has currently not officially started yet, there is no doubt that it will feature unprecedented performance, not only in terms of higher data rates per user and lower latency, but also in terms of "intelligence" of the network. To achieve this, 5G networks are expected to resort to solutions as, for instance, cell densification by means of massive small cell deployment (micro-, femto-, etc.), coordinated multi-cell and multi-antenna processing (CoMP, eICIC) and Centralized/Cloud RAN, possibly in a inter-optimized way. Such techniques will ultimately burden the wireline part of access networks and the aggregation networks, because around tens of Gb/s of data must be transported, due to backhaul, coordination and fronthaul traffic, with sub-ms latency constraints, and optimally adapting to the varying load condition over time. As a consequence, an optimized design of the optical access/aggregation networks, capable of satisfying all such requisites, is expected to have an increased importance, since its performance will directly influence the performance of the 5G network. In other words, it is possible that 5G will emerge more as a unified FMC access/aggregation network, rather than exclusively a mobile access network.

Today, it is very hard to predict if a massive fronthaul deployment will be the dominant solution for future 5G networks. Probably the answer depends on boundary conditions that are currently uncertain, or even unknown. It might happen that large-scale BBU centralization, as we define it in this work, will actually be inconvenient or unfeasible, due to the huge amount of combined fronthaul from macro- and small-cells, or because the new radio interface will further limit the maximum fronthaul latency. Nevertheless, it is unlikely that all the strong motivations behind C-RAN will cease to be effective. Perhaps fronthaul will have to be redefined (e.g., by a different splitting of base-station processing functions), becoming a sort of hybrid between raw baseband data and sophisticated coordination traffic. Also, a novel radio interface might be defined by taking into account the constraints dictated by fronthaul/coordination transport over the underlying optical network.

In any case, it is certain that an optimized design of the access/aggregation network supporting 5G will have more and more importance in the next years. For this reason, we do believe that the research presented in this thesis might have a relevant impact even in this new unpredictable, yet exciting context.

Bibliography

[1] *European Commission, Code of Conduct on Energy Consumption of Broadband Equipment, Version 4.1, Jan 2013.*

[2] J. Chen, L. Wosinska, M. Chughtai, and M. Forzati, "Scalable passive optical network architecture for reliable service delivery," *Optical Communications and Networking, IEEE/OSA Journal of*, vol. 3, no. 9, pp. 667–673, 2011.

[3] M. Nahas, A. Saadani, J. Charles, and Z. El-Bazzal, "Base Stations Evolution: Toward 4G Technology," in *19th International Conference on Telecommunications (ICT), 2012*, april 2012.

[4] A. Saadani, M. El Tabach, A. Pizzinat, M. Nahas, P. Pagnoux, S. Purge, and Y. Bao, "Digital radio over fiber for lte-advanced: Opportunities and challenges." in *ONDM*, 2013, pp. 194–199.

[5] *CPRI Specification V6.0; August 30, 2013; http://www.cpri.info.*

[6] *OBSAI; http://www.obsai.org.*

[7] M. Hadzialic, B. Dosenovic, M. Dzaferagic, and J. Musovic, "Cloud-ran: Innovative radio access network architecture," in *ELMAR, 2013 55th International Symposium*, Sept 2013, pp. 115–120.

[8] D. Sabella, P. Rost, Y. Sheng, E. Pateromichelakis, U. Salim, P. Guitton-Ouhamou, M. di Girolamo, and G. Giuliani, "Ran as a service: Challenges of designing a flexible ran architecture in a cloud-based heterogeneous mobile network," in *Future*

Network and Mobile Summit (FutureNetworkSummit), 2013, July 2013, pp. 1–8.

[9] *European Project "COMBO"; http://www.ict-combo.eu.*

[10] M. Ruffini, L. Wosinska, M. Achouche, J. Chen, N. Doran, F. Farjady, J. Montalvo, P. Ossieur, B. O'Sullivan, N. Parsons, T. Pfeiffer, X.-Z. Qiu, C. Raack, H. Rohde, M. Schiano, P. Townsend, R. Wessaly, X. Yin, and D. Payne, "Discus: an end-to-end solution for ubiquitous broadband optical access," *Communications Magazine, IEEE*, vol. 52, no. 2, pp. S24–S32, February 2014.

[11] P. Chanclou, A. Pizzinat, F. Le Clech, T.-L. Reedeker, Y. Lagadec, F. Saliou, B. Le Guyader, L. Guillo, Q. Deniel, S. Gosselin, S. Le, T. Diallo, R. Brenot, F. Lelarge, L. Marazzi, P. Parolari, M. Martinelli, S. O'Dull, S. Gebrewold, D. Hillerkuss, J. Leuthold, G. Gavioli, and P. Galli, "Optical fiber solution for mobile fronthaul to achieve cloud radio access network," in *Future Network and Mobile Summit (FutureNetworkSummit), 2013*, July 2013, pp. 1–11.

[12] A. Pizzinat, P. Chanclou, T. Diallo, and F. Saliou, "Things you should know about fronthaul," in *Optical Communication (ECOC), 2014 European Conference on*, Sept 2014, pp. 1–3.

[13] A. Pizzinat, P. Chanclou, F. Saliou, and T. Diallo, "Things you should know about fronthaul," *Lightwave Technology, Journal of*, vol. PP, no. 99, pp. 1–1, 2015.

[14] K. Ramantas, K. Vlachos, A. Bikos, G. Ellinas, and A. Hadjiantonis, "New unified pon-ran access architecture for 4g lte networks," *Optical Communications and Networking, IEEE/OSA Journal of*, vol. 6, no. 10, pp. 890–900, Oct 2014.

[15] C. Matrakidis, T. Orphanoudakis, A. Stavdas, J. Fernandez-Palacios Gimenez, and A. Manzalini, "Hydra: A scalable ultra long reach-high capacity access network architecture featuring lower cost and power consumption," *Lightwave Technology, Journal of*, vol. 33, no. 2, pp. 339–348, Jan 2015.

[16] *C-RAN - The Road Towards Green RAN; China Mobile White Paper, Version 2.6 (Sep 2013).*

[17] N. Carapellese, A. Pizzinat, M. Tornatore, P. Chanclou, and S. Gosselin, "An energy consumption comparison of different mobile backhaul and fronthaul optical access architectures," in *Optical Communication (ECOC), 2014 European Conference on*, Sept 2014, pp. 1–3.

[18] "OBSAI (Open Base Station Architecture Initiative); http://www.obsai.org."

[19] J. Lorca and L. Cucala, "Lossless compression technique for the fronthaul of lte/lte-advanced cloud-ran architectures," in *World of Wireless, Mobile and Multimedia Networks (WoWMoM), 2013 IEEE 14th International Symposium and Workshops on a*, June 2013, pp. 1–9.

[20] Z. Ghebretensae, K. Laraqui, S. Dahlfort, J. Chen, Y. Li, J. Hansryd, F. Ponzini, L. Giorgi, S. Stracca, and A. Pratt, "Transmission solutions and architectures for heterogeneous networks built as c-rans," in *Communications and Networking in China (CHINACOM), 2012 7th International ICST Conference on*, Aug 2012, pp. 748–752.

[21] *3GPP TS-36.213 (Physical layer procedures); http://www.3gpp.org.*

[22] A. Reid, P. Willis, I. Hawkins, and C. Bilton, "Carrier ethernet," *Communications Magazine, IEEE*, vol. 46, no. 9, pp. 96–103, September 2008.

[23] *ITU-T Recommendation G.709, "Interfaces for the Optical Transport Network (OTN)".*

[24] N. Carapellese, M. Tornatore, and A. Pattavina, "Placement of base-band units (bbus) over fixed/mobile converged multistage wdm-pons," in *Optical Network Design and Modeling (ONDM), 2013 17th International Conference on*, April 2013, pp. 246–251.

[25] A. SAADANI, M. El Tabach, A. Pizzinat, M. Nahas, P. Pagnoux, S. Purge, and Y. Bao, "Digital radio over fiber for lte-advanced: Opportunities and challenges," in *Optical Network Design and Modeling (ONDM), 2013 17th International Conference on*, April 2013, pp. 194–199.

[26] P. Rost, C. Bernardos, A. Domenico, M. Girolamo, M. Lalam, A. Maeder, D. Sabella, and D. Wübben, “Cloud technologies for flexible 5g radio access networks,” *Communications Magazine, IEEE*, vol. 52, no. 5, pp. 68–76, May 2014.

[27] B. Haberland, F. Derakhshan, H. Grob-Lipski, R. Klotsche, W. Rehm, P. Schefczik, and M. Soellner, “Radio base stations in the cloud,” *Bell Labs Technical Journal*, vol. 18, no. 1, pp. 129–152, June 2013.

[28] R. JAIN and S. Paul, “Network virtualization and software defined networking for cloud computing: a survey,” *Communications Magazine, IEEE*, vol. 51, no. 11, pp. 24–31, November 2013.

[29] P. Chanclou, A. Pizzinat, F. Le Clech, T.-L. Reedeker, Y. Lagadec, F. Saliou, B. Le Guyader, L. Guillo, Q. Deniel, S. Gosselin, S. Le, T. Diallo, R. Brenot, F. Lelarge, L. Marazzi, P. Parolari, M. Martinelli, S. O’Dull, S. Gebrewold, D. Hillerkuss, J. Leuthold, G. Gavioli, and P. Galli, “Optical fiber solution for mobile fronthaul to achieve cloud radio access network,” in *Future Network and Mobile Summit (FutureNetworkSummit), 2013*, July 2013, pp. 1–11.

[30] N. Carapellese, M. Tornatore, and A. Pattavina, “Energy-efficient baseband unit placement in a fixed/mobile converged wdm aggregation network,” *Selected Areas in Communications, IEEE Journal on*, vol. 32, no. 8, pp. 1542–1551, Aug 2014.

[31] R. S. Tucker, R. Parthiban, J. Baliga, K. Hinton, R. W. Ayre, and W. V. Sorin, “Evolution of wdm optical ip networks: A cost and energy perspective,” *Journal of Lightwave Technology*, vol. 27, no. 3, pp. 243–252, 2009.

[32] C. Lange, D. Kosiankowski, C. Gerlach, F. Westphal, and A. Gladisch, “Energy consumption of telecommunication networks,” in *Optical Communication, 2009. ECOC ’09. 35th European Conference on*, 2009, pp. 1–2.

[33] Y. Zhang, P. Chowdhury, M. Tornatore, and B. Mukherjee, “Energy efficiency in telecom optical networks,” *Communi-*

cations Surveys Tutorials, IEEE, vol. 12, no. 4, pp. 441–458, 2010.

[34] F. Musumeci, M. Tornatore, and A. Pattavina, “A power consumption analysis for ip-over-wdm core network architectures,” *Optical Communications and Networking, IEEE/OSA Journal of*, vol. 4, no. 2, pp. 108–117, 2012.

[35] C. Lange and A. Gladisch, “On the energy consumption of ftth access networks,” in *Optical Fiber Communication - incudes post deadline papers, 2009. OFC 2009. Conference on*, 2009, pp. 1–3.

[36] R. Wang, H. H. Lee, S. S. Lee, and B. Mukherjee, “Energy saving via dynamic wavelength sharing in wdm-pon,” in *Optical Network Design and Modeling (ONDM), 2013 17th International Conference on*, 2013, pp. 235–239.

[37] S. Aleksic, M. Deruyck, W. Vereecken, W. Joseph, M. Pickavet, and L. Martens, “Energy efficiency of femtocell deployment in combined wireless/optical access networks,” *Computer Networks*, 2013.

[38] B. Kantarci and H. Mouftah, “Energy efficiency in the extended-reach fiber-wireless access networks,” *Network, IEEE*, vol. 26, no. 2, pp. 28–35, 2012.

[39] *European Project "OASE"; Deliverable D4.2.2, Vers. 1.0, Apr 3 2012.*

[40] K. Zhu and B. Mukherjee, “Traffic grooming in an optical wdm mesh network,” *Selected Areas in Communications, IEEE Journal on*, vol. 20, no. 1, pp. 122–133, 2002.

[41] N. Laoutaris, G. Smaragdakis, K. Oikonomou, I. Stavrakakis, and A. Bestavros, “Distributed placement of service facilities in large-scale networks,” in *INFOCOM 2007. 26th IEEE International Conference on Computer Communications. IEEE*, 2007, pp. 2144–2152.

[42] *European Project "COMBO"; Deliverable D2.1; http://www.ict-combo.eu.*

[43] *European Project "COMBO"; Deliverable D5.1; http://www.ict-combo.eu.*

[44] D. Iida, S. Kuwano, J. ichi Kani, and J. Terada, “Dynamic twdm-pon for mobile radio access networks,” *Opt. Express*, vol. 21, no. 22, pp. 26 209–26 218, Nov 2013. [Online]. Available: http://www.opticsexpress.org/abstract.cfm?URI=oe-21-22-26209

[45] K. Sundaresan, M. Arslan, S. Singh, S. Rangarajan, and S. Krishnamurthy, “Fluidnet: A flexible cloud-based radio access network for small cells,” *Networking, IEEE/ACM Transactions on*, vol. PP, no. 99, pp. 1–14, 2015.

[46] E. G. Coffman, Jr., M. R. Garey, and D. S. Johnson, “Approximation algorithms for np-hard problems,” D. S. Hochbaum, Ed. Boston, MA, USA: PWS Publishing Co., 1997, ch. Approximation Algorithms for Bin Packing: A Survey, pp. 46–93. [Online]. Available: http://dl.acm.org/citation.cfm?id=241938.241940