Costa-Pérez, Xavier; García-Saavedra, Andrés; Li, Xi; Deiss, Thomas; Oliva, Antonio de la; Di Giglio, Andrea; Iovanna, Paola, and Mourad, Alain. 5G-Crosshaul: An SDN/NFV Integrated Fronthaul/Backhaul Transport Network Architecture. *IEEE Wireless Communications*, 24(1), pp. 38-45, February 2017

DOI: https://doi.org/10.1109/MWC.2017.1600181WC

5G-Crosshaul: An SDN/NFV Integrated Fronthaul/Backhaul Transport Network Architecture

Xavier Costa-Pérez, Andres Garcia-Saavedra, Xi Li, Thomas Deiss, Antonio de la Oliva, Andrea di Giglio, Paola Iovanna, and Alain Mourad

Abstract

This article proposes an innovative architecture design for a 5G transport solution (dubbed 5G-Crosshaul) targeting the integration of existing and new fronthaul and backhaul technologies and interfaces. At the heart of the proposed design lie an SDN/NFV-based management and orchestration entity (XCI), and an Ethernet-based packet forwarding entity (XFE) supporting various fronthaul and backhaul traffic QoS profiles. The XCI leverages widespread architectural frameworks for NFV (ETSI NFV) and SDN (Open Daylight and ONOS). It opens the 5G transport network as a service for innovative network applications on top (e.g., multi-tenancy, resource management), provisioning the required network and IT resources in a flexible, cost-effective, and abstract manner. The proposed design supports the concept of network slicing pushed by the industry for realizing a truly flexible, sharable, and cost-effective future 5G system.

5G-Crosshaul Architecture Design Rationale

The goal of the European H2020 5G-Crosshaul project is to build an adaptive, flexible and software-defined architecture for future fifth generation (5G) transport networks integrating multi-technology fronthaul and backhaul segments. The 5G-Crosshaul architecture thus aims to enable a flexible and software-defined reconfiguration of all networking elements through a *unified data plane and control plane* interconnecting distributed 5G radio access and core network functions, hosted on in-network cloud infrastructure.

The control plane needs to include a group of key functional elements (e.g., topology discovery, network monitoring, technology abstraction, provisioning of virtual infrastructure) and their main interfaces toward the applications (northbound interface) and toward underlying technologies (southbound interface). For the control plane design we leverage on software-defined networking (SDN) principles to have unified control, management, and configuration of the 5G multi-technology transport network, and apply network functions virtualization (NFV) to the infrastructure, enabling flexible function placement and cost-effective usage of the 5G-Crosshaul infrastructure resources. The SDN principle allows the separation of the data and control planes, fostering network and device programmability. NFV allows infrastructure and function virtualization, where the underlying physical infrastructure and network functions can be virtualized in such a way that they will be appropriately instantiated, connected, and combined over the underlying 5G-Crosshaul substrate.

The design of the data plane architecture needs to enable the integration of heterogeneous technologies for the fronthaul and backhaul links into a single SDN-based controlled network. The main challenge of the data plane is the need for extended flexibility to adapt to the new fronthaul and backhaul technologies arising with 5G as well as to incorporate legacy technologies through abstraction interfaces.

To achieve such a design, our approach is to leverage the state-of-the-art SDN and NFV architectures so as to avoid re-inventing the wheel and maximize the compatibility and integration of the system design with the existing standard frameworks and reference specifications.

So far the most well developed open source SDN controllers that provide carrier-grade features are Open Daylight (ODL) [1] and Open Network Operating System (ONOS). In the NFV case, the European Telecommunications Standards Institute (ETSI) NFV Industry Specification Group (ISG) currently gives the ability to deploy instances of network functions running in virtual machines (VMs) and provide network operators with the ability to dynamically instantiate, activate, and re-allocate resources and functions.

Based on these industry initiatives and standards, our 5G-Crosshaul architecture is designed to be compatible with the existing ODL/ONOS and ETSI NFV architecture frameworks. For the overall architecture design, we take a bottom-up approach to evolve from current management systems toward the integration of management and network orchestration (MANO) concepts.

Xavier Costa-Pérez, Andres Garcia-Saavedra, and Xi Li are with NEC Laboratories Europe.

Thomas Deiss is with Nokia, Germany.

Antonio de la Oliva is with Univ. Carlos III.

Andrea di Giglio is with Telecom Italia.

Paola Iovanna is with Ericsson, Italy

Alain Mourad is with Interdigital.

Digital Object Identifier: 10.1109/MWC.2017.1600181WC

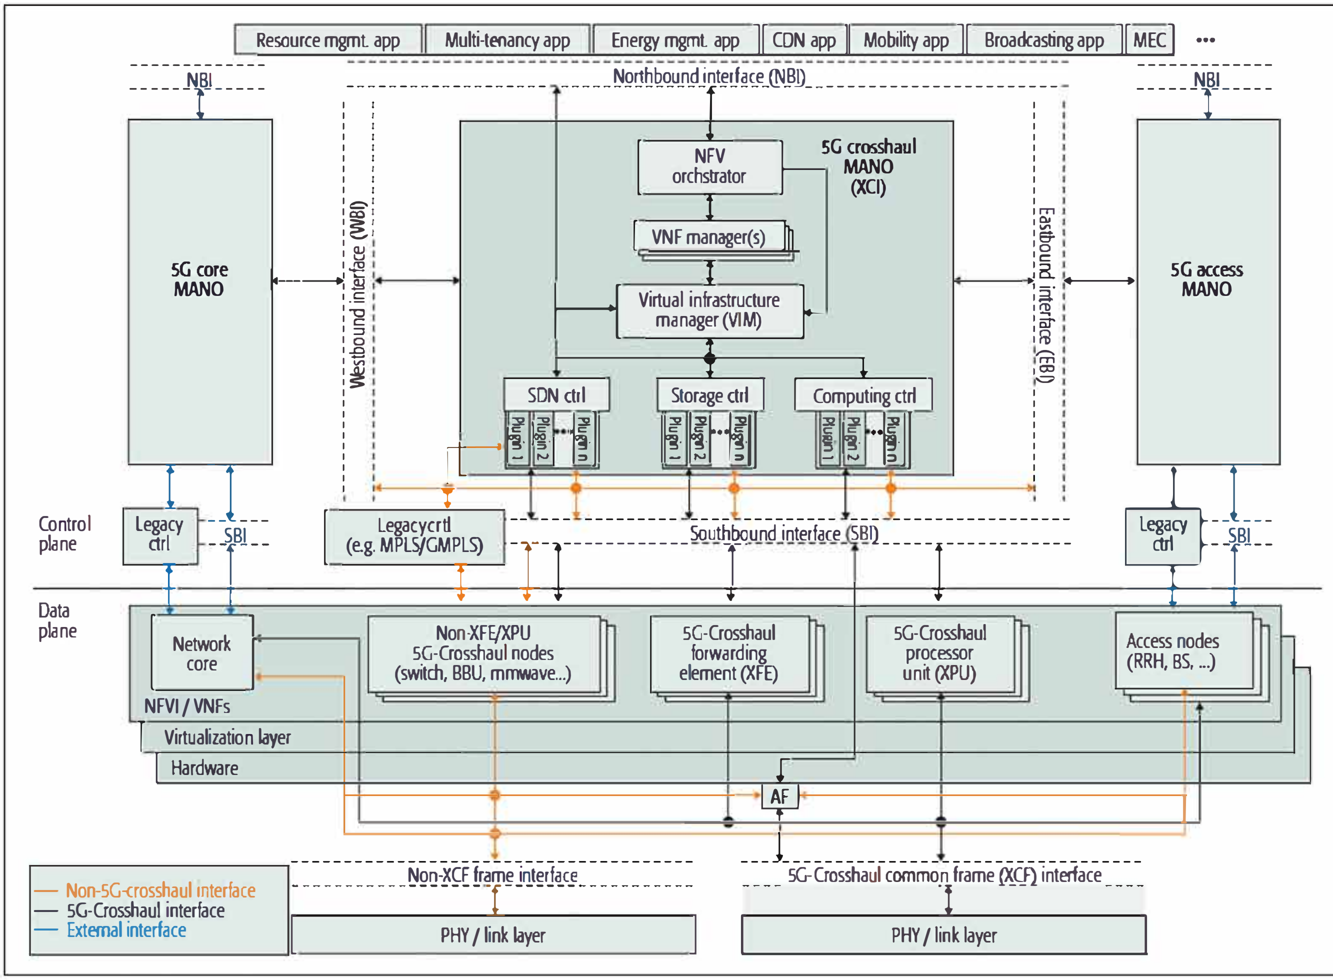

FIGURE 1. 5G-Crosshaul architecture illustration.

5G-Crosshaul Architecture Overview

Related Work

The advent of a new generation of mobile communications is fostering a wealth of research on 5G network architectures that cope with stringent KPI requirements. Practically all of the industry and academic communities agree that both SDN and NFV, as well as heterogeneous transport and access technologies, should play a key role in its design; see, for example, IMT-2020 (5G) Promotion Group's and Next Generation Mobile Network's (NGMN's) white papers [2, 3] as well as DoCoMo's paper in [4] on 5G architecture design principles. It is worth highlighting the work in [4], which presented some early results on the expected system-level gains of some of the novel access technologies. A comprehensive survey of 5G access technologies is published in [5].

However, all this work focuses almost exclusively on the access or core segment of the network, while the unification of fronthaul and backhaul segments to transport data between access and core has received little attention.

In the context of the 5G Public Private Partnership (5G-PPP), the 5G-Xhaul project [6] is addressing the convergence of an optical/wireless backhaul/fronthaul architecture. Their focus is on dynamic reconfigurability with a cognitive control plane for small cells and cloud radio access networks (RANs). In contrast, our goal in the 5G-Crosshaul project is rather to design a unified control and data plane for any type of backhaul and fronthaul traffic applying the SDN and NFV principles.

In our previous work [7] we focused on the benefits attainable to a unified backhaul/fronthaul system. This is in fact in line with the research activities carried out within the scope of the Next Generation Fronthaul Interface Alliance (NGFI), IEEE 1914, and IEEE 802.1, where a packetized version of fronthaul traffic is envisioned, compatible with existing backhaul deployments. This work is therefore highly relevant to bring along the flexibility, data plane interoperability, and system-wide management and optimization of heterogeneous technologies of the integrated 5G transport network, referred to here as 5G-Crosshaul.

System Design

Based on the design criteria exposed previously, we propose the 5G-Crosshaul architecture devised in Fig. 1. Our design follows the same principles of the SDN reference architecture as defined by the Open Networking Foundation (ONF) in [8]:

- Decoupled data plane and control plane
- Logically centralized control
- Exposure of abstract resources and state to applications

Control Plane: As illustrated in Fig. 1, we

5G-crosshaul integrates all communication links between remote radio heads/small cells and core network entities in a unified transport network, by designing a common data plane that enables the integration of heterogeneous technologies for the fronthaul and backhaul links into a single programmable, multi-tenant enabled packet-based network.

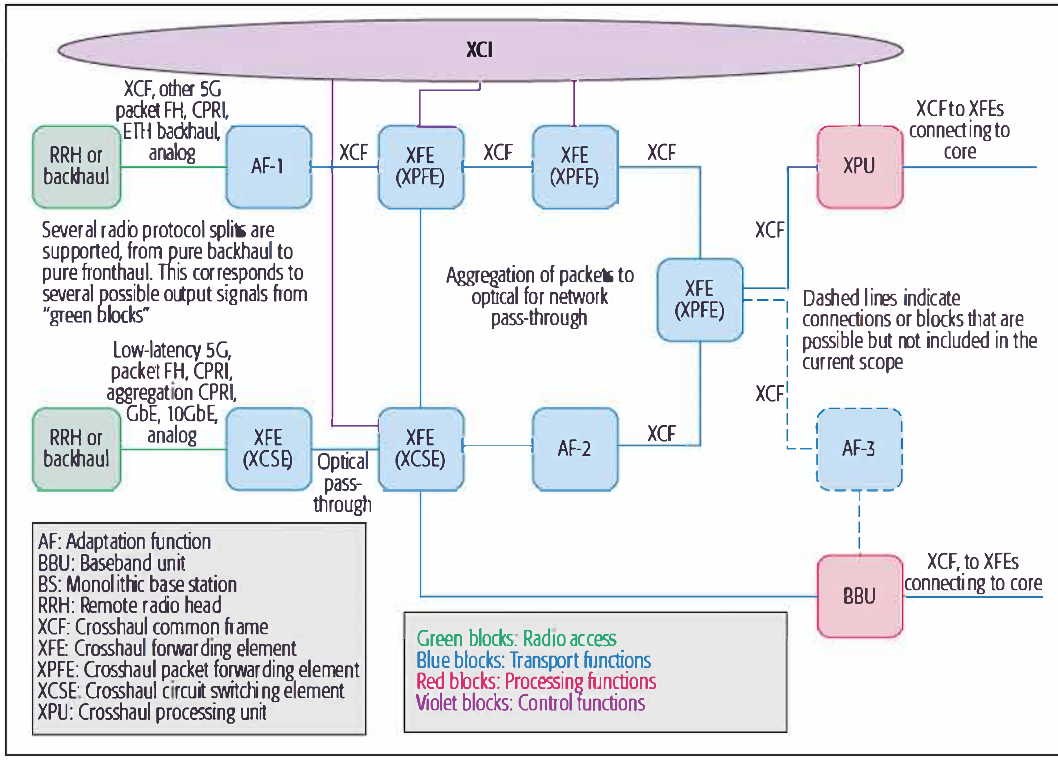

FIGURE 2. 5G-Crosshaul data path architecture.

divide the *control plane* into two clearly differentiated layers: a top layer for external applications and the *5G-Crosshaul control infrastructure* (*XCI*) below. An ecosystem of applications at the top-most part of the system architecture exploits 5G-Crosshaul resource orchestration functions to support diverse functionalities including planning, network and service monitoring/prediction, optimization of resources, energy management, multi-tenancy, content delivery networks, TV broadcasting, and so on.

In turn, the XCI is our 5G transport MANO platform that provides control and management functions to operate all available types of resources (networking and cloud). The XCI is based on the SDN/NFV principles and provides a unified platform that can be used by upper layer applications via a *northbound interface* (*NBI*) to program and monitor the underlying data plane by a common set of core services and primitives. XCI interacts with the data plane entities via a *southbound interface* (*SBI*) in order to:

- Control and manage the packet forwarding behavior performed by 5G-Crosshaul forwarding elements (XFEs) across the 5G-Crosshaul network
- Control and manage the physical layer (PHY) configuration of the different link technologies (e.g., transmission power on wireless links)
- Control and manage the 5G-Crosshaul processing units (XPU) computing operations (e.g., instantiation and management of virtual network functions [VNFs] via NFV)

Data Plane: 5G-Crosshaul integrates all communication links between remote radio heads/small cells and core network entities in a unified transport network, by designing a common *data plane* that enables the integration of heterogeneous technologies for the fronthaul and backhaul links into a single programmable, multi-tenant enabled packet-based network. To this aim, we use XFEs (Fig. 2). XFEs are switching units, based on packet or circuit technology, that interconnect a broad set of link and PHY technologies by means of a novel transport protocol which leverages the *5G-Crosshaul common frame* (XCF). The XCF is designed to handle fronthaul and backhaul traffic simultaneously, although they might have very diverse requirements. Note that this entails the definition of fields for handling traffic prioritization and timing.

In turn, XPUs carry out the bulk of the computing operations in the 5G-Crosshaul. These operations shall support C-RAN, thus hosting baseband units (BBUs) or medium access control (MAC) processors, but also 5G points of attachment (5GPoAs) functionalities that can be virtualized (VNFs) and a heterogeneous set of other services, such as content delivery network (CDN)-based services. In this manner, the *NFV infrastructure* (*NFVI*) comprises all data plane (software and hardware) components that build up the networking environment in which VNFs are deployed and connected.

Of course, with backward compatibility in mind, XCI can communicate with non-5G-Crosshaul-specific entities, such as legacy switches, BBUs, and millimeter-wave (mmWave) switches using proper *plugins*. 5G-Crosshaul-specific data plane elements (XFEs, XPUs) can communicate with non-XCF-compliant ones by means of an *adaptation function* entity (AF; Fig. 2) that acts as a translator between XCF and other protocols.

Interfaces: As mentioned above, an ecosystem of applications sits on top of the XCI to provide tools for optimization, prediction, energy management, multi-tenancy, and others. The XCI is the means to achieve the application goals, and the interface, typically based on REST, NETCONF, or RESTCONF application programming interfaces (APIs) that interconnect both domains, is an NBI.

The configuration of network resources, computing resources, and storage resources is directly executed on each of the required data plane elements by the XCI by means of the SBI. Candidates for SBI are OpenFlow, OF-Config, OVSDB, Simple Network Management Protocol (SNMP), and/or an ecosystem comprising several of them.

The scope of operation of the XCI is limited to (physical/virtual networking/storage/computing) resources within the 5G-Crosshaul transport domain. However, given that proper optimization of the data plane elements may require knowledge of the configuration and or other information from the core network and/or RAN domains, our system design contemplates a *westbound interface* (WBI) to communicate with the 5G core MANO and an *eastbound interface* (EBI) to interact with the 5G access MANO.

5G-Crosshaul Architecture Main Components

In the following we describe in detail the 5G-Crosshaul main architecture building blocks briefly introduced in the previous section.

5G-Crosshaul Control Infrastructure

The XCI is the brain controlling the overall operation of the 5G-Crosshaul. The XCI part dealing with NFV comprises three main functional blocks: NFV orchestrator, VNF manager(s), and virtual infrastructure manager (VIM) (following the ETSI NFV architecture [9]).

The **NFV orchestrator (NFVO)**: This is the functional block that manages a network service (NS) life cycle. It coordinates the VNF life cycle (supported by the VNFM) and the resources available at the NFV infrastructure (NFVI) to ensure an optimized allocation of the necessary resources and connectivity to provide the requested virtual network functionality.

The **VNF managers (VNFMs)**: These functional blocks are responsible for the life cycle management of VNF instances (e.g., instance instantiation, modification, and termination).

The virtualized infrastructure manager (VIM): This functional block is responsible for controlling and managing the NFVI computing (via *Computing ctrl*), storage (via *Storage ctrl*), and network resources (via *SDN ctrl*).

In addition to these modules, which are in charge of managing the different VNFs running on top of the 5G-Crosshaul, the XCI includes a set of specialized controllers to deal with the control of the underlying network, storage, and computation resources.

SDN controller: This module is in charge of controlling the underlying network elements following the conventional SDN paradigm. 5G-Crosshaul aims at extending current SDN support of multiple technologies used in transport networks (e.g., microwave links[1]) in order to have a common SDN controlled network substrate that can be reconfigured based on the needs of the network tenants.

Computing/storage controllers: Storage and computing controllers are included in what we call a cloud controller. A prominent example of this kind of software framework is OpenStack.

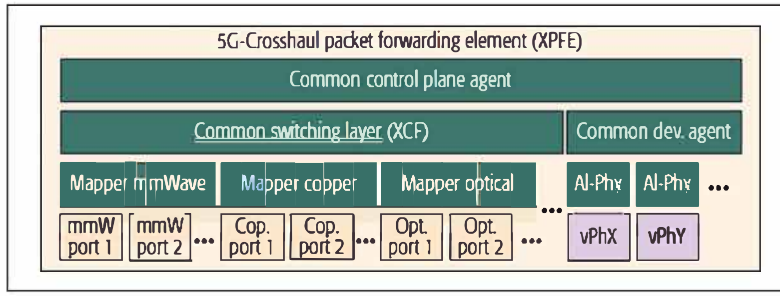

Figure 3. XPFE functional architecture.

Note that the *SDN/computing/storage controllers* are functional blocks with one or multiple actual controllers (hierarchical or peer-to-peer structure) that centralize some or all of the control functionality of one or multiple network domains. We consider the utilization of legacy network controllers (e.g., MPLS/GMPLS) to ensure backward compatibility for legacy equipment.

5G-Crosshaul Forwarding Element

XFEs are switching units that support single or multiple link technologies (mmWave, Ethernet, fiber, microwave, copper, etc.). A key part of the envisioned solution is a common switching layer in the XFEs for enabling unified and harmonized transport traffic management. This common switching layer supports the 5G-Crosshaul common frame (XCF) format across the various traffic flows (of fronthaul and backhaul) and the various link technologies in the forwarding network. The common switching layer in the XFEs is controlled by the XCI, which is foreseen to have a detailed (as per the abstraction level defined) view of the fronthaul and backhaul traffic and resources, and to expose this detailed view through a further abstraction to the orchestration layer to enable intelligent resource, network function, and topology management across the two domains.

As depicted in Fig. 2, XFEs include packet switching elements (XPFEs) and circuit switching elements (XCSEs). Two paths are defined in this framework: a packet switching path (upper part) and an all-optical circuit switching path (lower part). The packet switching path is the primary path for the transport of most delay-tolerant fronthaul and backhaul traffic, whereas the circuit switching path is there to complement the packet switching path for those particular traffic profiles that are not suited for packet-based transport (e.g., legacy common public radio interface [CPRI] or traffic with extremely low delay tolerance). This two-path switching architecture is able to combine bandwidth efficiency through statistical multiplexing in the packet switch, with deterministic latency ensured by the circuit switch. The modular structure of the 5G-Crosshaul switch, where layers may be added and removed, enables various deployment scenarios with traffic segregation at multiple levels, from dedicated wavelengths to virtual private networks (VPNs), which is particularly desirable for multi-tenancy support.

Figure 3 depicts an initial functional architecture for the 5G-Crosshaul XPFE. It includes the following key functions:

[1] ONF is actively working toward the definition of a southbound interface for micro-wave links: http://5g-crosshaul.eu/wireless-transport-sdn-proof-of-concept/

Use cases	Description
Dense urban society	This use case addresses the connectivity required at any place and at any time by humans in dense urban environments, considering both traffic between humans and the cloud, and direct information exchange between humans and/or the environment.
Mobile edge computing	This use case is focused on the deployment of IT and cloud computing capabilities toward the edge of the network. Content, service, and application providers can leverage on such distributed computing capabilities to serve high-volume and latency-sensitive traffic on dense areas with a high number of users.
Media distribution: CDN/TV broadcasting	This use case addresses the distribution over 5G networks of media contents, especially video traffic, and TV broadcasting, which are expected to be the dominant contributors to the mobile data traffic demand.
Vehicle mobility	This use case addresses the support of 5G communication in vehicles during motion, such as passengers using 5G services as real-time video on a very high-speed train (500 km/h) and messages among vehicles for traffic control, emergency, and safety.
Multi-tenancy/network slicing	This use case addresses the dynamic allocation of backhaul/fronthaul network slices across multiple tenants. It is a key enabler to maximize the utilization of 5G-Crosshaul infrastructure resources in a cost-efficient manner.

TABLE 1. 5G-Crosshaul use cases.

- A common control-plane agent to talk to the common control infrastructure (XCI).
- A common switching layer based on the common frame (XCF) to forward packets between technology-independent interfaces. The switching engine is technology-agnostic and relies on an abstract resource model (i.e., bandwidth, latency, bit error rate, jitter, latency,) of the underlying interfaces (i.e., mmWave, optical, etc.); and on traffic requirements (i.e., fronthaul/backhaul, jitter tolerance, packet loss, etc.) that can be carried in the XCF.
- A common device agent to talk with system peripheral. This agent exposes device-related information, including CPU usage, RAM occupancy, battery status, GPS position, and so on, to the control infrastructure.
- Mappers for each physical interface.
- Physical interfaces to transmit the data on the link. Multiple physical interfaces of different technologies can coexist in the unit.

The common control plane and device agents are relevant for both packet- and circuit-switched forwarding elements of the XFE. In the XPFE, the common abstraction of the heterogeneous data plane provides a technology-independent data plane and allows dynamic reconfiguration of the transport resources. It also allows interworking with transport legacy technology. That function is enabled by the SBI, which allows exposing legacy domains to the XCI.

5G-Crosshaul Common Frame

The XCF is the frame format used by the XPFE. Ideally, the XCF is supported by all physical interfaces where packets are transported. Circuit-switched forwarding is independent of the XCF. Where necessary, the frame format of the endpoints is mapped to the XCF for forwarding by the XPFEs. As an example, CPRI over Ethernet would have to be mapped to the XCF. Mapping functions are also used among XPFEs and legacy switches.

The XCF is based on Ethernet, utilizing MAC-inMAC (or provider backbone bridged network) [10]. MACinMAC, or alternatively QinQ (or provider bridged network) [10], provides a more flexible separation of tenants compared to VLANs. Networks of different tenants can be separated via the outer MAC header; nevertheless, within one tenant there can be several virtual customer networks. The priority bits of the Ethernet header are used to indicate the priorities of the different traffic flows. Basing the XCF on Ethernet eases reuse of legacy switches and increases synergies with the development of more generic switches.

5G-Crosshaul Processing Unit

While the SDN control platform is responsible for the configuration of the network elements of the 5G-Crosshaul physical infrastructure (i.e., the XFEs), the cloud and storage control platform of the XCI handles the 5G-Crosshaul IT components (computing and storage resources) in the XPUs. Virtual infrastructure is instantiated, configured, and operated by XCI in XPUs, where VNFs can be deployed to run the 5G-Crosshaul services in a proper and efficient manner.

5G-Crosshaul Innovative Apps

Use Cases

The 5G fronthaul and backhaul integration enables a new set of use cases that are summarized in Table 1.

5G-Crosshaul Applications

Based on the above mentioned use cases, in the following we provide a set of relevant examples of the novel applications under development by the 5G-Crosshaul project partners.

Resource Management Application (RMA): This application provides logically centralized and automated management of 5G-Crosshaul resources to promptly provision transport services according to their service level agreements (SLAs) while ensuring effective resource utilization. The RMA can operate over physical or virtual network resources on a per-network or per-tenant basis, respectively. Essentially, the RMA has two main functional pillars:

1. Dynamic resource allocation and (re-) configuration (e.g., new routes or adaptation of physical parameters) as the demand and network state changes
2. Dynamic NFV placement (e.g., enabling multiple cloud-RAN functional splits flexibly allocated across the transport network)

Multi-Tenancy/Network Slicing Application (MTA): This application is designed to enable a generalized, dynamic network slicing of the 5G-Crosshaul infrastructure by multiple network operators or service providers (i.e., multiple tenants), each operating on a slice of the physical resources by virtualization techniques. The target of this application is to significantly reduce capital expenditure (CAPEX) and operational expenditure (OPEX) by jointly exploiting the infrastructure resources in a cost-efficient manner. The MTA is envisioned to be used not only by mobile virtual network operators (MVNOs) but also by over-the-top (OTT) service providers to quickly deploy novel services. MTA allows network slice resources to be dynamically allocated to tenants on demand providing per-tenant monitoring of network quality of service (QoS) and resource usage. The main challenges here are for example to ensure isolation across tenants; and manage (instantiate, reconfigure, remove) tenants at small time-scales in a seamless manner.

Content Delivery Network Management Application (CDNMA): A CDN is a combination of a content-delivery infrastructure (in charge of delivering copies of content to end users), a request routing infrastructure (which directs client requests to appropriate replica servers), and a distribution infrastructure (responsible for keeping an up-to-date view of the content stored in the CDN replica servers). This application is designed to manage the transport resources for a CDN infrastructure, controlling load balancing over several replica servers strategically placed at various locations to deal with massive content requests while improving content delivery based on efficient content routing across the 5G-Crosshaul fronthaul and backhaul network segments and the corresponding user demands.

5G-Crosshaul Resource Management

The presented system architecture has been designed to support all 5G-Crosshaul applications and use cases in an adaptive and flexible manner. In this section, we analyze how to realize 5G-Crosshaul resource management based on the proposed system architecture, as an example to show the method of implementation by means of leveraging the existing open source projects.

Functional Blocks for Resource Management

Resource management is one of the most important and fundamental functions to provide central and automated management of the 5G-Crosshaul transport network to support different applications and services. In the scope of this work, resources include not only networking but also computing and storage resources. Moreover, the resource manager can operate over physical as well as virtual resources, on a per-network or a per-tenant basis. Hence, to access and control these resources it requires the support of active functional elements of the 5G-Crosshaul MANO (XCI) including different controllers to properly operate them. We define four main functional blocks in the XCI that constitute the resource management ecosystem, as shown in Fig. 4.

Resource Management Application (RMA): The RMA is the decision entity in charge of making optimized decisions on the control and management of the underlying network, computing, and storage resources. The RMA collects different types of information from the underlying network infrastructure via the NBI, and runs optimization algorithms for context-aware system-wide resource allocation. Such decisions, for example, routing and re-routing, coordinated transmission power control, and algorithms to configure local scheduling, as programmatic control of the abstracted resources, will be conveyed through the NBI to the XCI, which is responsible for enforcing them.

Resource Management Orchestrator (RMO): The RMO is a building block inside the NFV orchestrator (NFVO) that allows the orchestration of virtualized resources to support the instantiation of VNFs when required by upper applications (e.g., OTT services). The RMO receives requests from the applications (not necessarily only RMA) in the form of VNF templates (CPU, memory, IP, policies, service function chain templates to interconnect a set of VNFs, etc.). Upon receiving such requests, the RMO will first evaluate the VNF policies, provision the required resources if they are available, and accept or reject the request. Once the request is granted by the RMO, it in turn will request the VIM to instantiate the allocation and provision of required network, computing, and storage resources.

Resource Management VNF Manager (RMVM): The RMVM is a building block inside the VNF managers to deal with life cycle management of VNF instances. This is needed when the RMA decides to run one or multiple VNF instances in the network and even connect different VNF instances of a service function chain (e.g., allowing flexible RAN functional splits as a resource management decision). The deployment and operational behavior of each VNF is captured in a template called a virtualized network function descriptor (VNFD) that is stored in the VNF catalog. A VNFD is used to create instances of the VNF it represents, and to manage the life cycle of those instances.

Resource Management Provisioner (RMP): The RMP is a building block inside the VIM as the decision enforcing entity to do the actual provision and allocation of the requested resources by talking to different controllers (i.e., SDN controller, storage controller, and computing controller) depending on the type of required resources. Correspondingly, the SDN controller will compute the paths and provision the required network resources to connect between the VM endpoints. The storage and computing controller will allocate the required IT resources (CPU, memory) to instantiate the VMs.

Leveraging on Open Source Projects

One of the main goals of the 5G-Crosshaul architecture is to enable the reuse of ongoing open source projects as much as possible to facilitate its deployability and compatibility, and to minimize the implementation costs. Figure 4 illustrates the rich ecosystem of such projects that we can exploit in our system.

When dealing with physical resources, RMA can orchestrate its optimizations by interacting directly with a controller. In the case of networking resources, open source projects such as ODL [1] and ONOS [11] can be used for infrastructure discovery, event reporting, monitoring, and execution of RMA's optimizations into underlying networks by means of protocols such as OpenFlow and NETCONF that interact with physical elements through the SBI. Similarly, storage and computing resources can be controlled with well-known tools, as shown in Fig. 4.

Virtualized resources can be managed via an NFVO like OpenMANO, OpenStack's Tacker, or OpenBaton. These are ongoing projects that

The deployment and operational behavior of each VNF is captured in a template called Virtualized Network Function Descriptor (VNFD) that is stored in the VNF catalogue. VNFD is used to create instances of the VNF it represents, and to manage the lifecycle of those instances.

The RMA is the decision entity in charge of making optimized decisions on the control and management of the underlying network, computing and storage resources. The RMA collects different types of information from the underlying network infrastructure via NBI, and runs optimization algorithms for context-aware system-wide resource allocation.

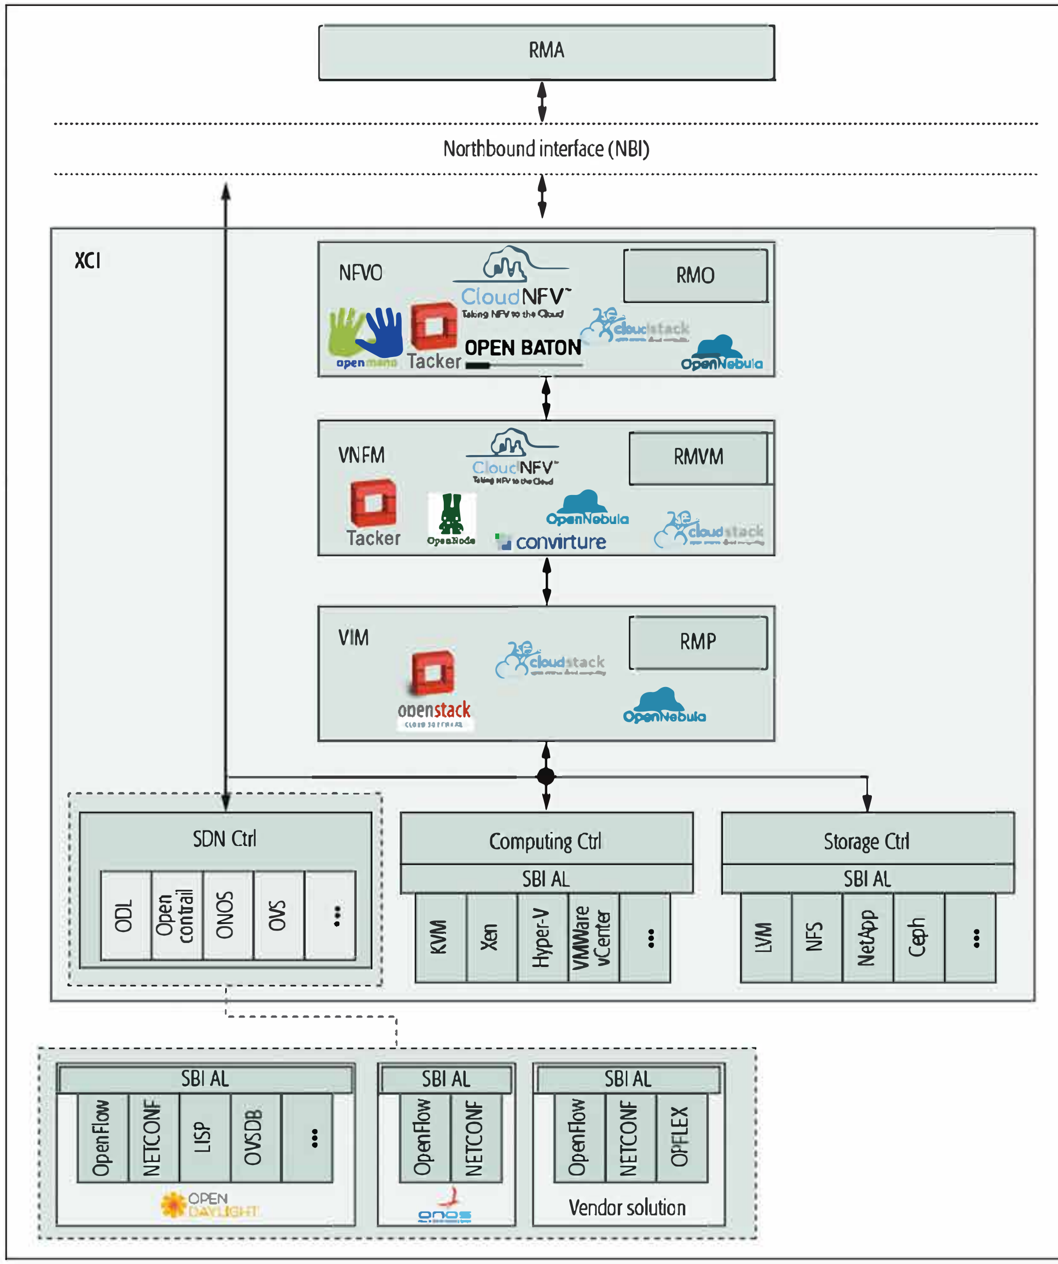

FIGURE 4. 5G-Crosshaul resource management: functional mapping to the system architecture.

closely follow the reference architecture of NFV MANO and integrate easily with OpenStack. An updated survey on MANO projects can be found in [12]. In the case of Tacker, a VNF manager is built in, while OpenBaton can interact with different VNFM solutions. Because most of the cloud platforms use it, we advocate relying on OpenStack as our VIM to orchestrate and manage virtual networking resources (via OpenStack's Neutron), computing resources (using OpenStack's Nova), and storage resources (integrating OpenStack's Cinder, Glance, and Swift). For instance, in the case of virtual networking, Neutron interfaces can be found in ODL for the control of OpenVSwitch virtual switching infrastructure (a de facto standard) via OVSDB to manage it and OpenFlow to configure its forwarding behavior. Successful integration between OpenStack and OpenDaylight has been demonstrated, for instance, in [13].

SUMMARY AND CONCLUSIONS

This article presents an innovative architecture for integrating existing and new fronthaul and backhaul networks into a flexible unified 5G transport solution. The architecture defines key building blocks and interfaces in the data, control, and application planes. These included: (1) an SDN/NFV-based management and orchestration (MANO) entity, referred to as XCI; (2) an Ethernet-based packet forwarding entity, referred to as XPFE; and (3) an NFV-enabled processing entity, referred to as XPU. The XCI leveraged the ETSI NFV architecture framework for resource orchestration and instantiation, as well as open source initiatives for SDN control such as OpenDaylight and ONOS. The XCI opens multiple interfaces, northbound toward the network application layer, southbound toward the data forwarding layer, eastbound toward the 5G access MANO, and westbound toward the 5G core MANO. An instantiation of the XCI is also provided showing the underlying functional blocks for the management of network and IT resources in such a way as to cater for the various applications on top.

The data forwarding plane features a packet switching entity (the XPFE) along with a circuit switching entity (the XCSE) to support extremely low-latency requirements. The XPFE features a common switching layer based on a common frame format (the XCF) supporting various existing and new fronthaul and backhaul traffic profiles. The XCF is based on Ethernet to lower costs and enable economies of scale, while it utilizes MACinMAC extensions to deliver carrier-grade QoS including multi-tenancy/network slicing support. The support of non-XCF (e.g., legacy or proprietary) switching infrastructure is also anticipated through an adaptation function that adapts to the common XCF domain.

The XCF is based on Ethernet to lower costs and enable economies of scale, whilst it utilizes MACinMAC extensions to deliver carrier-grade QoS including multi-tenancy/network slicing support. The support of non-XCF switching infrastructure is also anticipated through an adaptation function that adapts to the common XCF domain.

Acknowledgment

This work has been funded by the EU H2020 project "5G- Crosshaul: The 5G Integrated Fronthaul/Backhaul" (Grant no. 671598).

References

[1] J. Medved *et al.*, "OpenDaylight: Towards a Model-Driven SDN Controller Architecture," *Proc. 2014 IEEE 15th Int'l. Symp. a World of Wireless, Mobile and Multimedia Networks*, June 2014, pp.1-6.

[2] IMT-2020 (5G) Promotion Group, 5G Network Technology Architecture White Paper; www.imt-2020.cn/en/documents/download/64.

[3] NGMN Alliance. 5G White Paper; https://www.ngmn.org/uploads/media/NGMN_5G_White_Paper_V1_0.pdf.

[4] P. Agyapong *et al.*, "Design Considerations for a 5G Network Architecture," *IEEE Commun. Mag.*, vol. 52, no. 11, Nov. 2014, pp. 65–75.

[5] A. Gupta and R. K. Jha, "A Survey of 5G Network: Architecture and Emerging Technologies," *IEEE Access*, vol. 3, 2015, pp. 1206–32.

[6] Horizon 2020, 5GPPP: 5G Xhaul Project; https://5g-ppp.eu/5g-xhaul/.

[7] A. De La Oliva *et al.*, "Xhaul: Toward an Integrated Fronthaul/Backhaul Architecture in 5G Networks," *IEEE Wireless Commun.*, vol. 22, no. 5, Oct. 2015, pp. 32–40.

[8] ONF SDN Architecture 1.0 (TR-502).

[9] ETSI GS NFV 003 V1.2.1.

[10] IEEE Std. 802.1Q-2014, "Bridges and Bridged Networks."

[11] ONOS overview white paper, ONOS project; http://onosproject.org/wp-content/uploads/2014/11/Whitepaper-ONOS-final.pdf.

[12] R. Mijumbi *et al.*, "Management and Orchestration Challenges in Network Functions Virtualization," *IEEE Commun. Mag.*, vol. 54, no. 1, Jan. 2016, pp. 98–105.

[13] A. Csoma *et al.*, "Multi-Layered Service Orchestration in a Multi-Domain Network Environment," *Proc. 2014 Third Euro. Wksp. Software Defined Networks*, 1–3 Sept. 2014, pp. 141–42.

Biographies

Xavier Costa-Pérez is head of 5G Networks R&D at NEC Laboratories Europe, where he manages several projects focused on 5G mobile core, backhaul/fronthaul, and access networks. He is a 5G-PPP Technology Board member and the Technical Manager of the 5G-Crosshaul project. His team contributes to projects for NEC products roadmap evolution as well as to European Commission R&D collaborative projects and has received several NEC R&D Awards for successful technology transfers. In addition, the 5G team contributes to related standardization bodies: 3GPP, BBF, ETSI NFV, ETSI MEC, and IETF. He has served on the Program Committees of several conferences (including IEEE Greencom, WCNC, and INFOCOM), published at top research venues, and holds over 25 patents. He received both his M.Sc. and Ph.D. degrees in telecommunications from the Polytechnic University of Catalonia (UPC) and was the recipient of a national award for his Ph.D. thesis.

Andres Garcia-Saavedra received his M.Sc and Ph.D. from University Carlos III of Madrid (UC3M) in 2010 and 2013, respectively. He then joined the Hamilton Institute, Ireland, as a research fellow until the end of 2014 when he moved to Trinity College Dublin. Since July 2015 he has been working as a research scientist at NEC Laboratories Europe. His research interests lie in the application of fundamental mathematics to real-life computer communications systems, and the design and prototyping of wireless communication systems and protocols.

Xi Li received her M.Sc in electronics and telecommunication engineering from Technische Universität Dresden in 2002 and her Ph.D. from the University of Bremen in 2009. Between 2003 and 2014 she worked as a research fellow and lecturer at the Communication Networks Group in the University of Bremen, leading a team working on several industrial and European R&D projects on 3G/4G mobile networks and future Internet design. From 2014 to 2015 she worked as a solution designer in Telefonica Germany GmbH & Co. OHG, Hamburg. Since March 2015 she has been a senior researcher in 5G Networks R&D at NEC Laboratories Europe.

Antonio De La Oliva (aoliva@it.uc3m.es) received his telecommunications engineering degree in 2004 and his Ph.D. in 2008 from UC3M, where he has been working as an associate professor since then. His current line of research is related to networking in extremely dense networks. He is an active contributor to IEEE 802 where he has served as Vice-Chair of IEEE 802.21b and Technical Editor of IEEE 802.21d. He has also served as a Guest Editor of *IEEE Communications Magazine*. He has published more than 30 papers on different areas of networking including mobility and wireless networks.

Paola Iovanna received her degree in electronics engineering from the University of Roma "Tor Vergata" in 1996. Since 2014 she has led a research team to define transport networking and control solutions for 5G. She is actively involved in European projects and is on the Technical Program Committees of international conferences like ECOC. She holds more than 40 patents in routing, traffic engineering systems, and PCE solutions for packet-opto networks based on GMPLS, and multi-domain SDN transport, fronthaul, and backhaul solutions for 5G, and is an author of several tens of publications in either international scientific journals or conferences.

Thomas Deiss received his degree in computer science in 1990 and his Ph.D. in 1999 from the University of Kaiserslautern. He joined Nokia in 1999. He has contributed to standardization in the area of automated testing and worked in requirements engineering for backhaul functionality of WCDMA and LTE base stations with a focus on backhaul sharing among radio technologies.

Andrea Di Giglio received a Dr.Ing. degree in electronic engineering from the University of Pisa, Italy, and Scuola Santa Anna. He is with Telecom Italia Lab. His fields of research are Internet security and optical networks architecture. He has been involved in several European Projects: the architectural Work Package of the IST Project NOBEL, DISCUS as Work Package leader, and ICT STRONGEST as project coordinator.

Alain Mourad is currently leading the R&D of 5G wireless systems at InterDigital Europe Ltd, United Kingdom. Prior to joining InterDigital, he was a principal engineer at Samsung Electronics R&D, United Kingdom, and previously a senior engineer at Mitsubishi Electric R&D Centre Europe, France. Throughout his career, he has been active in the research and standardization of recent wireless communication and broadcasting systems (3GPP LTE, WiMAX 16m, DVB-T2/NGH, ATSC 3.0). He has several inventions and has authored numerous peer-reviewed scientific publications. He received the Inventor of the Year Award from Samsung Electronics R&D UK in 2012 and 2013.