Orchestration of End-to-End Network Services in the 5G-Crosshaul Multi-Domain Multi-Technology Transport Network

Jorge Baranda, Josep Mangues-Bafalluy, Iñaki Pascual, José Núñez-Martínez, Juan Luis De la Cruz, Ramon Casellas, Ricard Vilalta,
Josep X. Salvat, and Charles Turyagyenda

The authors validate the flexibility, scalability, and recovery capabilities of the 5G-Crosshaul architecture in a testbed distributed geographically. More specifically, the central component of the validation is the hierarchical 5G-XCI, conceived to handle multi-domain multi-technology transport network resources. Its performance is characterized through two experimental case studies.

ABSTRACT

Upcoming 5G mobile networks are addressing ambitious KPIs not just in terms of capacity and latency, but also in terms of network control and management. In this direction, network management schemes need to evolve to provide the required flexibility, and automated and integrated management of 5G networks. This also applies to the 5G-Crosshaul transport network, which provides an integrated fronthaul and backhaul. Software defined networking and NFV are seen as key enablers for that. This article validates the flexibility, scalability, and recovery capabilities of the 5G-Crosshaul architecture in a testbed distributed geographically. More specifically, the central component of the validation is the hierarchical 5G-XCI, conceived to handle multi-domain multi-technology transport network resources. Its performance is characterized through two experimental case studies. The first one illustrates the automated provisioning of all network resources required to deploy a complete LTE virtual mobile network featuring fronthaul and backhaul configurations. This takes 10.467 s on average for the network under test. The second one exploits the flexibility of the hierarchical XCI to apply local or centralized service recovery in the event of link failure depending on the desired path optimality vs. recovery time trade-off. On average, recovery takes 0.299 s and 6.652 s, respectively. Overall, the proposed solution contributes to attaining the target set for 5G networks of reducing service setup from hours to minutes.

INTRODUCTION

Upcoming fifth generation (5G) networks are addressing ambitious key performance indicators (KPIs) in terms of not only overall/per user capacity and latency, but also network control and management. Diverse traffic profiles must be served, which are generally classified as enhanced mobile broadband (eMBB), massive machine type communications (mMTC), and ultra-reliable low-latency communications (URLLC). Furthermore, 5G is not only about the redefinition of the air interface, but it also requires coordinated handling of heterogeneous network segments (access, transport, and core) toward more efficient operation in the context of declining average revenues per user (ARPU).

In this article, we focus on the transport segment, where the increasing stress on radio access performance requires the evolution of current backhaul (BH) and fronthaul (FH) transport networks, traditionally isolated from each other from a management perspective. The 5G-Crosshaul project [1] integrates FH and BH through common control and data planes to enable a flexible, automated, and homogeneous handling of the interconnection of 5G radio access and core network functions. In turn, its open interface enables system-wide optimizations and network cost reductions. In order to do so, 5G-Crosshaul embraces software defined networking (SDN) and network functions virtualization (NFV) principles, since traditional network management approaches based on the configuration of proprietary devices are cumbersome, error-prone, and inefficient [2]. In fact, SDN principles allow unified and dynamic control, management, and configuration of integrated 5G multi-technology transport elements, while NFV allows infrastructure and function virtualization to enable flexible and cost-effective usage of the underlying physical infrastructure resources.

There are other approaches that also address 5G transport networks. For instance, the 5G-XHaul project [3] focuses on dynamic reconfigurability with a cognitive control plane for small cells and cloud radio access networks (RANs). On the other hand, our focus is the validation of the unified control and data plane for BH and FH.

In this article, we deploy a complete instance of the 5G-Crosshaul network in a testbed composed of three geographically distributed sites. The core of this system resides in the control plane, where a hierarchical approach, exploiting a developed application programming interface (API) providing common abstraction models, has been adopted to automate network resource orchestration in the deployed multi-domain multi-technology transport network. This contrasts with current management schemes of telco pro-

---

Jorge Baranda, Josep Mangues-Bafalluy, Iñaki Pascual, Juan Luis De la Cruz, Ramon Casellas, and Ricard Vilalta are with Centre Tecnològic de Telecomunicacions de Catalunya (CTTC/CERCA) – Castelldefels; Josep X. Salvat is with NEC Labs Europe; Charles Turyagyenda is with InterDigital Europe; José Núñez-Martínez is currently with Telefonica Product Innovation, but was with CTTC/CERCA when this work was done..

Digital Object Identifier:
10.1109/MCOM.2018.1701329

0163-6804/18/$25.00 © 2018 IEEE

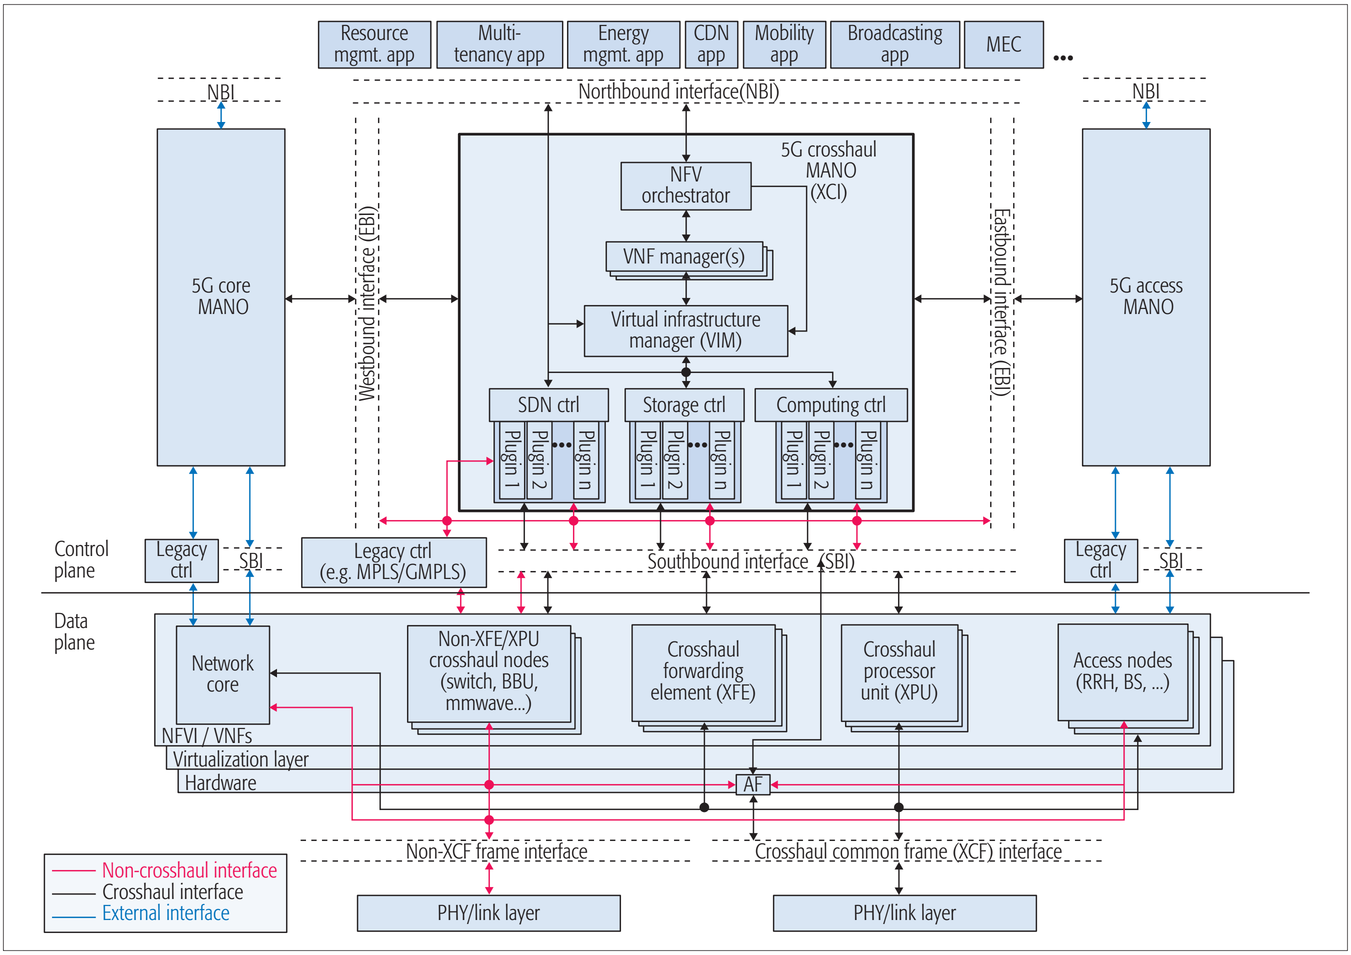

Figure 1. 5G-Crosshaul architecture.

vider transport networks, where the management is done manually and in isolation at each domain with very limited integration of SDN control, hence incurring long times (up to the order of days), mainly due to administrative overheads. Per-vendor/per-technology transport domains, each with its isolated control employing proprietary extensions, hinders the end-to-end (E2E) global management of transport networks.

The concept of hierarchical distributed control of SDN networks has been proposed in the context of data centers [4] (as opposed to wide area network environments considered herein), mainly to pursue a high degree of scalability. In [5, 6], preliminary design/architectural ideas behind the joint orchestration of wireless and optical transport domains and the initial characterization of part of this setup was presented. This article extends the complex and heterogeneous system of [6] with another wireless transport domain, which benefits from our scalable hierarchical design and characterizes its operation from a service management perspective. In particular, we present the assessment of the automatic setup and flexible recovery of a complete transport network deployed to fulfill the needs of a virtual LTE mobile network featuring FH and BH. The obtained average values for the system under test are 10.467 s for service setup, and 0.299 s or 6.652 s for local and centralized recovery, respectively. This confirms the validity of the 5G-Crosshaul architecture to orchestrate E2E network resources on the scale of seconds to meet the 5G target of substantially reducing the average service creation time.

This article is organized as follows. The next section provides an overview of the global 5G-Crosshaul architecture. Then we present our instantiation of this architecture in a multi-technology testbed. After that, we present the description and the quantitative experimental assessment of the deployed system for two case studies. The final section concludes the article.

5G-CROSSHAUL ARCHITECTURE OVERVIEW

The 5G-Crosshaul architecture proposed in [1] for the management of a mobile transport network is depicted in Fig. 1. The core and access segments are also depicted for completeness, although they are out of the scope of 5G-Crosshaul architecture. This design considers state-of-the-art NFV [7] and SDN [8] architectures to maximize the compatibility and integration of the system design with existing standard frameworks and reference specifications, as described in [1]. This consideration allows reusing open source projects, hence facilitating its deployment while reducing implementation costs. The main components of this design are below.

At the Data Plane: The 5G-Crosshaul forwarding element (XFE) supports FH and BH traffic profiles simultaneously. XFEs are switching units supporting single or multiple link technologies, including millimeter-wave (mmWave), WiFi, Ethernet, fiber, microwave, copper, and so on. A key part of the XFE is a common switching layer implementing a

The data plane combines several domains using different transport technologies. The mmWave/WiFi transport network (in Barcelona) and the mmWave transport network (in London) represent the wireless edge packet-switched domains of the transport network. All of the XFEs in these domains are equipped with wireless Gigabit based interfaces based on mmWave (IEEE 802.11ad) links.

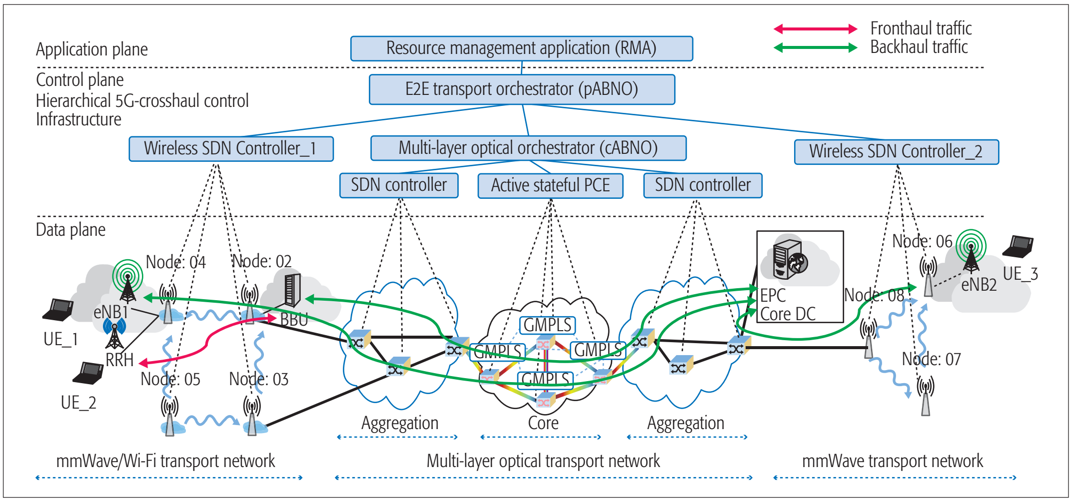

Figure 2. Experimental 5G-Crosshaul network. The arrows represent the type of traffic and the different flows required to establish the LTE mobile service described below.

layer 2 common frame encapsulation for enabling unified and harmonized transport traffic management. Non-XFE nodes may also be integrated. As for IT resources, 5G-Crosshaul processing units (XPUs) are in charge of hosting the various virtual network functions (VNFs) of the deployed services.

At the Control Plane: The 5G-Crosshaul control infrastructure (XCI) is the main element controlling the overall operation of the 5G-Crosshaul network. It is a management and orchestration (MANO) entity, which also supports the management of networking, storage, and processing resources through dedicated controllers. The XCI part follows the European Telecommunications Standards Institute (ETSI) NFV architecture [7].

At the Application Plane: Innovative network applications exploit 5G-Crosshaul resource orchestration functions to support the most diverse functionalities such as planning, network and service monitoring/prediction, optimization of resources, energy management, multi-tenancy, or media distribution.

BUILDING THE EXPERIMENTAL 5G-CROSSHAUL NETWORK

Figure 2 presents the architecture of the deployed system, in which we map the 5G-Crosshaul architecture presented in Fig. 1. In particular, the aim of the deployed system is to show the capabilities of a hierarchical deployment of the XCI to manage heterogeneous transport domains for the automatic orchestration of E2E network resources controlled from the application plane.

The system is composed of building blocks coming from three different geographically distributed sites throughout Europe: Barcelona, Heidelberg, and London, which are connected by dedicated virtual private network (VPN) tunnels. The following subsections present the different deployed components to build an instance of the 5G-Crosshaul architecture.

DATA PLANE

The data plane combines several domains using different transport technologies. The *mmWave/WiFi transport network* (in Barcelona) and the *mmWave transport network* (in London) represent the wireless edge packet-switched domains of the transport network. All of the XFEs in these domains are equipped with wireless-gigabit-based interfaces based on mmWave (IEEE 802.11ad) links. Moreover, some of them also come with WiFi (IEEE 802.11ac) interfaces, and hence support multiple wireless technologies. The *multi-layer optical transport* network (in Barcelona) represents the core transport domain, which features two layers separated into different sub-domains: first, two aggregation packet-switched Ethernet networks with tunable optical interfaces at edge nodes, and second, an all-optical dense wavelength-division multiplexing (DWDM) mesh network with two colorless reconfigurable optical add/drop multiplexer (ROADM) and two optical cross-connect (OXC) nodes, interconnected through five bidirectional optical links with a total of 610 km of optical fiber. Each optical node has multiple DWDM transceivers up to 2.5 Gb/s and one at 12.5 Gb/s with fully tunable laser sources.

CONTROL PLANE

As mentioned before, the core of the 5G-Crosshaul system resides in the XCI, and in this system, the SDN controller component has been deployed following a hierarchical approach of three layers for which a *parent-child* relationship is established between contiguous layers. The hierarchical approach not only enables orchestration in a multi-technology, -domain, -vendor environment typical of envisioned 5G network deployments, but also improves scalability, modularity, and security as explained in [8]. Within the transport SDN community, it is commonly accepted that a single integrated controller for a large or complex network may present several technical issues, or may not be doable in practice. Three main reasons are:

- Network size, in terms of controllable elements, which has a direct impact on the controller requirements
- Network complexity, in terms of having a network combining multiple technology layers
- Security, in terms of exposure to all network equipments toward the same control instance, which could easily be hijacked.

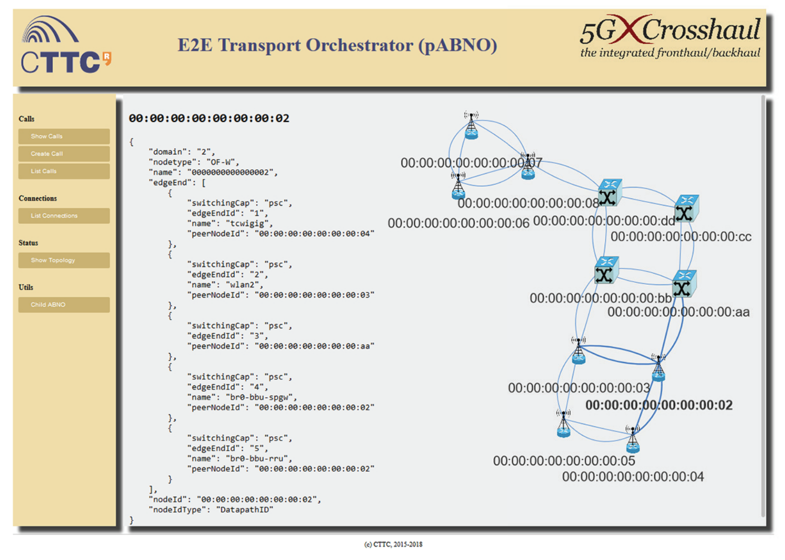

00:00:00:00:00:00:00:02

```
{
    "domain": "2",
    "nodetype": "OF-W",
    "name": "0000000000000002",
    "edgeEnd": [
        {
            "switchingCap": "psc",
            "edgeEndId": "1",
            "name": "tcwigig",
            "peerNodeId": "00:00:00:00:00:00:00:04"
        },
        {
            "switchingCap": "psc",
            "edgeEndId": "2",
            "name": "wlan2",
            "peerNodeId": "00:00:00:00:00:00:00:03"
        },
        {
            "switchingCap": "psc",
            "edgeEndId": "3",
            "peerNodeId": "00:00:00:00:00:00:00:aa"
        },
        {
            "switchingCap": "psc",
            "edgeEndId": "4",
            "name": "br0-bbu-spgw",
            "peerNodeId": "00:00:00:00:00:00:00:02"
        },
        {
            "switchingCap": "psc",
            "edgeEndId": "5",
            "name": "br0-bbu-rru",
            "peerNodeId": "00:00:00:00:00:00:00:02"
        }
    ],
    "nodeId": "00:00:00:00:00:00:00:02",
    "nodeIdType": "DatapathID"
}
```

Figure 3. GUI of E2E transport orchestrator (pABNO) presenting the topological vision of the parent controller based on the information provided by the different child controllers. Nodes in the figure correspond to those labeled in Fig. 2, except for the optical domain, abstracted by the cABNO.

These three main issues can be addressed by the introduction of a hierarchical architecture, as presented in this article.

At the top of the hierarchy, the *parent controller* orchestrates the different *child controllers*, which handle the specificities of the different equipment and link technologies at the underlying XFEs. In this setup, all the control plane entities depicted in Fig. 2 reside in the Barcelona site except *Wireless SDN Controller_2* of the *mmWave transport network*, which is placed in London.

In the deployed system, the parent controller is the *E2E Transport Orchestrator*, which is based on the Internet Engineering Task Force (IETF) application-based network operations (ABNO) architecture [9]. It is referred to as the parent ABNO (*pABNO*). The *pABNO* is in charge of composing the multi-domain topology and generating the corresponding connectivity requests from/toward the child controllers. The ABNO architecture supports hierarchical deployments with arbitrary depth, allowing recursive and scalable deployments thanks to the use of a unified southbound interface (SBI) and northbound interface (NBI). Hierarchical deployments imply no direct communication between neighboring controllers (i.e., those controllers that are responsible for adjacent network segments or those at lower or equal hierarchy levels). In this case, the interface used in this system is the Control Orchestration Protocol (COP),[1] which offers topology, connectivity, and path computation services. COP, initially developed in the context of the STRAUSS project,[2] is further extended in the 5G-Crosshaul project, notably to account for wireless technologies. COP precedes similar efforts later carried out at standards development organizations (SDOs) such as the Open Networking Foundation (ONF) Transport API (T-API).

COP defines an information model using the YANG modeling language [10], and it employs RESTconf [11] as the underlying protocol to exchange JSON objects between the different control entities. In addition to this, as a REST API does not allow notifications, websockets have been introduced (as described in RESTconf) in order to manage notifications about nodes, link status, and updates on the established paths.

An example of recursion in the deployed XCI of the experimental system under description is present in Fig. 2, where an ABNO-based controller acts as a *multi-layer optical orchestrator*, referred to as a *child ABNO* (*cABNO*). A recursive hierarchical deployment can help accommodate multiple underlying transport technologies, SDN controller types, geographical/administrative domains, or heterogeneous network segments. In the deployed system, the hierarchy serves multiple administrative and transport technology domains. The data plane components placed in Barcelona come from two testbeds located at the Technological Centre of Telecommunications of Catalonia (CTTC), namely EXTREME[3] and ADRENALINE.[4]

In turn, the *cABNO* also has child controllers. In fact, the *cABNO* provides E2E multi-layer (packet and optical) and multi-domain topology and network resource provisioning capabilities across its different child domains as *pABNO* does for the whole infrastructure. Two child SDN controllers serve the packet-switched domains represented at both sides of the all-optical network and a general multiprotocol label switched (GMPLS)/active stateful (AS) path computation element (PCE) manages the optical circuits of the core segment. Within the ABNO architecture [9], the *cABNO* also includes the abstraction manager (AM), which is able to provide several types of abstraction levels, such as node or link abstraction of the underlying network resources. Such abstractions serve to hide the unnecessary particularities of this domain to the *pABNO*, hence easing its job toward more modular and scalable deployments, as depicted in Fig. 3.

The *wireless SDN controllers* are the child controllers for their respective transport domains. As such, they control the edge wireless transport segments to attend the topology and connectivity requests coming from the *pABNO*. In addition, in this setup, these controllers also come with path computation capabilities for path recovery purposes, as explained below. At its SBI toward the XFEs, they use OpenFlow (OF) protocol for configuration of forwarding behavior. Furthermore, the child controller of the *mmWave/WiFi* domain also offers a RESTconf interface for configuration and management of its wireless interfaces. This interface has also been developed within the context of the 5G-Crosshaul project as a way to provide an abstract network information model (control parameters and system status metrics) of the wireless data plane technology. Thanks to these kinds of abstractions, network management applications (application plane) can exploit 5G-Crosshaul resources to support diverse functionalities.

Additionally, the elements of the control plane present a graphical user interface (GUI), where useful information about the network state under its vision can be extracted. For instance, the *pABNO* GUI, in Fig. 3, can show the abstracted representation of the underlying transport network and the set of installed service calls accord-

---

[1] https://github.com/5G-Crosshaul/COP, accessed 20 April, 2018

[2] www.ict-strauss.eu, accessed 20 April, 2018

[3] http://networks.cttc.cat/mobile-networks/extreme_testbed/, accessed 20 April, 2018

[4] http://networks.cttc.cat/ons/adrenaline, accessed 20 April, 2018

The Resource Management Application is the application that resides on top of the system and is in charge of managing E2E network resources. Physically, it resides in Heidelberg and makes use of the services offered by the hierarchical control infrastructure by means of the previously described COP protocol via a VPN tunnel.

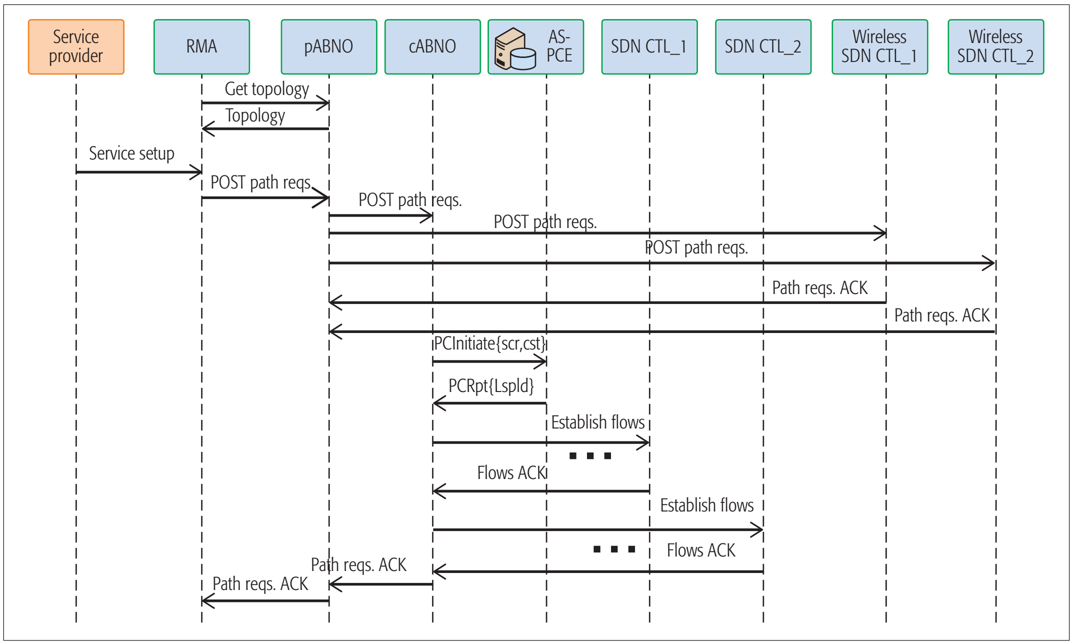

Figure 4. Message exchange workflow for provisioning connectivity services.

ing to the model defined in the COP API. The *Wireless SDN Controller_1* GUI shows the configuration of the wireless interfaces (also allowing its reconfiguration), the *mmWave/WiFi* transport topology and a dynamic representation of the relative load of the wireless transport nodes and its associated links.

Application Plane

The resource management application (RMA) is the application that resides on top of the system and is in charge of managing E2E network resources. Physically, it resides in Heidelberg and makes use of the services offered by the hierarchical control infrastructure by means of the previously described COP protocol via a VPN tunnel. When connecting with the *pABNO*, the RMA retrieves the current topology from the *pABNO*, which is updated constantly thanks to the websocket notification system. Upon reception of a service request, the RMA computes optimal paths for the request, where source, destination, and traffic profiles (bandwidth, delay) are defined. Then the RMA generates the appropriate JSON objects communicating to the *pABNO* the computed paths satisfying the requirements of the different traffic profiles received in the request. When the service is no longer required, the RMA orders the *pABNO* to delete all the forwarding rules at the XFEs associated with this service. The RMA also has a GUI that shows the abstracted view of the transport network and the currently installed paths in the system.

5G-Crosshaul Network Assessment Description

In [6], we find the characterization from the control plane perspective of the different transport domains placed in Barcelona, which use different transport technologies, namely optical and mmWave/WiFi. The aim of this analysis was to validate the individual transport domains while measuring the time contribution of each domain at the time of setting up a multi-domain E2E path with the hierarchical control infrastructure. Here, we further contribute to the previous characterization of the deployed system from a service perspective level, including another wireless transport domain (both control and data plane) deployed at another physical site, showing the scalability and flexibility of the built system. In particular, the case studies used to characterize the deployed system are the experienced service setup time and the service recovery time under two strategies: a local recovery procedure performed by a child SDN controller and a centralized recovery procedure performed at the RMA.

Different control plane designs — especially if centralized — need to account for heterogeneous recovery requirements, and it is well known that a pure centralized solution may not be optimal in all cases (e.g., in terms of recovery delay). Hybrid approaches (i.e., that favor local decisions where possible and escalating to a higher level based on predefined policies and rules or when recovery is not possible) are commonly conceived to design for the shortcomings. In any case, there are important trade-offs to be considered, and an operator policy should decide when to apply local or centralized recovery. In addition to this, the type of failure can impose a specific type of recovery: an inter-domain failure implies corrective actions in multiple domains, which cannot be handled locally, or multi-layer/multi-technology connection constraints may limit applicability of local recovery (colored optical interfaces, asymmetrical ROADM nodes). Macroscopically, local recovery may be optimal in terms of minimizing traffic disruption time, but the lack of topology visibility associated with multi-domain networks and the fact that local recovery is commonly constrained to maintain the original (local) endpoints preclude end-to-end optimality in terms of resource usage or the fulfillment of traffic requirements, such as delay. On the other hand, centralized recovery may potentially be optimal, but at the expense of higher restoration times, control bottlenecks, and scalability issues (local recovery is to some extent distributed); and finally, topology abstraction of hierarchical deployments may also limit optimality.

Herein, the aim of assessing different recovery strategies is to illustrate the flexibility of the presented hierarchical system to allow both recovery approaches.

The deployed (and then recovered) service is an LTE mobile network layer based on the OpenEPC platform [12] composed of several virtual machines (VMs) deployed as endpoints of the transport network. Two enodeBs and a remote radio head (RRH)-baseband unit (BBU) pair featuring fronthaul packet data convergence protocol/radio link control (PDCP/RLC) split are placed at the edge wireless transport domains, as depicted in Fig. 2. All of them connect to the Evolved Packet Core (EPC) entities placed at the *multi-layer optical transport network*, namely the mobility management entity (MME) and serving/packet data network (PDN) gateway (SPGW). Hence, the RMA will determine the appropriate paths among the different mobile network entities for the FH and BH traffic profiles, represented by the different arrows in Fig. 2.

In order to establish the LTE mobile network service, eight paths (a total of 16 flows to provide bidirectionality to the different paths) are required: from RRU to BBU, from BBU to MME, from BBU to SPGW, from eNodeB1 to MME, from eNodeB1 to SPGW, from eNodeB2 to MME, from eNodeB2 to SPGW, and from SPGW to MME, as shown in Fig. 2.

Figure 4 shows the workflow followed to provision the described LTE mobile network service. Based on the received service request and the abstracted view of the network topology provided by the *pABNO*, the RMA generates the set of path provisioning requests to provide the desired connectivity between mobile network layer entities satisfying the requirements, in terms of throughput/delay, of the different traffic profiles. The *pABNO* processes these path requests, decomposing them into different requests for each network domain, which are then sent to the underlying child controllers. Child controllers process these path requests and enforce the required forwarding rules in its corresponding network elements under control. A chain of acknowledgment answers validates the correct orchestration of the requested resources.

5G-CROSSHAUL NETWORK EXPERIMENTAL RESULTS

This section presents a quantitative evaluation of the case studies presented earlier.

CASE STUDY I: SERVICE SETUP TIME

We define the service setup time as the interval between the computation of required paths at the RMA and the reception of the confirmation from the pABNO that all the requested paths have been installed at the different transport domains. Figure 5 presents the histogram and the cumulative distribution function (seen from the RMA) of the time to set up the described LTE mobile network service 100 times.

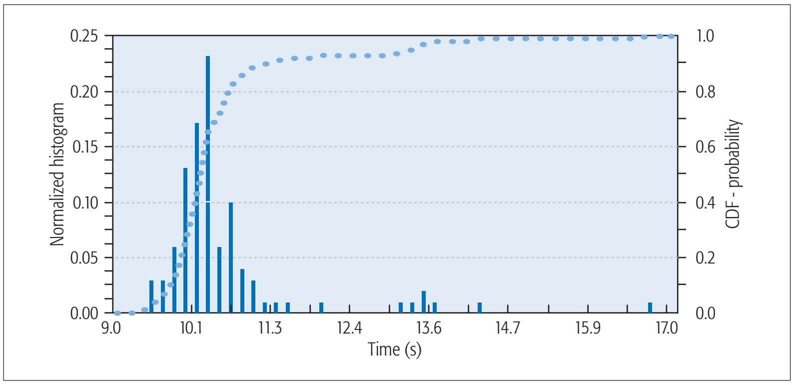

Figure 5. Normalized histogram and CDF of the service setup time seen by the RMA.

As shown in Fig. 5 , the time it takes for the RMA to calculate and validate the installation of the whole set of path requests to interconnect the different deployed mobile network entities is lower than 11 s in 90 percent of the samples, with an average value of 10.467 s and a minimum value of 9.343 s. Hence, the achieved system verifies that the built 5G-Crosshaul network meets the 5G KPI target of enabling the introduction/provisioning of new network services on the order of magnitude of minutes/hours (even seconds in this case), improving current management schemes as stated previously.

From a deeper analysis of the measured experimental time values at each layer of the control hierarchy, the most time consuming operations in terms of setup time are those requests involving setting up a bidirectional multi-domain path traversing the multi-layer optical network. This is due to the amount of actions that have to be done in the set of network elements hidden under the abstraction provided by the cABNO controller to the RMA. In particular, the first of such path requests is always the one that requires more time, around 2.9 s. As described in [6], this obeys the need to tune laser interfaces at the nodes labeled as GMPLS in Fig. 2 to set up the initial lightpath and the Ethernet service on top. After this first path is established, all the remaining path requests traversing the multi-layer optical network reuse the same lightpath and are established much faster, around 1 s.

CASE STUDY II: SERVICE RECOVERY TIME

Next, we evaluate the system's configurable approaches to perform service recovery in the event of a link failure. In both cases, the link failure happens in the direct link connecting Node:04-Node:02, depicted in Fig. 2. In order to re-establish the mobile network service, three paths (six unidirectional flows) have to be re-established: from enodeB1 to MME, from enodeB1 to SPGW, and from RRU to BBU.

Local Service Recovery: *Wireless SDN Controller_1* starts the recovery process upon the detection of an OF PORT_STATUS message. In parallel, it sends an update topology notification up the control hierarchy (which will eventually arrive to the application plane, RMA) to report the topology update and that the recovery process will be done locally. Since the child SDN controller maintains a database with all the installed flows,

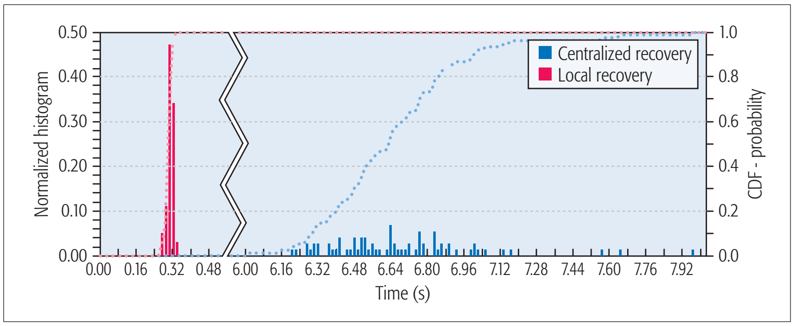

Figure 6. Normalized histogram and CDF of local and centralized service recovery procedures.

and based on the current "domain" network state, it checks which flows have been affected and proceeds to remove affected paths following a "break-before-make" strategy and to calculate an alternative path, if possible, to install it in the underlying XFEs. In this case, the alternative path is Node:04-Node:05-Node:03-Node:02 to maintain the original endpoints and circumscribe the recovery within a single domain. If the path is not possible, the child SDN controller sends a notification up the hierarchy to inform that the paths are definitively removed from the system. Once *Wireless SDN Controller_1* has received confirmation that the alternative flows have been installed in the corresponding XFEs, it notifies the RMA through the *pABNO* of the alternative paths set up in the network. Then the RMA can update its own database to have the current network vision to be able to provide the appropriate paths for subsequent service requests. In Fig. 6, we can see the histogram and the CDF of performing 100 operations of local recovery of the deployed LTE mobile network service.

The measured recovery time values, with an average of 299 ms, are similar to those reported in [13] when using OF Fast Failover (FF) group tables and using Bidirectional Forwarding Detection (BFD) protocol for the same environment and even when repairing several flows. By using the BFD plus OF-FF approach, a link down event can be detected, and re-routing decisions are triggered locally without the need to involve any SDN controller (if backup paths have been configured in advance). However, in this case, the BFD plus OF-FF would not be a good solution since no control plane entity is involved in the recovery process; hence, RMA or another control plane entity would not have a consistent or updated view of the network in terms of topology and installed paths. In addition to this, the OF-FF failover solution can potentially present additional drawbacks. A short BFD monitoring interval could lead to increased traffic and processing overhead or triggering false link transitions, which could affect the link performance for data transmission [14]. In order to provide backup paths, more forwarding rules have to be installed in the nodes (the more rules, the denser the transport network deployment). Finally, the requirement of installing additional rules to perform crankback forwarding to send packets back toward a transport node with an alternative active path toward the destination impacts the experienced latency [15].

Centralized Service Recovery: In this case, after the link failure event, *Wireless SDN Controller_1* reports up the hierarchy the change in the topology and that no local recovery is going to be performed. The RMA receives this notification and proceeds to detect the affected flows to send the corrective actions back to the *pABNO*. First, it deletes those previous requests with affected paths. Then it computes the new paths for the affected ones and re-installs them if they continue fulfilling the specified requirements at the service request; otherwise, the path is not valid, and the service is not recovered. Figure 6 shows the histogram and the CDF of performing 100 operations of centralized recovery at the RMA.

The measured times present more variance, increasing to the order of seconds (average value of 6.652 s) compared to the local recovery case evaluated previously. This is due to several factors. First, climbing up and down in the hierarchy introduces delay, which is also increased due to the VPN connections. Second, the processing time at the RMA to update the topology and check the affected rules will be higher than that in the lower control level entities. Third, and most relevant, are the corrective actions the RMA orders. Due to its global view, these corrective actions involve deleting and creating actions in multiple domains; namely, *mmWave/WiFi* and *multi-layer optical transport*. In the case of local recovery, corrective actions only involve the *mmWave/WiFi* transport network. As pointed out previously, actions performed in the multi-layer optical transport domain (for both creating and deleting) may potentially require more time due to the bigger amount of involved network elements at both the control and data planes. In spite of the presented centralized recovery time values, we validate the flexibility of our control plane infrastructure to have different levels of reaction in front of network events.

SUMMARY, CONCLUSIONS, AND FUTURE WORK

The ambitious KPIs pursued by upcoming 5G networks will require an integrated management of diverse network segments (access, transport, and core) enabling increased flexibility and dynamicity. In the case of transport networks, the 5G-Crosshaul project [1] proposes the integration of FH and BH through common control and data planes to interconnect distributed access functions with core network functions embracing the use of SDN and NFV principles. This article validates the flexibility, scalability, and recovery capabilities of the 5G-Crosshaul architecture in a real testbed distributed among three geographically distributed sites. Particular focus is put on the proposed hierarchical control plane. Compliance of the 5G-Crosshaul network with 5G service deployment and operation requirements is evaluated through two case studies. In the first one, we show that end-to-end orchestration of network resources across a multi-domain multi-technology transport network can be done in seconds. More specifically, we attain service deployment times around 10.5 s for provisioning all the networks paths required by a complete virtualized LTE mobile network layer featuring FH and BH configurations. In this way, this work contributes to achieving the envisioned 5G KPI target of reducing service deployment time from days to minutes. In the second one, we experimentally show that the 5G-Crosshaul control infrastructure (XCI) offers the required recovery capabilities in the event of link failures. In particular, it is shown how the automatic service recovery procedure can be done at the controller closer to the physical infrastructure, hence providing faster reaction times at the cost of potentially suboptimal paths. Alternatively, it can be centralized at the network management application, which results in higher delays but potentially optimal global paths.

Following up on this research, we intend to evolve the XCI with slicing capabilities to enable 5G telco infrastructure sharing between different vertical industries. Additionally, work on federation will enable the further development of the unified network orchestration concept and virtualization of heterogeneous 5G mobile transport networks.

Following up on this research, we intend to evolve the XCI with slicing capabilities to enable 5G telco infrastructure sharing between different vertical industries. Additionally, work on federation will enable the further development of the unified network orchestration concept and virtualization of heterogeneous 5G mobile transport networks.

ACKNOWLEDGMENTS

This work was supported by MINECO grants TEC2017-88373-R (5G-REFINE) and TEC2015-69256-R (DESTELLO), Generalitat de Catalunya grant 2017 SGR 1195, and the 5G-Crosshaul project (H2020-671598).

REFERENCES

[1] X. Costa-Prez *et al.*, "5G-Crosshaul: An SDN/NFV Integrated Fronthaul/Backhaul Transport Network Architecture," *IEEE Wireless Commun.*, Feb 2017, pp. 38-45.

[2] W. Xia *et al.*, "A Survey on Software-Defined Networking," *IEEE Commun. Surveys & Tutorials*, vol. 17, 2015, pp. 27–51.

[3] Horizon 2020, 5GPPP: 5G Xhaul Project; https://5g-ppp.eu/5g-xhaul , accessed 20 Apr. 2018.

[4] L. Fang *et al.*, "Hierarchical SDN for the Hyper-Scale, Hyper-Elastic Data Center and Cloud," *Proc. ACM SIGCOMM Symp. Software Defined Networking Research*, Santa Clara, CA, 2015.

[5] R. Vilalta *et al*, "Hierarchical SDN Orchestration of Wireless and Optical Networks with E2E Provisioning and Recovery for Future 5G Networks," *Proc. OFC*, Anaheim, CA, 2016.

[6] J. Mangues et al., "Experimental Evaluation of Hierarchical Control over Multi-Domain Wireless/Optical Networks," *Proc. EUCNC Euro. Conf. Networks and Commun.*, Oulu, Finland, 2017.

[7] ETSI NFV, "Network Functions Virtualisation (NFV); Architectural Framework," ETSI GS NFV 002 v1.1.1, Oct. 2013.

[8] ONF, SDN Architecture 1.0 (TR-502), 2014.

[9] D. King and A. Farrell, "A PCE-Based Architecture for Application-Based Network Operations," IETF RFC 7491, Mar. 2015.

[10] M. Björklund, "YANG — A Data Modeling Language for the Network Configuration Protocol (NETCONF)," IETF RFC 6020, Oct. 2010.

[11] A. Bierman *et al.*, "Restconf Protocol," IETF RFC 8040, Jan. 2017.

[12] The OpenEPC project, Core Network Dynamics GmbH; http://www.openepc.com/, accessed 20 Apr. 2018.

[13] R. Santos and A. Kassler, "Small Cell Wireless Backaul Reconfiguration Using Software-Defined Networking," *2017 Proc. IEEE Wireless Commun. Networking Conf.*, San Francisco, CA, 2017.

[14] J. Vestin and A. Kassler, "Low Frequency Assist for mmWave Backhaul — The Case for SDN Resiliency Mechanisms," *Proc. 2017 Wksps. IEEE ICC*, Paris, France, 2017.

[15] N. L. M., van Adrichem, B. J. Van Asten, and F. A. Kuipers, "Fast Recovery in Software-Defined Networks," *Proc. 2014 3rd Euro. Wksp. Software Defined Networks*, Budapest, Hungary, 2014.

BIOGRAPHIES

Jorge Baranda (jorge.baranda@cttc.es) is currently a researcher in the Mobile Networks Department at CTTC, Barcelona, Spain. He received an M.Sc. (2008) degree in telecommunications engineering from the Technical University of Catalonia (UPC). At CTTC, he participates in national, European, and industrial projects related to the development of novel wireless communication systems and 5G mobile networks. His current research interests include wireless communications, wireless backhaul, network optimization, software-defined networking, and network functions virtualization.

Josep Mangues-Bafalluy (josep.mangues@cttc.cat) is a senior researcher and head of the communication networks division (http://networks.cttc.cat) of CTTC, where he participates in EU, industrial, Spanish, and Catalan research projects in various roles, including seven as Principal Investigator. He has published over 100 international journal and conference publications and was General Vice-Chair of IEEE WCNC '18. He received his M.Sc. (1996) and Ph.D. (2003) in telecommunications from UPC. His research interests include software-defined networking and network functions virtualization in self-organized mobile networks.

Iñaki Pascual (inaki.pascual@cttc.es) holds a degree (2014) in informatics engineering from UPC. He has collaborated with the research group N3CAT (UPC) and the IN3 research group of the Open University of Catalonia. In 2012, he moved to Sillicon Valley to work at PLUMgrid, an SDN company adquired by VMWare in 2016. In 2015, he joined the CTTC where he conducts research in SDN and NFV applied to mobile networks.

José Núñez-Martínez (jose.nunezmartinez@telefonica.com) is currently a senior technology expert within Telefonica Product Innovation. He received his M.Sc. (2005) in computer science and his Ph.D. (2014) in telecommunications, both from UPC. From 2004 to 2007, he worked as a network engineer in the Advanced Broadband Communication Center (CCABA) of UPC. From 2007 to 2017, he worked as a researcher in CTTC, participating in national, European, and industrial projects. He has published more than 40 research papers.

Juan Luis De la Cruz (juan.luis.delacruz@cttc.es) received his B.Sc. (2016) in telecommunications engineering from the University of Zaragoza (EINA-UNIZAR). He was awarded by the COGITT as the third best final degree thesis nationwide in the telematics area. He joined CTTC in January 2017 as a research assistant, where he collaborates in projects related to software-defined networking and network functions virtualization in data centers and optical transport networks.

Ramon Casellas (ramon.casellas@cttc.es) has been with CTTC since 2006. Before, he worked as an associate professor at ENST France, obtaining his Ph.D. in 2002. At CTTC, he has worked in European, national, and industry research projects, on topics related to traffic engineering, network optimization, and network control and management of optical and multi-layer transport networks. His research interests include the GMPLS/PCE architectures and protocols, and SDN/NFV. He has coauthored over 150 papers, 4 book chapters, and 4 IETF RFCs.

Ricard Vilalta (ricard.vilalta@cttc.cat) has a telecommunications engineering degree (2007) and Ph.D. degree (2013) from UPC, Spain. He is a senior researcher at CTTC in the Communication Networks Division. His research is focused on SDN/NFV, network virtualization, and network orchestration. He has been involved in international, EU, national, and industrial research projects, and published more than 170 journal and conference papers and invited talks. He is also involved in standardization activities in ONF, IETF, and ETSI.

Josep Xavier Salvat (josep.Xavier.Salvat@neclab.eu) received his M.Sc (2016) from UPC. He currently works as a research associate in NEC Laboratories Europe. Since 2017, he is pursuing a Ph.D. degree at Kaiserslautern University, Germany. His research interests are computer networks, 5G transport networks, and the design and prototyping of SDN/NFV solutions.

Charles Turyagyenda (charles.Turyagyenda@InterDigital.com) holds a Ph.D. in telecommunications from the University of Sheffield, United Kingdom. He currently works at InterDigital International Labs, London. Throughout his career, he has been active in the research, development, and standardization of communication network technologies, namely multi-access edge computing, software-defined networking, and network functions virtualization. He is an inventor, developer, and scientific writer who has invented multiple patents and authored several peer-reviewed publications.