

A Service-based model for the Hybrid Software Defined Wireless Mesh Backhaul of Small Cells

José Núñez-Martínez Jorge Baranda Josep Mangues-Bafalluy

Centre Tecnològic de Telecomunicacions de Catalunya (CTTC)
Av. Carl Friedrich Gauss, 7 08860 Castelldefels (Barcelona), Spain
e-mail: [jose. nunez, jorge.baranda, josep.mangues]@cttc.cat

Abstract—The backhaul of upcoming dense 5G Small Cell (SC) networks needs a full redesign, but there is no clear consensus on how to deploy such infrastructure. A low-cost backhaul solution is that provided by a wireless mesh network. Additionally, Software Defined Networking (SDN) is being considered as an alternative to distributed approaches to lower network infrastructure costs while enabling programmability and flexibility mostly for wired networks. This paper evolves the canonical SDN model by presenting a service-based hybrid SDN (hSDN) model that alleviates the problems caused by the unreliability of the in-band control channel formed by a wireless mesh backhaul between the SDN controller and the SCs. At the infrastructure level, we propose a wireless mesh backhaul combining sub-6GHz and millimeter wave links with long range microwave links. This architecture pursues the coexistence of network services located on top of a centralized SDN controller, with distributed network services, such as routing. To show the robustness of our proposed model, we compare a service-based hSDN model and a canonical SDN model under perfect control plane communication channel. In our service-based hSDN model, we further compare two distributed routing schemes used as fallback control plane mechanisms when the SDN controller is unreachable. Simulation results with ns-3 show improvements of up to 1.5x and 6x in terms of throughput and latency, respectively.

I. INTRODUCTION

The mobile data growth is expected to continue for emerging 5G networks. Given the limited spectrum availability, reducing the cell radii by deploying extremely dense Small Cell (SC) networks, which we will refer to as DenseNet, represent and effective way to increase capacity. Being unlikely that fiber reaches every SC site, the backhaul of emerging DenseNet deployments will have to be redesigned to attain the required capacity at a low cost. A low-cost backhaul solution for DenseNets is to wirelessly inter-connect all its SCs, thus forming a wireless mesh backhaul as in [1]. On the other hand, the evolution of commodity network infrastructure offers the opportunity to evolve the management of these networks to lower costs in totally new ways. A key player enabling this transformation is Software Defined Networking (SDN) [2]. SDN can improve the management and control of new and existing network services¹, allowing their evolution and programmability to be agnostic from the underlying infrastructure. However, SDN technologies have been hitherto constrained to the presence of a wired infrastructure, such as the wired data center [3]. The adoption of SDN for the wireless backhaul of DenseNets, currently under discussion ([4,5]), raises major concerns. The SDN paradigm implies the centralization of all the control plane, which requires of a protocol to control the

forwarding plane of network elements (in our case SCs), such as OpenFlow (OF) [6]. Thus, the control plane relies on a high speed network communication channel from the SDN controller to the SCs to enforce reliability and low response delay. A wireless mesh backhaul challenges the existence of a reliable in-band network communication channel between SC devices and the SDN controller, thus, compromising the efficiency of a centralized SDN model. The centralization of critical network services requiring a quick response could often generate network failures, especially for delay-sensitive network services, such as the procedures in charge of updating/modifying the routing table of SCs. Interestingly, there is a wide range of robust distributed wireless routing protocols designed to cope with the nature of wireless mesh backhails, such as Backpressure for DenseNets (BDN) [1].

The main contribution of this paper is to sum the benefits of centralizing and distributing network services by presenting a service-based hybrid SDN [7] (hSDN) architecture to manage the backhaul of DenseNets. At the infrastructure layer (see Section II), we present a two-level wireless mesh backhaul infrastructure, in which sub-6GHz and millimeter wave technologies cooperate with microwave links to aggregate capacity at a low cost. As for the control layer, we present a model based on centralizing delay-tolerant services and distributing critical non-delay tolerant services. As for the application layer, we focus on the placement of routing and energy efficiency network services in the resulting architecture. Section III evaluates the hSDN architecture versus a fully centralized SDN architecture using ns-3 simulations [8]. Results show the better adaptability of the service-based hSDN architecture to dynamic backhaul reconfigurations compared to the centralized SDN, showing improvements of up to 6x in terms of latency. Besides, within a service-based hSDN architecture we compare two distributed control plane: BDN [1] and GPSR [9], a State of the Art (SoA) routing protocol for wireless networks, as fallback distributed routing mechanisms when the SDN controller is down. As fallback mechanism, BDN showed improvements of up to 6x and 1.5x in terms of latency and throughput compared to GPSR. Section IV concludes the paper.

II. THE DENSENET ARCHITECTURE

In this section, we first describe the infrastructure proposed to form a wireless mesh backhaul for DenseNets. Second, we provide an overview of the proposed service-based hybrid SDN (hSDN) architecture. Third, we show how some canonical network services fit into the proposed architecture.

¹In this paper, the terms application and service are used interchangeably.

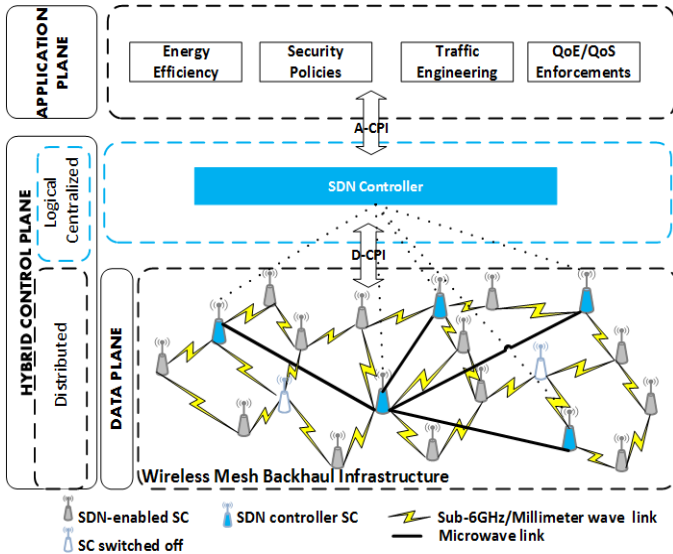


Fig. 1. The proposed service-based Hybrid SDN (hSDN) architecture.

A. The DenseNet Infrastructure

There is no clear consensus on the backhaul technologies of choice for DenseNets, since fiber may be unsuitable in an economic manner for extremely dense SC deployments. Due to installation expenses (e.g., digging the trench) fiber is a more expensive CAPEX than equipping each SC with several wireless backhaul radios. As Figure 1 illustrates, we propose a two-level wireless backhaul infrastructure: the first level consists of sub-6GHz and millimeter wave radios, whereas the second one is based on long-range microwave links.

The first infrastructure level, present in all SCs, uses Non-line-of-sight (NLOS) low-cost technology to decrease the planning cost invested for deploying each SC while increasing the resiliency level by means of Point to Multipoint (PTMP) wireless links. To tackle peak rate demands between short distances, this level also includes high capacity millimeter wave links requiring LOS, such as 60GHz 802.11ad WiGig links. These millimeter wave links can be deployed between SCs located in lampposts, traffic lights, at high heights where LOS can be feasible. The second infrastructure level consists of long range wireless links to communicate network elements located at distant points. We propose to locate some SCs above roof height on top of towers and endow them with high-gain directional antennas featuring a Point to Point (PTP) LOS microwave links. Due to the longer range of microwave links (see Figure 1), compared to sub-6GHz and millimeter wave wireless links, this complementary technology offers high-capacity low-latency shortcuts between distant points.

B. Control Architecture View

In the envisioned infrastructure, network services (e.g., routing, energy efficiency, QoS/QoE, and security policies) in the application plane (see Figure 1) need to be deployed at a low cost, while attaining the required flexibility. Though the deployment of a centralized SDN model simplifies the hardware of SCs, its adoption in the presented wireless mesh backhaul presents a major issue: *the unreliability of the*

wireless medium together with the assumed low-cost deployment compromises the communication channel amongst SDN-enabled SCs and the SDN controller. The data-controller plane interface (D-CPI or southbound interface in Figure 1) includes among others functions, the programmatic control of the rules (e.g., forwarding rules) installed in SC devices [10]. The in-band control channel amongst SCs and the SDN controller is unlikely to provide a quick and reliable response when frequently updating the Forwarding Information Base (FIB) of SC devices, which could be needed for certain critical delay-sensitive network services, such as routing. Such an in-band control channel can cause data plane disruptions due to, for instance, the additional delays (or losses) experienced by OpenFlow (OF) control messages under congestion.

We propose a service-based hSDN model [7] that combines the pure centralized SDN approach with a distributed control logic approach based on the robustness of existent distributed protocols. The main idea is to centralize delay-tolerant and specific network services, while distributing those delay-sensitive network services. The infrastructure's unreliability especially affects delay-sensitive network services (e.g., those based on providing network connectivity and routing amongst network nodes). Indeed, these network services have traditionally been provided by spreading the control plane amongst the different network elements, operating efficiently in a distributed way. However, the application of a centralized SDN model seems to cope with fine-grained and complex emerging backhaul network services generating network actuation with a milder time scale, such as network services aiming to alleviate energy consumption. In this way, the SDN controller would act as the brain for all those delay-tolerant network services commanding the set of rules to apply to the rest of the SCs. In turn, all the SDN-enabled SCs (gray SCs in Figure 1) would receive orders from the SDN controller, but also would embed distributed control logic to perform delay-sensitive operations based on a scalable distributed protocol such as BDN [1].

There are also two more issues addressed by our proposed architecture. First, *a single SDN controller constitutes a point of failure and presents scalability problems for big DenseNet deployments.* A way to counteract this is to build a logically centralized SDN controller as in [11], hence spreading the SDN controller amongst different physically separated network elements. Second, *next generation deployments should guarantee the interoperability with already deployed DenseNet backhauls.* The transition to a full SDN control architecture will imply the coexistence of legacy SCs with SDN-enabled SCs. We consider that legacy SC devices, with the time, could support SDN protocols, such as OpenFlow (OF) [6], via a software upgrade [7]. Thus, all SCs would be SDN-enabled (gray SCs in Figure 1).

C. Network Services

This subsection focuses on two network services with different requirements: routing and energy efficiency.

Routing: We decompose the routing service into two main components. *The first component is a centralized control logic* that receives the different routing requests, and may apply different traffic engineering policies. To implement routing policies, the SDN application integrates topology discovery and path computation modules to provide non-technology

dependent fined-grain routing policies through the Application-Control Plane Interface (A-CPI in Figure 1). However, the infrastructure herein considered lacks from a fast and reliable communication network for the control plane between SDN-enabled SCs and the SDN controller. For this reason, *the second component involves a distributed control logic* that will accommodate routing and load balancing while the communication with the SDN controller is lost, or when SCs experience congestion due to low communication bandwidth between the SCs and the SDN controller.

Energy Efficiency: The regulation of sleep cycles of SCs is taken periodically at long timescales (i.e., in the order of tenths of minutes or hours). Given also that global network information is required, energy efficiency will run centralized as an SDN application. The SDN controller will send an abstraction of the relevant network information (i.e., topology information, wireless link type, and traffic patterns) through the Northbound interface (A-CPI interface in Figure 1) periodically, or under the request of the energy efficiency application. Based on this information, the SDN energy efficiency application will take network actions based on switching on/off SCs, hence, reconfiguring the SC wireless backhaul.

III. DENSENET: PERFORMANCE EVALUATION

We use ns-3 simulations [8] to model a 7x7 grid wireless mesh backhaul of hybrid SCs. In particular, it embeds a single SDN controller in the SC located in the center of the grid. The SC embedding the SDN controller includes not only routing functionalities but an energy efficiency network service. Hybrid SCs also embed a distributed control plane running BDN [1] and a SoA routing for distributed wireless networks, namely GPSR [9]. All the SC devices endow a 54Mbps PTMP WiFi interface and 20% of the SC devices include an additional 311Mbps PTP microwave interface.

A. SDN Controller Up: Centralized vs Hybrid Control Plane

Here we compare a fully centralized and a hybrid control plane. The fully centralized architecture counts with an SDN routing application based on distance to the destination in terms of number of hops, whereas the service-based hybrid architecture embeds the BDN distributed routing component in each SC. Common to both architectures, there is a centralized energy efficiency SDN application oriented to achieve energy costs savings by dynamically reconfiguring the wireless backhaul topology. This scenario reveals the achieved data plane performance under ideal control plane conditions for the centralized model, since the control plane channel experiences no losses and null latency between SCs and the SDN controller.

We simulate fluctuating traffic patterns by modifying the number and spatial distribution of the set of injected 2Mbps UDP CBR flows. In the centralized approach, the SC encapsulates the first packet of a flow with no matching entry in its flow table in an OF *OFPT_PACKET_IN* request and forwards it to the SDN controller. The controller, then, inserts the proper matching rules in SCs by means of OF *OFPT_FLOW_MOD* messages forming the shortest path between source and destination SCs. Based on OF *OFPT_PACKET_IN* requests, the energy efficient service can abstract the traffic patterns and dynamically reconfigure the wireless backhaul topology every 40 minutes (i.e., non-delay sensitive network actuation). This reconfiguration is done by communicating over the A-CPI

interface to the SDN controller, which in turn communicates with SC devices sending the proper OF *OFPT_FLOW_MOD* message [2]. We repeat the aforementioned process 10 times for a total simulation time of 400 minutes. Whereas the centralized approach determines fixed paths of the same length in terms of number of hops requiring global network topology information, BDN exploits local information to determine in each SC the next-hop on a per-packet basis.

Figure 2 compares the data plane performance of both approaches in terms of throughput and latency. BDN presents remarkable latency improvements experiencing up to 6x of improvement for 15 traffic flows even though it merely uses local information. This is because of its per-packet load-balancing capabilities to deal with time-varying network conditions without entailing communication with the SDN controller. With the centralized approach, latency increases approx 2x for 12 and 15 traffic flows due to queuing experienced by used SCs for routing traffic. The problem with the centralized approach is that, due to its static nature, the trajectory followed by packets of a same traffic flow is equivalent. Even though there is a perfect communication channel between SCs and the controller, BDN obtains slight throughput gains due to its better usage of scarce wireless resources.

B. SDN Controller Down: Recovery

In this scenario, the connection to the SDN controller is lost during a significant amount of time (i.e., 40 minutes). We compare the performance of two distributed routing schemes that could act as fallback mechanisms for resiliency and availability purposes, namely BDN [1] and GPSR [9]. The SCs and the SDN controller periodically exchange keep-alive OF *OFPT_HELLO* messages. We consider that SCs detect that the communication with the SDN controller is lost after losing the reception of two consecutive OF *OFPT_HELLO* messages. At this stage, SCs delete all forwarding entry rules issued by the SDN controller and changes its operation mode by activating the distributed routing service (BDN or GPSR) in all SCs. We consider that upon the establishment of a new handshake between the SCs and a SDN controller (i.e., OF handshake initiated by the SDN controller) the communication with the SDN controller is reestablished and the SCs can therefore change its operation mode to receive OF forwarding rules from the SDN controller. In the simulations, we inject different sets of thirty fluctuating traffic flows where we vary the source-destination pairs of each traffic flow every minute. We repeat the previous experiments but increasing the rate of the thirty injected traffic flows from 0.5Mbps up to 2Mbps, ranging the workload from 15Mbps up to 60Mbps. Furthermore, we have evaluated different network topologies, i.e., 10 different sets of SCs equipped with a microwave backhaul link are tested, to check the independence of these fallback distributed mechanisms with the network topology.

Figure 3 shows significant data plane performance improvements by using BDN in the hSDN model while the SDN controller is down, especially under high loads. In particular, a hSDN with BDN outperforms a hSDN with GPSR by up to 1.5x and 6x for an offered load of 60Mbps in terms of throughput and latency, respectively. The jump experienced by latency curves for 30 Mbps, 45Mbps, and 60Mbps case is due to the inefficiency of the distributed component GPSR to exploit non congested data plane paths (e.g., those enabled by

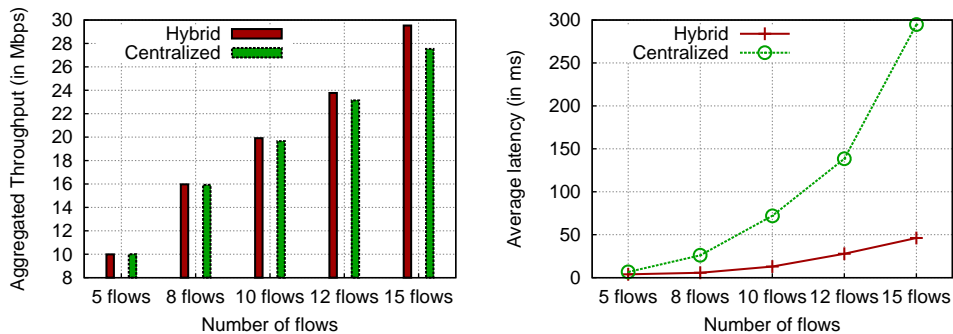


Fig. 2. Data plane performance of centralized SDN-based architecture versus a service-based hybrid SDN architecture using BDN.

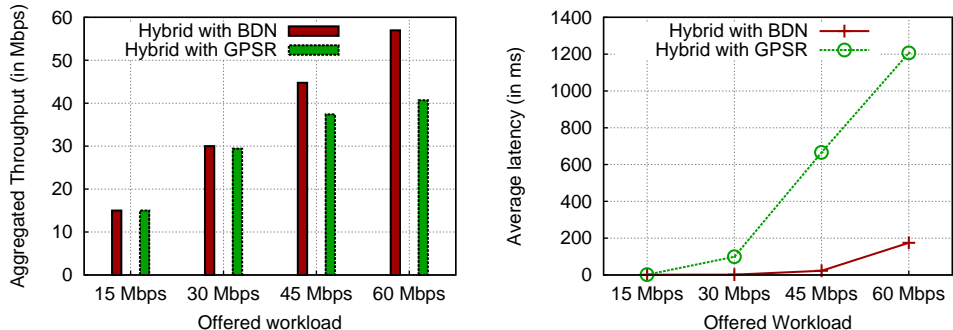


Fig. 3. Data plane performance of service-based hSDN model under SDN controller failure during a 40 minute period.

SCs with microwave links) to reach their intended destinations. While GPSR restricts the backhaul resources among those that geographically steer packets closer to the destination, BDN dynamically increases/reduces the set of backhaul resources, and so the number of non-congested paths, to route traffic flows according to traffic demands. Under congestion, hSDN with BDN is able to bias the shortest path trajectory by choosing less loaded wireless backhaul resources, whereas hSDN with GPSR is unaware of the load, hence suffering from throughput and latency degradation. These results confirm the benefit of a service-based hSDN architecture under the interruption of the communication channel between SCs and the controller. The particular benefit comes from the fact that the flow table of SCs can be controlled by a distributed routing protocol.

IV. CONCLUSIONS

This paper presents our service-based hSDN model to manage a wireless mesh backhaul that shares the flexibility and programmability given by the canonical SDN centralized approach and the robustness of distributed control planes, hence offering the proper compromise between centralization and distribution of network services. We showed how our service-based hSDN model has much higher performance with respect to a canonical SDN model, showing improvements of up to 6x in terms of latency. These results confirm the suitability of our service-based hSDN model for the management of dense wireless mesh backhauls. We observed how our hSDN model can obtain good performance results even when the SDN controller is down, especially with scalable distributed protocols that can perform adaptive load balancing such as BDN, showing improvements of up to 1.5x and 6x with respect to GPSR in terms of throughput and latency, respectively. We expect this paper could shed some light on the role of SDN in low-cost, unreliable, and dynamic wireless mesh backhaul

deployments. Thanks to this work, we have first performance figures of a service-based hSDN model managing a dense wireless mesh backhaul.

ACKNOWLEDGEMENTS

This work was carried out in part within Xhaul project (H2020-ICT-671598), and by the Spanish Ministry of Economy and Competitiveness under grant TEC2014-60491-R.

REFERENCES

- [1] J. Núñez-Martínez, J. Baranda, and J. Manges-Bafalluy, "A self-organized backpressure routing scheme for dynamic small cell deployments," *Ad Hoc Networks, Elsevier*, 2015.
- [2] Open Network Foundation, "Software defined networking: The new norm of networks," *White Paper*, 2012.
- [3] S. Jain, A. Kumar, S. Mandal, J. Ong, L. Poutievski, A. Singh, S. Venkata, J. Wanderer, J. Zhou, M. Zhu *et al.*, "B4: Experience with a globally-deployed software defined wan," in *ACM SIGCOMM Computer Communication Review*. ACM, 2013.
- [4] IEEE P802.16r, "View of connection-oriented software-defined networking for wireless backhaul of small cells, July 2013," July 2013.
- [5] Open Network Foundation, Wireless & Mobile Working Group.
- [6] Open Network Foundation, "Openflow switch specification v.1.4.0," *White Paper*, 2013.
- [7] S. Vissichio, L. Vanbever, and O. Bonaventure, "Opportunities and research challenges of hybrid software defined networks," *ACM SIGCOMM Computer Communication Review*, 2014.
- [8] "The ns-3 network simulator," available at: <http://www.nsam.org>.
- [9] B. Karp and H.-T. Kung, "GPSR: Greedy perimeter stateless routing for wireless networks," in *Proc. of the 6th annual Int. Conf. on Mobile computing and networking*. ACM, 2000.
- [10] Open Network Foundation, "SDN Architecture," June 2014, available at <http://www.opennetworking.org>.
- [11] A. A. Dixit, F. Hao, S. Mujherjee, T. Lakshman, and R. R. Kompella, "Towards an elastic distributed SDN controller," in *Procs of the 2nd ACM SIGCOMM Workshop on Hot Topics in Software Defined Networking (HotSDN 2013)*. ACM, 2013.