Integrating WMN Based Mobile Backhaul with SDN Control

Kari Seppänen(✉), Jorma Kilpi, and Tapio Suihko

VTT Technical Research Centre of Finland, Vuorimiehentie 3, Espoo, Finland
{kari.seppanen,jorma.kilpi,tapio.suihko}@vtt.fi

Abstract. Resilient high capacity and low delay millimeter wave wireless mesh networks (WMNs) can provide suitable backhaul connections for future broadband mobile networks. The WMN solution is best suited in cases where base stations are installed in locations without optical fiber connection to transport network, e.g., small-cell deployment to hot spots in dense urban areas. Recently software defined network (SDN) concept has become popular in many networking areas including mobile networks. One of the key promises is to provide an efficient way for network operators to extend and create new services. As the whole network is controlled by a single central entity that is based on software code, it would be easy to make large scale network upgrades without need to wait that updates are available for all network elements (NEs). There is, however, a clear conflict between SDN ideas and WMN operation. The performance and reliability of the latter one is heavily depended on fast local reactions to, e.g., link degradations. Centralized control would introduce longer delays in reactions. In this paper, we are proposing a concept which solves these problems and allows for combining the best features of both WMN and SDN.

Keywords: Wireless mesh · SDN · Mobile backhaul · Network abstraction · Network virtualization

1 Introduction

Increasing capacity demands in broadband mobile networks call for new solutions especially in dense urban areas. To meet these needs, the traditional macro-cell architecture has to be augmented with small cells that cover the hot-spot areas. One of the major problems in small-cell deployment is the lack of suitable fixed wireline backhaul connections in many potential installation sites. Furthermore, this development can easily lead to multiplication of the number of cells by a factor of 10. This means that the cost of small-cell installation must be brought down as low as possible with, e.g., zero-configuration [1].

Installing new fiber optic network connections for micro cells is very costly in dense urban areas and, in many cases, also very time consuming because of the required planning and official permits. Thus, the use of wireless backhaul connections is a natural choice for these kind of small-cell scenarios. However,

R. Giaffreda et al. (Eds.): IoT360 2014, Part II, LNICST 151, pp. 222–233, 2015.
DOI: 10.1007/978-3-319-19743-2_32

the capacity requirements for LTE-A and forthcoming 5G are such (≥1 Gbit/s per base station) that they are hard to meet with current wireless systems. Millimeter wave (mmW) RF systems (e.g., 60 GHz or 71–88 GHz) can provide ample capacity to meet the requirements.

Using mmW RF technology makes it necessary to apply narrow-beam directed point-to-point links between stations to provide sufficient link budgets for usable link spans. This is actually an advantage as it increases the total system capacity compared to omni-directional transmissions. However, narrow mmW beams are rather vulnerable to disturbances and thus the reliability of an end station with only one point-to-point link could be quite low. The solution for this problem is to have mesh connectivity between end stations (from now on WMN nodes) and gateways. Moreover, using WMN for backhaul connectivity allows for having reliable multihop paths between micro cells and WMN gateways and thus extending the area that can be covered with a single WMN.

Software-defined networking (SDN) consists of techniques that facilitate the provisioning of network services in a deterministic, dynamic, and scalable manner. SDN currently refers to approaches of networking in which the control plane is decoupled from the forwarding functions and assigned to a logically centralized controller. The SDN architecture, with its software programmability, provides agile and automated network configuration and traffic management that is vendor neutral and based on open standards. Network operators are able to dynamically adjust the network's traffic flows to meet the changing needs while optimizing the network resource usage. An OpenFlow-based SDN is formed by switches that forward data packets and communicate with one or more controllers using the OpenFlow protocol. An OpenFlow controller configures the forwarding behavior of the switches by setting packet processing rules in their flow tables. A rule is composed of match criteria and actions. The match criteria are multi-layer traffic classifiers that inspect packet headers and identify the set of packets to which the actions will be applied. The actions may involve modification of the packet and forwarding through a defined output port, for example.

In this paper, we are proposing a concept that will integrate an existing mmW WMN backhaul solution with SDN-based centralized transport network control. The main goals of our concept are to resolve the conflict between local and centralized control as well as to provide "plug-and-play" style incremental network extension. Furthermore, the capability of sharing the network resources among multiple mobile network operators (MNOs) is a very important target.

The structure of the paper is following: in Sect. 2 we provide background information about using WMN to provide backhaul connectivity and about related work in SDN; in the next Sect. 3 we present our solution for controlling a WMN-based backhaul with SDN; in Sect. 4 we discuss about the potential problems we have identified this far; and, finally, Sect. 5 ends this paper with some conclusions.

2 Background

First in this section, we explain some of the most critical requirements that are imposed on WMN by broadband mobile backhaul (MBH) and then we describe the main characteristics of our WMN backhaul concept. Then we will go through some related work about SDN and WMN as well as SDN and mobile networks. Finally, we discuss about the trade-offs between centralized versus local control.

2.1 WMN Based Small-Cell Backhaul

Small-cell backhaul should be seen as a part of the whole mobile network infrastructure and WMN portion as a last mile segment of the backhaul connection [1]. Thus, events in WMN, like failures, can affect the rest of the network by, e.g., triggering handovers between base stations and fixed network side protection switching. This means that, to get best advantages from alternative backup paths that WMN provides, the fault recovery mechanisms at that level should operate, in the most of the cases, faster than "normal" telecom grade protection (50 ms).

There are already some concepts that integrate mmW radio links in backhaul and SDN ideas, e.g., hybrid wireless optical MBH described in [2]. In our concept, we are utilizing a novel mmW WMN MBH system that has been developed in various earlier projects [3,4]. This WMN concept is not limited to repeater (or relay) configurations but it supports meshed multihop paths allowing better coverage and more alternative routes. As a single WMN can be fairly large, multiple gateways to fixed network are also supported.

Packet forwarding in the proposed WMN concept is done at flow level. Flows are identified by inspecting L2 and/or L3 headers, e.g., Ethernet MAC addresses together with VLAN Id and PCP (Priority Code Point), or IP addresses and DSCP (Differentiated Services Code Point) field. Each WMN flow can be assigned to a separate path and these assignments can change dynamically based on network state. In case of congestion or link failure, high priority traffic can get better service while the best effort traffic suffers most of the damage. It is also possible to split one traffic flow to multiple paths as long as the paths end at the same gateway. Due to related processing overheads, this is usually applied only for "fat" non-realtime traffic flows.

2.2 Related Work

The applicability of SDN in carrier networks has been analyzed in [5]. The identified key challenges are performance vs. flexibility, scalability, security, and interoperability (backward compatibility) with existing networking technologies. Our work aims at finding solutions for these challenges in a software-defined WMN-based mobile backhaul.

Especially, in WMNs that are formed of commodity devices, node isolation and network fragmentation may occur frequently, which makes the application of centralized control problematic. To exploit the benefits of SDN while mitigating the drawbacks, [6] proposes to use the combination of a distributed routing protocol (OLSR) and OpenFlow in a Wireless Mesh Software Defined Network (wmSDN). In this solution, OpenFlow is used for balancing traffic load among Internet gateways of the mesh. On the other hand, our WMN system applies autonomous routing, load balancing, fault recovery, and other traffic management functions. Still, the WMN requires coordinated topology management and resource allocation. To that aim, we present the WMN to the SDN controller through an aggregating network abstraction, provided by a "mediator" function that implements a Hardware Abstraction Layer (HAL). The mediator translates between the SDN and WMN's operation models.

In general, a HAL is required to "hybridize" mixed networks, which contain pure legacy and SDN devices, in order to hide the idiosyncrasies of legacy network equipment such that on the outside the equipment looks like one or more SDN switches. The HAL concept [7] proposed by EU FP7 ALIEN project encompasses three types of integration model for different use cases depending on the tightness of the coupling between HAL and the legacy equipment and the multiplicity of devices that the HLA covers. One of the HAL implementations is xDPd (eXtensible OpenFlow DataPath daemon) developed within the ALIEN project.

When the WMN is considered as a legacy network segment partition in a mixed SDN network, the WMN is abstracted as a virtual SDN switch, among the native SDN switches. This is in contrast to the technology migration approaches, like HybNET [8], in which a legacy network is abstracted as a virtual link between the SDN switches.

2.3 Local vs Centralized Control

As explained earlier in this section, current and future mobile backhaul requirements are such that it is of paramount importance to hide all WMN impairments as perfectly as possible. In the best case, the WMN portion of the backhaul connection would be seen as a reliable bit-pipe – with somewhat elastic capacity. To achieve this kind of performance, it is necessary that fault management mechanisms inside WMN react to failures and other events much faster than other fault management mechanisms in the network. In practice, this means just few 10 s of ms time scales. As the delays inside WMN could be something like 1 ms per hop, the only viable way to achieve the required reaction speeds is to use local protection and recovery mechanisms.

One of the common main ideas of SDN is the optimization of the whole network configuration as the network state is (in principle) known by a single centralized control entity. This would make the usage of network resources more efficient and, at the same time, make it possible to utilize simple and cheap network equipment.

The main problem with WMN based backhaul and centralized control is that WMN is potentially very dynamic environment. Moreover, optimizing its operation requires lots of quite specific information from each link as well as detailed system specific understanding. Thus, centralized control for WMN would require transferring a considerable amount of status information and configuration commands between WMN nodes and centralized control entity. Furthermore, the centralized controller would have to, in practice, replicate the WMN control plane to provide the necessary functionality.

All this would cause extra traffic in the network and additional delays to fault protection operations. As the same functions can be handled locally, this is hard to justify. However, centralized control is very attractive alternative for configuring and controlling end-to-end traffic flows and backhaul connections. Thus, it would be quite beneficial if the local and centralized control could be made to live together by utilizing the best features from both.

3 SDN Configurable WMN

Our solution to the problem of integrating WMN MBH with SDN control is to leave the most of the WMN functions as they are in our WMN concept and to use SDN only to configure end-to-end connections. However, SDN controller cannot be allowed to configure routing inside the WMN as that would mess up the fault recovery, load-balancing and other WMN operations — and, *vice versa*, any self-configuration action taken by WMN would confuse the SDN controller. Thus, a key part of the solution is to hide the WMN internal structure and operations from SDN layer.

The WMN backhaul solution proposed in this paper will be a part of larger backhaul system that includes also fixed legacy and SDN transport network portions and covers the whole backhaul connection from base stations to MNO's mobile core network (e.g., EPC). One of the main ideas in this backhaul system is to provide virtualized network slices to multiple operators. While the virtualization concept *per se* is out of scope of this paper, the support for virtualization is included in some of the WMN backhaul features. Most clearly this can be seen in network infrastructure extension procedures (see Sect. 3.2).

3.1 WMN Abstraction

Network (or topology) abstraction is a powerful tool that allows construction of hierarchies in SDN [9]. The main idea is the same as with abstractions in programming in general: to give the programmer an access to information he needs while hiding all the internals from accidental manipulations. This abstraction principle is exactly what we need in hiding WMN operations from SDN controller.

The key elements of our abstraction model are that the whole WMN domain is represented as a single virtual SDN switch and that each WMN node port connected to a micro cell is shown as a separate (virtual) port in that switch (see Fig. 1). This effectively hides all the WMN functionality from the upper layers while it, at the same time, offers full control to configure all traffic flows from and to the micro cell. This abstraction model can also be used to hide the

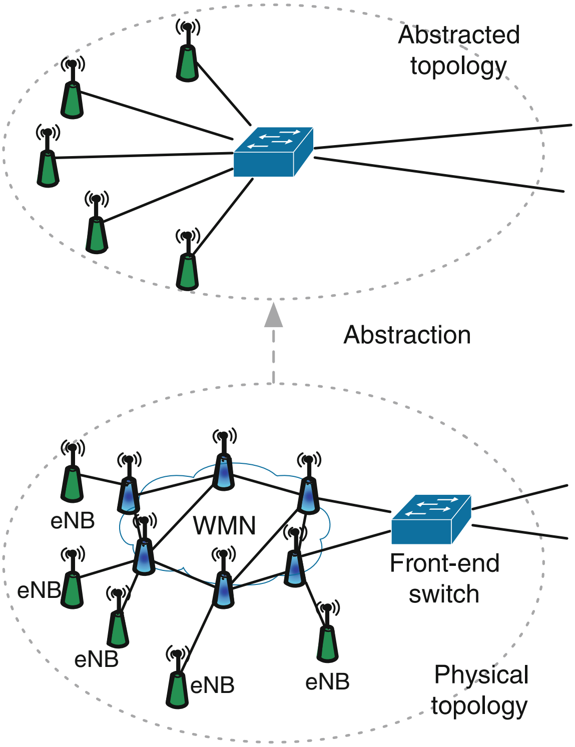

Fig. 1. Hiding physical WMN topology using network abstraction. The SDN controller layer sees only a single switch and base stations seem to be attached directly to it.

existence of multiple WMN gateways (GW) in a single domain and, which is quite useful, hide all such protection mechanisms inside WMN that could cause moving traffic flows from one gateway to another.

Hiding the existence of multiple GWs and traffic flow rerouting from one GW to another requires some extra functionality between WMN and fixed transport network. In our concept, we are using WMN Front-end Switch (WFS) to handle this functionality, i.e., switching the WMN traffic flows so that they appear to originate from a single virtual port. At the control plane, WMN Mediator Function (WMF) takes the responsibility of providing the SDN control interface and interpreting SDN commands given to WFS and translating them into WMN's control operations (see Fig. 2). In practice, the WMN abstraction is done by WMF.

In our WMN system, the traffic flows are, in practice, tunneled between WMN nodes and GWs, and the paths that these tunnels take are changing dynamically (down to ms scale). Even the target GW for a flow can change if necessary. In earlier configurations with legacy networks, the GWs terminated these tunnels and only customer payload was forwarded to fixed network. However, in this abstraction model, we have to hide some peculiarities of WMN from

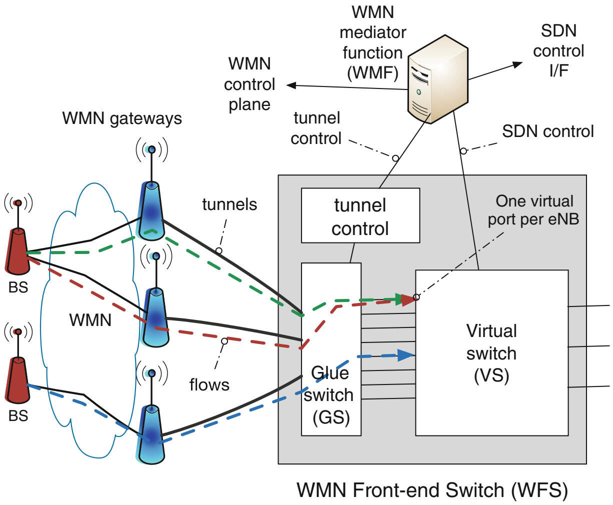

Fig. 2. WMN Front-end switch structure.

the SDN control plane, e.g., flows moved from a GW to another and flows from one base station passing through different GWs. For this purpose, we use additional tunneling between GWs and WFS. The main purpose of this tunneling is to carry information about flow identification (source WMN node) to WFS.

The changes in traffic flow routing over WMN can be detected in two ways: the WMN control plane sends a notification about path reselection to WMF or WFS detects that a traffic flow has moved from one GW-WFS tunnel to another. The latter case can be handled in OpenFlow (OF) like manner: WFS sends the "unknown" packet to WMF that makes the WFS reconfiguration after inspecting the packet header (in this case, tunnel specific header). In our WMN system, downstream and upstream traffic can have separate paths and even via separate GWs. However, in this case, we are forcing each flow to use the same GW in both directions. This allows that also downstream GW change can be triggered by WMN node. Thus, when the change of the GW for the upstream traffic is detected, WFS is reconfigured also to reroute downstream to the same GW. The GWs automatically adapt to this change and the paths for upstream and downstream traffic between GW and WMN node can still be different.

We will also support direct base station to base station connectivity (e.g., for LTE X2 traffic). The idea is, that when the mediator identifies a SDN command that tries to configure a connection between two WMN side ports in WFS, it asks WMN control plane to configure a direct connection inside WMN. Thus, the traffic can take the shortest route instead of being hauled over WFS. However, this causes also some problems with port statistics: it is not sufficient to just return the counters from WFS but we have to merge these values with intra-WMN traffic counters.

Our current design is based on two virtual switch instances inside WFS (as shown in Fig. 2). The Glue switch (GS) is taking care of routing WMN flows between GWs and WFS virtual ports. These virtual ports are, in fact, WMN side ports of the second switch instance, Virtual switch (VS). While it would be possible to create the same functionality using only one switch, this two-switch approach has some clear advantages. In practice, we can use some existing virtual switch as VS, e.g., Open virtual switch (OVS) or Indigo virtual switch (IVS). In such case, WMF can pass most of the OF commands directly to VS and VS has all the required functionality to provide (in this configuration) the abstracted network view for SDN controllers. Furthermore, all effects of the changes in traffic flows in WMN side are limited to GS.

3.2 Incremental Network Infrastructure Extension

In our network extension scenario, a new micro cell is installed to a hot-spot location and after power switch-on, the new micro cell should be brought into active state automatically. The new micro-cell base station can be connected to an existing WMN node or the WMN node can be installed at the same time (as a separate co-located unit or integrated to the cellular base station). In the latter case, WMN self-configuration procedures will first initialize the WMN node, which can then provide transport network connectivity for the micro-cell base station. In any case, the network infrastructure extension should not need any human interaction besides the actual physical installation procedure and powering up the new equipment.

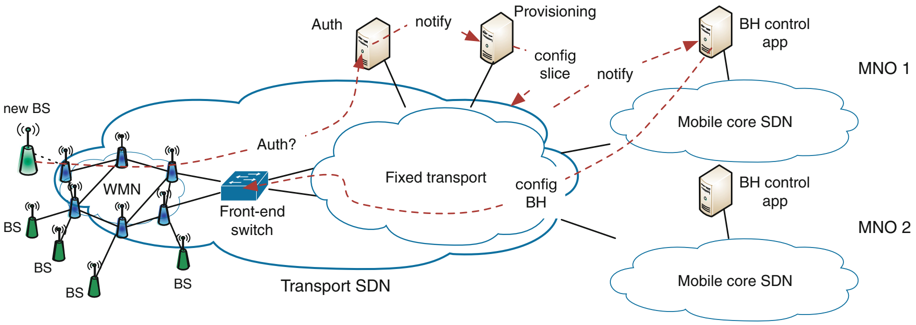

Fig. 3. Infrastructure extension use case.

During the WMN node configuration, the WFS is also configured to facilitate network extension. WMN control plane notifies WMF about the new WMN node and its configuration. Using this information, WMF adds new virtual ports to WFS virtual switch and configures each port so that all micro-cell authentication related messages are forwarded to authentication function ("Auth" in Fig. 3). All other traffic can be dropped by default until further configuration.

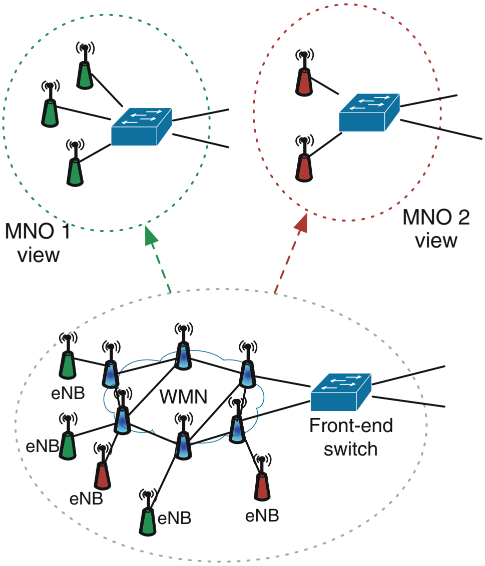

Fig. 4. Virtualized and abstracted network views for MNOs.

When the new micro-cell base station is installed and switched on, it should first try to authenticate itself and get some basic network configuration information. At first phase, authentication packets are received by Transport Network Operator's (TNO's) "Auth" service that identifies the owner of the new base station. If the identification and authorization is successful, TNO's provisioning element will reconfigure network virtualization so that the virtual port, to which the new base station is attached, will be added to the network view of correct MNO. As a result, each MNO should have its own virtualized view to the abstracted WMN (see Fig. 4).

When MNO's network controller is notified about the new port and thus about the new base station, MNO can continue with its own authentication procedures. When authentication has been passed, MNO can activate, e.g., backhaul control application that configures connections between the new base station and mobile core network.

In certain configurations, there can be a site switch owned by MNO and connected to WMN node port. If the site switch is fully transparent, there should be no specific new requirements for the network extension procedure — in MNO's network view, there can be just more than one base station attached to each port. However, if we want to make it possible for MNO to control the site switch as a SDN switch, things get a little bit more complicated. One possibility is that site switch can authenticate itself in the same manner as base stations. Yet this could be quite unrealistic scenario as it is likely that MNOs want to use low cost off-the-shelf equipment. An alternative approach is that one of base stations can detect and configure the site switch. In this case, the site switch should start its operation in transparent mode and change to SDN mode only after MNO's control applications have made the required configurations.

At the moment, we are planning to use FlowVisor (FV) for network virtualization. It should provide sufficient functionality for WMN virtualization but it is still unclear if it can meet the needs of the whole backhaul virtualization.

4 Discussion

There are some issues about reliability in the current WMN abstraction scheme. Only one WFS between WMN and fixed transport network is a clear single point of failure. It is true that failure rates at fixed network side hardware are much lower compared to, e.g., WMN link failures – especially if high-availability equipment is used. However, this situation is not satisfactory if, above all, WMN is used as a part of mobile backhaul network. This problem cannot be solved just by adding a second WFS in parallel as it would break down the WMN abstraction. One possibility is to mimic some kind of MC-LAG (Multi-Chassis Link Aggregation Group) functionality (similar constructions are already used in current network edge realizations) and hide that functionality inside WMN abstraction.

Network virtualization in SDN is quite commonly understood simply as just slicing physical resources (e.g., OpenFlow switches) to provide somewhat isolated network slices for multiple SDN controllers. In this simple virtualization model, all controllers can see the actual physical topology of their network slice. In our concept, the network abstraction hides the actual physical topology and, in this sense, it is not directly compatible with the current "controller should know everything" models. Our ideas of network virtualization are closer to, e.g., ITU-T Y.3011 model of Logically Independent Network Partitions (LINPs) and we are studying if we can steer the development of our backhaul concept more towards that direction [10].

When the WMN is shown as a virtual switch to the MNO, each switchable connection (which is actually a path or even a collection of paths) has some resiliency metric (path availability) due to mesh connectivity and some path E2E delay metric(s) due to predictable link scheduling. Given the routing and the scheduling of the WMN, these metrics have some computable optimal targets and measurable realizations.

One potential problem with our WMN abstraction model is that there is no clear capacity concept inside the WMN. The link throughputs are not the only varying factor as the dynamic path selection is also changing path allocations. Thus, the capacity one flow "sees" might fluctuate all the time. Furthermore, the flows traveling between base station and virtual port seen by SDN control can have separate paths and thus they do not share the same fate. The first problem is with SDN flow configuration: if the MNO wants do capacity reservations, then what capacity value should be given to each virtual port. The second problem is with capacity fluctuations: if MNO tries to optimize SDN flow routing by monitoring flows, it is necessary to quickly identify which impairments are due to WMN (can't do anything) and which ones caused by rest of the MBH.

When two or more MNOs have slice of the same WMN the capacity problem has another dimension: temporarily the capacity of the WMN can be lower than the sum of the slices sold to the MNOs. There should be some business oriented but fair approach to diminish the available capacity of all MNOs.

A possible solution is to define the marketable capacity of the WMN as a function of the (total) available capacity of the gateway(s') links. This means that capacity fluctuations of the gateway links only are taken into account. Capacity fluctuations of all other links are ignored. This information would be readily available at the gateway(s), without signalling delays. As far as the mesh topology assumption holds this approximative approach is actually quite justified, but in real life there will also be non-mesh topologies, e.g., tail sites and ignoring the capacity fluctuations of these unprotected tail links is less justified.

This WMN capacity and network abstraction problem is one thing that we have to study further in the later development phases. At the moment, it is quite difficult to advance with the analysis of these challenges without any hands on experimenting.

5 Conclusions

In this paper, we presented our solution to the problem of how to integrate WMN with SDN in mobile backhaul context. Instead of trying to figure out how to control WMN directly with a centralized SDN controller, we are proposing that network abstraction should be used to hide the whole WMN from the SDN layer. This abstraction means that the SDN controller sees the whole WMN as a single SDN switch that can be controlled normally as any other switch. A key component in our concept is the WMN Front-end switch that allows us to hide also the GW changes from the controller. Another important feature besides network abstraction is the support for incremental network infrastructure extension that is made possible by a combination WMN self-configuration and WMN node client authentication. This feature supports also network virtualization and thus the MBH infrastructure can be shared with multiple MNOs.

We are currently working to demonstrate these functions in our WMN testbed. Moreover, we are continuing to study the open items that are discussed about in the previous section. Therefore, we are looking forward to be able to present the real hands on results from the testbed as well as a more complete concept.

Acknowledgments. This work has been performed in the framework of the CELTIC-Plus project C2012/2-5 SIGMONA. The organisations on the source list would like to acknowledge the contributions of their colleagues to the project, although the views expressed in this contribution are those of the authors and do not necessarily represent the project.

References

1. NGMN Alliance: Small cell backhaul requirements (2012)
2. Bojic, D., Sasaki, E., Cvijetic, N., Wang, T., Kuno, J., Lessmann, J., Schmid, S., Ishii, H., Nakamura, S.: Advanced wireless and optical technologies for small-cell mobile backhaul with dynamic software-defined management. IEEE Commun. Mag. **51**(9), 86–93 (2013)
3. Taipale, T.: Feasibility of wireless mesh for LTE-advanced small cell access backhaul. Master's thesis, Aalto University School of Electrical Engineering (2012). http://lib.tkk.fi/Dipl/2012/urn100686.pdf
4. Wainio, P., Taipale, T.: Wireless mesh access backhaul for small cell base stations. In: Costa-Requena, J. (ed.) Innovative Solutions for Mobile Backhaul, CELTIC/CP7-011 MEVICO (2012). http://www.mevico.org/D32.pdf
5. Sezer, S., Scott-Hayward, S., Chouhan, P.K., Fraser, B., Lake, D., Finnegan, J., Viljoen, N., Miller, M., Rao, N.: Are we ready for SDN? Implementation challenges for software-defined networks. IEEE Commun. Mag. **51**(7), 36–43 (2013)
6. Detti, A., Pisa, C., Salsano, S., Blefari-Melazzi, N.: Wireless mesh software defined networks (wmSDN). In: 2013 IEEE 9th International Conference on Wireless and Mobile Computing, Networking and Communications (WiMob), pp. 89–95, October 2013
7. Report on hardware abstraction models, Deliverable D2.1, EU FP7 Project ALIEN, January 2013. http://www.fp7-alien.eu/files/deliverables/D2.1-ALIEN-updated.pdf
8. Lu, H., Arora, N., Zhang, H., Lumezanu, C., Rhee, J., Jiang, G.: Hybnet: network manager for a hybrid network infrastructure. In: Proceedings of the Industrial Track of the 13th ACM/IFIP/USENIX International Middleware Conference, Middleware Industry 2013, pp. 6:1–6:6. ACM, New York (2013)
9. Monsanto, C., Reich, J., Foster, N., Rexford, J., Walker, D.: Composing software-defined networks. In: Proceedings of NSDI 2013: 10th USENIX Symposium on Networked Systems Design and Implementation (2013)
10. Recommendation ITU-T Y.3011 Framework of network virtualization for future networks (2012)